



Textual Assemblages and Transmission:

Unified models for (Digital) Scholarly Editions and
Text Digitisation

Richard Hadden M.A.

Submitted for the qualification of PhD

Maynooth University

September 2018

Supervisor: Professor Susan Schreibman, An Foras Feasa

Head of Department: Professor Susan Schreibman, An Foras Feasa

Table of Contents

Table of Contents.....	3
Abstract	7
Acknowledgements.....	9
Introduction.....	11
1. Literature Review.....	29
1.1. Introduction	29
1.1.1. Basic terminology: document, work, text	35
1.2. Scholarly Editing Traditions	39
1.2.1. Critical editing.....	39
1.2.2. Anglo-American Editing.....	41
1.2.3. ‘Critique...’	45
1.2.4. Documentary editing.....	51
1.2.5. French and German perspectives	55
1.3. Digital Scholarly Editions	62
1.3.1. The digital turn.....	62
1.3.2. Scale and heightened representation of the digital	69
1.3.3. The digital encoding of texts	72
1.3.4. Digital documentary editing.....	76
1.3.5. Digital critical editing.....	82
1.3.6. Digitisation for other purposes	87
1.4. Conclusion	93
2. Textual Transmission and Validity	99
2.1. Introduction	99
2.2. Transmission of textual material by reproduction	108
2.3. Assemblages, intensities and extensities	114
2.4. Towards stability: the creation of extensities.....	121
2.5. Decoding and deterritorialisation	128
2.6. Recoding and reterritorialisation in the output.....	136
2.7. The User–Output Interaction	142
2.8. Constraining Interpretation	148
2.9. Regimes of signs.....	157
2.10. Conclusion	160
3. Transmission models and parameters.....	165
3.1. Introduction	165
3.2. A transmission model	170

3.2.1.	Decomposing the transmission process.....	170
3.2.2.	Linear and recursive connections.....	174
3.2.3.	Sequential and parallel paths.....	183
3.2.4.	Convergence.....	187
3.3.	Model specificity and criticality.....	189
3.4.	Selection of documents.....	197
3.5.	Delimitation regimes.....	203
3.5.1.	Sequential delimitation	204
3.5.2.	Parallel delimitation	212
3.5.3.	Disjunctive delimitation.....	216
3.5.4.	Hybrids and subdivisions	226
3.6.	Points of synthesis and merging.....	227
3.7.	The output, spatiality and movement.....	233
3.8.	Conclusion.....	242
4.	(Digital) Scholarly Editions and Digitisation	245
4.1.	Introduction	245
4.2.	Scholarly Editions	245
4.2.1.	Editions of works.....	247
4.2.2.	Editions of documents	255
4.2.3.	Genetic editions.....	266
4.3.	Defining the scholarly edition.....	267
4.4.	Textual digitisation	268
4.4.1.	Digital archives.....	269
4.4.2.	Critical digitisation	272
4.4.3.	Large-scale digitisation	275
4.4.4.	Mass digitisation.....	284
4.5.	Conclusion.....	294
5.	<i>Letters of 1916</i>	295
5.1.	Introduction.....	295
5.2.	Project overview.....	297
5.3.	Selecting source material.....	299
5.3.1.	Conceptual delimitation.....	299
5.3.2.	Pragmatics of selection	305
5.4.	The transmission process	310
5.4.1.	Photographing and collecting preliminary metadata	311
5.4.2.	Uploading letters	314
5.4.3.	Crowd transcription and mark-up.....	316
5.4.4.	Automated processing and creating XML documents	321
5.4.5.	Final editing and validation	324
5.5.	Outputs and spatiality	328
5.6.	Conclusion.....	336

6. Towards a Rhizomatic Edition	339
6.1. Introduction	339
6.2. Rhizomatic rupture and the internalisation of connections.....	344
6.3. Rhizomes as graphs.....	352
6.3.1. Graph visualisation.....	357
6.4. Building a network interface	361
6.4.1. Preliminary considerations	361
6.5. Graphs, data models and data storage	364
6.6. Building the system architecture.....	366
6.7. Graph manipulation.....	375
6.8. Citation mechanisms.....	387
6.9. Using the <i>LetterGraph</i> system.....	393
6.9.1. Possible Criticisms.....	400
6.10. Conclusion	404
Conclusion.....	407
Appendix A: On Documents and Selection.....	419
Appendix B: List of Abbreviations and Technical Terminology.....	431
Works cited.....	436

Abstract

Scholarly editing and textual digitisation are typically seen as two distinct, though related, fields. Scholarly editing is replete with traditions and codified practices, while the digitisation of text-bearing material is a recent enterprise, governed more by practice than theory. From the perspective of scholarly editing, the mere digitisation of text is a world away from the intellectual engagement and rigour on which textual scholarship is founded. Recent developments have led to a more open-minded perspective. As scholarly editing has made increasing use of the digital medium, and textual digitisation begins to make use of scholarly editing tools and techniques, the more obvious distinctions dissolve. Such criteria as ‘critical engagement’ become insufficient grounds on which to base a clear distinction. However, this perspective is not without its risks either. It perpetuates the idea that a (digital) scholarly edition and a digitised text are interchangeable.

This thesis argues that a real distinction can be drawn. It starts by considering scholarly editing and textual digitisation as *textual transmissions*. Starting from the ontological perspective of Deleuze and Guattari, it builds a framework capable for considering the *processes* behind scholarly editing and digitisation. In doing so, it uncovers a number of critical distinction. Scholarly editing creates a regime of representation that is self-consistent and self-validating. Textual digitisation does not. In the final chapters, this thesis uses the crowd-sourced *Letters of 1916* project as a test-case for a new conceptualisation of a scholarly edition: one that is neither globally self-consistent nor self-validating, but which provides a conceptual model in which these absences might be mitigated against and the function of a scholarly edition fulfilled.

Acknowledgements

I would like to thank my supervisor, Professor Susan Schreibman, for her support, advice and tireless work in reading endless drafts of this thesis (especially as, for most its writing, it was far from clear just where it was heading). Also, all the staff in *An Foras Feasa* and, in particular, everyone involved with the *Letters of 1916* project, without whom I would now be looking at an empty page. Thanks likewise to Franz Fischer, Patrick Sahle and everyone in the Centre for eHumanities in Cologne for a wonderful secondment. This thesis would not have been possible without the extremely generous funding and outstanding programme provided by the Marie Skłodowska-Curie DiXiT Initial Training Network, including the invaluable advice provided by other supervisors in the programme.

I would also like to thank, for their warmest friendship, wise council, Guinness, and putting up with my endless wittering: Neale Rooney (*good coffee...!*), Vinayak Das Gupta, Alex Hughes, Richard Breen, Saman Fegghi, (still awaiting my *Catan* revenge on you lot), Jack Kavanagh, Chris Beausang, Linda Spinazzè, Roman Bleier, Hannah Healy, Emma Clarke and Karolina Badzmierowska. It's not easy moving to another country on top of completing a PhD, and you all made it so much more than bearable. To my friends in London, especially Daniel Turner (thanks for all the Deleuze books) and Elinor Bailey (for telling me to ignore the Deleuze books and submit the damn thing). Thanks also to all the DiXiT fellows ("you guys") for the intense, enterprising, endlessly fun camps and conferences spent doing Digital Humanities "at the point of a smoking pint", and for providing such an amazingly supportive network of friends only a message away.

So many thanks to my parents-in-law, Bob and Sue Emmerson, for supporting me during the final stages of writing this thesis; and to Paul Watkins, for a lovely little cottage-cum-writing-retreat. To my parents, Jane and Frank Hadden, for all they have done for me, and who have supported me in absolutely everything. And, especially, to Lucy Emmerson: a bit, of course, for proofreading this beast; but mostly for her love, support, encouragement, sanity, and everything, everything else (including, y'know, getting married to me).

Introduction

The aim of this thesis is to consider developments in text digitisation — specifically in light of traditional understandings of textual scholarship and scholarly editing. In taking this approach, there is something of a risk: that of, to use the common English idiom, ‘comparing apples and oranges’. The problem (to run with this analogy for a moment) is not that apples and oranges *cannot* be compared. There are similarities. They are both fruit. They both have pips. In contrast, we can also say that oranges are orange, and that apples are red or green; that oranges have a thick and inedible skin, unlike apples; and so forth. Several (evidently tongue-in-cheek) scientific articles (Sandford 1995; Barone 2000) have gone further, with approaches that either involve a selection of arbitrary characteristics (measuring size and perception of sweetness in the case of Barone) or a highly reductive operation (Sandford’s dissolving both in potassium bromide and performing spectroscopy on the resulting solution). The problem here is that we are really just enumerating (and then juxtaposing) the defining properties of said fruits.

This is just the kind of problem that the philosopher Gilles Deleuze sees with comparison rooted in a system of categories. A comparison of this sort only serves to perpetuate the categorical distinctions that were already established, to repeat established ‘defining properties’ of apples and oranges in perpetuity (oranges are x , apples are y). It is a replication of “good sense”, which “affirms that in all things there is a determinable sense or direction” (Deleuze 1969, 3) and “common sense”, the “assignation of fixed identities” (Deleuze 1969, 5). The problems with this are twofold. Firstly, it closes off potential for saying anything new. (Adkins 2015, 22–3) Secondly, it asserts a stability, a stasis, that is fundamentally

ahistorical: oranges are oranges, apples are apples; we must leave aside the fact that both evolved, and from a common ancestor. (DeLanda 2006 discusses more fully the relation between a Deleuzian ontology and evolution.)

As I said, the starting point of this thesis is a comparison of scholarly editing and text digitisation. And the problem — as with apples and oranges — is that such a comparison simply perpetuates distinctions already made: in other words, juxtaposing definitions. Just as significantly, it views each from the perspective of stasis — things must be as defined — which precludes viewing them as continually evolving enterprises.

On the face of it, of course, such differences are manifest. Scholarly editing has a long history, originating in biblical and classical textual scholarship: “the most ancient of scholarly activities in the West.” (Greetham 1992, 297) Text digitisation is something of a new-fangled upstart in comparison, founded evidently on digital technologies. Scholarly editing is replete with theoretical positions and practices, while for digitisation these are thin on the ground: practice leads the way. Scholarly editing is a scholarly practice, founded on “rigor” and “expertise”, that “produces knowledge”, while “mere digitizing produces information”. (Price 2007) Scholarly editing is concerned with ‘establishing’ a text, whereas digitisation reproduces a text. We could go on. While none of the above is untrue— scholarly editing and text digitisation are as different as apples and oranges — starting from such a perspective is problematic for a number of reasons.

Firstly, as Deleuze would suggest, any comparison dependent on such dividing lines will only reproduce them. (Scholarly editing is scholarly; digitisation is not. The former is critical; the latter is not. Digitisation must be digital; scholarly editing might employ the digital medium, but has shown itself to function very well without it.) A second, more fundamental, problem

is that none of these distinctions really stand up to scrutiny. Digitisation is dependent on the digital medium, but is it really defined by it? Or, put another way, is it the basis for a stable dividing line? A scholarly edition can be digital and digitisation must be, but what is the intrinsic difference between digital and print? As Mak argues, the digital medium and print are both technologies that enable reading: “Digitizations are not ... primarily technological, or, at least, no more or less technological than medieval illuminated manuscripts or printed newspapers.” (Mak 2013, 4) Of course, digital and print are different, but these differences are only evident when the two are framed in opposition to each other. If we compared the digital, print, and soft cheese with regards to their text-bearing capabilities, print and digital look like brothers-in-arms.

The same can be said of other dividing lines. One such distinction, between the critical establishment of a text and “uncritically reproducing the text of existing documents” (Eggert 2013, 104) was questionable even before the digital medium. Many scholars have convincingly argued that ‘uncritical reproduction’ of a text is hardly devoid of critical activity. (Kline 1995; Dahlström 2012; Pierazzo 2011, 2014) ‘Critical editing’ is, in other words, distinguished only by a particular definition of ‘critical’ (on the whole, the definition given by critical editors). Moreover, as Pierazzo says, ‘non-critical’ documentary editions have proliferated thanks to the affordances of the digital medium. (Pierazzo 2014) Of course, some digitisation might well be regarded as highly *uncritical* by any reasonable standard (mass digitisation, most obviously), but other approaches (such as that which Dahlström, 2012, describes as ‘critical digitisation’) clearly are critical. Indeed, drawing a line between digital documentary editions and these kinds of digitisation is, I would contend, all but impossible.

A similar case can be made for saying that scholarly editing is ‘scholarly’ and text digitisation is not. As intrinsically tautological, such a definition is either useless, or requires an

articulation of what is meant by ‘scholarly’. Taken as *carried out with scholarship*, the question then arises, ‘how much?’ Once we have reached the point of placing a characteristic on a continuum, its functioning as a dividing line becomes impossible. There is, finally, the (perhaps most insidious) suspicion that scholarly editing equals good, digitisation bad. I recognise I am perhaps on shaky ground here. That said, I think a case can be made that a lot of the criticism of digitisation — and here I am thinking of mass digitisation especially — that comes from textual scholars is made from the perspective of textual scholars. Not that these criticisms are invalid: accuracy is often a problem (Coyle 2006, 643), and Nunberg’s description of Google Books as a “metadata train wreck” (Nunberg 2009) is certainly well-substantiated.¹ However, it could be argued from the other side of the fence that making a lot of text available in digital form is useful in itself; that quantity is, in essence, another quality: it allows scholarship of a different kind, predicated on scale. This is not an argument I intend to pursue further, but it raises a fundamental point: that considering digitisation in terms of ‘what it lacks compared to scholarly editing’ is a position necessarily rooted in textual scholarship.

A third difficulty with direct comparison has already, to an extent, been suggested: that neither scholarly editing nor text digitisation are singular activities. Scholarly editing, indeed, might be regarded as an overarching label for many activities: critical editing, documentary editing, genetic editing; and various traditions: Lachmannian stemmatics, Greg’s copy-text theory, social editing, the German school of textual criticism (itself hardly singular). Textual digitisation ranges from mass digitisation, as practiced by Google, to more considered large-scale digitisation, crowd-sourcing, and critical-digitisation. Even when such pluralities of

¹ Various political and ethical questions — most obviously, is it right for large, private companies to govern access to knowledge? — are also, of course, of profound importance.

practice can be grouped together, there is some danger in doing so: not least that it tends to gloss over the development of these pluralities.

Fourthly, and perhaps most significantly, drawing distinctions in such a way ignores the co-evolution of text digitisation and scholarly editing since the advent of the digital medium. The fields, in practice, have not been isolated. Textual scholars used computers, for instance for collation, long before the digital was a viable medium for dissemination or preservation. (Hockey 2004) ‘Digitisers’ in turn have borrowed concepts and tools from textual scholars, not least the Text Encoding Initiative, which began as an open standard for describing digital texts but has received substantial input from various schools of scholarly editing. (Hockey 2004; Schreibman 2013) The two fields can therefore be seen as forming what Deleuze and Guattari describe as a “rhizome” — a system with no hierarchy, that proceeds instead by lateral connections among myriad entities and concepts. (Deleuze & Guattari 1987, “Rhizomes”; Adkins 2015, 23) What is therefore under consideration is a more complex interplay of systems that does not easily decompose into (even) two distinct fields. To attempt to do so by pre-emptively defining scholarly editing and textual digitisation is to insist, externally, on arbitrary division, and, internally, on unwarranted homogenisation.

Accordingly, we should resist the urge to ‘unite’ scholarly editing and digitisation under some new banner. A rhizomatic system, as Adkins suggests, resists takeover by hierarchies. (Adkins 2015, 23) This would only serve in turn to create a larger category, which would be just as exclusive and unsustainable in the long run. Instead, a rhizome should be laid out on what Deleuze and Guattari call a “plane of consistency”. (Deleuze & Guattari 1987, 4) Not a higher category, but a kind of bedrock.

What is the “plane of consistency” here? In this thesis, I suggest, it is *textual transmission*. Every kind of scholarly edition and digitisation transmits a text; that is, it takes a text that already exists and transmits it to the user via *remediation*. Not a spectacular definition, but one that operates as a *sine qua non*. If a scholarly edition or digitisation does not transmit a text, it has failed at its task. In fact, it is not a definition at all, as (evidently) a scholarly edition or a digitisation are not the only vehicles for transmitting an existent text. Every printed book transmits an existing text (a *non*-scholarly edition, if you like). But the aim is to not build a higher category — which would involve introducing yet more arbitrary distinctions — but rather to consider what has been laid out.² Textual transmission, I would suggest, is what is left when the specificities of particular kinds of digitisation or edition are subtracted: “Subtract the unique from the multiplicity to be constituted: write at $n - 1$ dimensions.” (Deleuze & Guattari 1987, 7) Such an approach also avoids the problems arising from considering textual transmissions on any kind of continuum. A continuum is nothing more than the extraction of one particular dimension of variation, and its elevation to a higher category. When Sandford dissolved apples and oranges in bromide solution for spectroscopic analysis, he was in essence reducing the difference between apples and oranges to that which could be measured by spectroscopy. (As if spectroscopy could explain why an apple pie was a thing and an orange pie not!)

Adopting this position, I would argue, represents the first key insight presented in this thesis. It serves as a point of departure for a less restrictive, and less normative, analysis, as well as forcing a reconsideration of how various textual transmissions might be distinguished. The

² That said, it is certainly not the aim of this thesis to stray beyond its parameters (mostly). Its concentration is on scholarly editions and different shades of text digitisation, though it occasionally nods to other textual transmissions.

reason for this is that textual transmission is a dynamic process. It does not assume as a prerequisite that, for instance, a scholarly edition of a document automatically ‘represents’ that document — indeed, it starts from the more plausible notion that they are completely different — instead, asking what it is about the process of transmission that gives rise to representation. Accordingly, differences between various kinds of textual transmission can be seen as arising from differences in process.

The first chapter of this thesis aims to more fully set out this starting position. In doing so, it covers, first, textual scholarship from a historical perspective, examining key debates within the field, and investigating the varying premises behind the many practices that fall under the rubric of scholarly editing. The second section is concerned with the advent of the digital medium, its early uses for studying texts, and gradual utilisation for the creation of scholarly editions. The final section is concerned with other kinds of textual digitisation, and some of their rationales. In doing so, it elaborates on some of the themes articulated above.

Chapter Two focuses on developing a theoretical basis for textual transmission. As suggested above, transmission must be considered as a process (not a product): that a transmission occurs at all depends on its *success*. A failed transmission is not a transmission. In developing this model, I draw again on the work of Deleuze and Guattari, considering the various ‘actors’ in the transmission — the text itself, but also the editor, the reader, and the ‘output’ — as *assemblages*. An assemblage is, for Adkins, Deleuze and Guattari’s answer to the question “What is a thing?” (Adkins 2015, 10) Most importantly, it allows for consideration of ‘things’ not as stable identities, but as dynamic systems of interconnections. In the third chapter of *A Thousand Plateaus* (1987), Deleuze and Guattari develop an approach to understanding the interaction of assemblages. (Deleuze & Guattari 1987) This is particularly useful here, as

textual transmission is, I would strongly suggest, a dynamic process. As Drucker says, “No text is ‘transferred’ wholesale, like a bucket of coal being moved along a conveyor.” (Drucker 2013) Rather, transmission is achieved by the interaction of assemblages.

Deleuze and Guattari’s perspective here can be seen as an extension of the ideas developed by Deleuze in *Difference and Repetition* (1968). In this work, he argues that *difference* should be considered ontologically prior to *identity*. Developing what he describes as a “metaphysics of differential calculus” (Deleuze 1968, 280), Deleuze argues that it is difference that creates identity, rather than difference existing as a secondary effect as the ‘difference between’ stable things. It is a metaphysics that allows for continuous variation (Adkins 2015, 6), with identities emerging as the temporary slowing-down of change.

This perspective undoubtedly creates a problem for textual transmission: for a transmission to succeed, there must be identity between the source and the output of the transmission — instead, we have difference. The same view has been articulated in many ways by textual scholars, especially since the emergence of post-structuralist philosophy. In this chapter, I draw direct comparisons between text as a Deleuzian assemblage and McGann’s consideration of text as a Bakhtinian space. (McGann 2003, 3) The problem, as Pierazzo suggests (2011), is one of objectivity of interpretation: chiefly, that it does not exist. If there is no objectivity, and entities are fundamentally different, how can a transmission claim to represent that which it transmits? The notion of rigour is also thereby brought into question. Rigour itself cannot be seen as some abstract property, but rather as ‘rigour *with respect to something*’. As text can have no ideal ‘essence’, it cannot be the subject of any ‘rigorous’ practice. The question posed here is, therefore, what *creates* ‘rigour’?

Here, I think, this thesis makes an original contribution to knowledge. I take Eggert's notion of the literary work as a "regulative idea" (Eggert 2009, 235) and extend it as a general 'regulator' for all textual transmissions. The regulative idea produces what Deleuze and Guattari call a "signifying regime" (Deleuze & Guattari 1987, 124), one in which identity operates: a small domain of localised objectivity, something against which rigour can be measured. The most fundamental point is that this is necessarily an editorial creation, produced by reflexive consideration of a textual assemblage.

Chapter Three aims to further develop the logic of textual transmission as a process, proposing that any transmission can be considered as a series (or recurring loop) of interpretative functions. It is, continuing the argument of the previous chapter, this repetition that enables the reflexive production of a regulative idea. The chapter then turns to discuss selection as an interpretative process. It explores the various connections between documents that give rise to certain forms (for instance, a work). These forms, it is suggested, map onto Deleuze and Guattari's three 'thought images', the tree, the radicle and the rhizome. (Deleuze & Guattari 1987, "Rhizomes") Only a tree, which designates a signifying regime, allows for identity; other forms produce different degrees of representation.

This is particularly important when considering transmissions of multiple texts, especially in large, digital collections. To illustrate this, we might take two scholarly editions, say of *Ulysses* and *Leaves of Grass*. These are governed by their own regulative ideas; they are separate works. Now bind the two together in a single book. The obvious question is, what *is* this book? What does it represent? Two things joined together represents not only the component parts but something more: the 'rationale' for the two combined, and how they are combined. This chapter ends by considering the implications of such situations by exploring the spatiality of

the transmission output (how it organises text, and enables or restricts movement) — in other words, asking not just what comprises the collection, but how the fact of being collected, and the organisation of the collection, in turn affects the component parts.

In Chapter Four, I undertake a brief survey of various kinds of scholarly edition and textual digitisation, using the approaches and models considered in the previous two chapters. Chapter Five applies the same methodology to the *Letters of 1916* project (on which I worked in conjunction with this thesis during my time at Maynooth University). *Letters* is most interesting as a case-study, as it combines many different approaches drawn from textual scholarship, as well as digitisation practice, and crowd-sourcing. By considering the development of the output as a process, it is possible to see the effects of these various approaches and how they operate in practice. This analysis suggests that *Letters* is not a scholarly edition (not, indeed, that it claims to be) in a global sense, but that it is valid to consider it a ‘collection’ of scholarly editions at the level of the individual component letter. This in turn, however, places fundamental restrictions on how a user can validly interact with it.

Chapter Six comprises a second original contribution to knowledge. Using the *Letters of 1916* collection, it argues that a dynamic, rhizomatic system can be approximated in the digital medium as an interactive graph. I argue that such a system can be considered a scholarly edition. Rather than achieving validity by the creation of a single *regulative idea*, it instead models the contingencies of production as part of the ‘edition-text’. This presents a very different paradigm. Instead of a certainty provided by the editor, the user is instead given enough information to consider both individual letters, and, more importantly, historical networks of letters, while taking into account the contingent bases for their representation in

the collection. It can be the object of what Deleuze and Guattari call “nomad science” — an interaction not with “discrete and ideal essences” (Adkins 2015, 14) but with the messiness of actual things.

*

In this thesis (as is apparent by now), I have made great use of the philosophy of Gilles Deleuze, often considered a part, albeit a highly idiosyncratic one, of the French post-structuralist school (I also utilise his collaborations with psychoanalyst Félix Guattari, particularly *A Thousand Plateaus*). This is for many reasons. Most notably, I find great appeal in Deleuze’s philosophy being grounded in a coherent metaphysics. “Unabashedly doing metaphysics” (Adkins 2015, 1), as Smith and Protevi note, puts him somewhat at odds with other philosophical tendencies, not least the Heideggerian notion of the “end of metaphysics.” (Smith & Protevi 2017) Deleuze’s metaphysics, aimed to be a metaphysics adequate to deal with modern science and mathematics (Smith & Protevi 2017), is characterised by “continuity ... a continuity of the sensible and the intelligible” (Adkins 2015, 1-2) or of “univocity of being” (Smith 2001). Everything exists in the same way; there is no ontological distinction between the real (the sensible) and the ideal (the intelligible). As such, there is no need to appeal to analogy, which for Deleuze fails to explain the link between the sensible and intelligible. For example, to say that ‘God is good’ and that ‘man is good’ requires an analogical relation between God’s goodness and human goodness (Smith 2001, 169); we are not talking about the same kind of goodness at all, and the connection between these two kinds of goodness cannot, for Deleuze, be adequately explained.

This puts Deleuze at odds with the dominant trend of Western metaphysics, which, beginning with Plato, has insisted on an ontological divide between the real and the ideal.

(Adkins 2015, 3-4) (Deleuze traces in his early monographs an alternative lineage of philosophers of “univocity”, including Duns Scotus — who coined the term —, Spinoza, Hume and Nietzsche: Smith 2001.) For Deleuze, *identity* — that is, stability — is subordinate to *difference*: again, an inversion of traditional Western metaphysics. Difference, as ontologically prior to identity, is what drives dynamic systems; identity is a secondary phenomenon, a slowing down or stabilisation of change.

Furthermore, Deleuze’s philosophy – unlike other post-structuralists, most obviously (later) Barthes, Foucault, and Derrida – is not primarily concerned with language. Rather, with its focus on materiality and, above all, process, it can be seen as providing a kind of underpinning for more obviously linguistic-oriented post-structuralism. This makes his work an ideal fit for investigating textual scholarship, a field that can be seen as residing at the intersection of materiality (documents as material objects) and language.

Starting from such a perspective is useful in considering textual scholarship (as already suggested), as it allows the sidestepping of questions concerning, for instance, the nature of a literary work and how the ‘essence’ of such a work might be perpetuated by scholarly editors. In particular, it inverts the hierarchy of *work* (at the top) and its physical manifestations (all of which are necessarily imperfect representations of the ideal — being, really, only analogical). The idea of the work is just as real as any physical manifestation, and, moreover, the idea of the work is a kind of emergent phenomenon arising from physical manifestations. As such, it is not fixed, but (as Eggert also argues: 2009, 235) constantly unfolding through interaction with material manifestations. This is not to say that the ‘idea of the work’ (the work *really* taken up in thought) does not have any effect: the work-in-thought organises, in turn, the physical manifestations, making, say, the manuscripts of

Ulysses what they are (as opposed to simply pieces of paper that Joyce wrote on).³ The physical and the ideal are both real, and together constitute a heterogeneous system (an assemblage). Unlike Eggert's exploration (2009), this is not a theory therefore of what a work *is*; such a thing, as Foucault states, does not exist (Foucault 1969, 104). Rather, it points to a theory of how we (editors) 'know' a work and thus transmit it.

This has even more implications for textual transmission, as I suggested earlier. It is not possible to view the source of a transmission and its output (say, the manuscripts of a literary work and the resulting critical edition) as simply manifestations of the work: not that they are *not* manifestations of the work in one sense; rather that, in the absence of an Ideal work, there is nothing to point to in order to underpin the relationship between one and the other. (And if there were an ideal work, it would be fundamentally inaccessible.) As such, the relationship between the source and output, and the validity of that relationship, must be grounded in process (or performance: Drucker 2014).

In *A Thousand Plateaus*, Deleuze and his collaborator Félix Guattari provide many tools for considering such processes. Considering 'things' as assemblages is one step towards this, allowing a focus on instability and change rather than on stasis and identity. Assemblages are discussed at greater length in Chapter Two. They also articulate various mechanisms or

³ It also, incidentally, answers Eggert's 'riddle': "If the *Mona Lisa* is in the Louvre, where is *Hamlet*?" The answer from a Deleuzian perspective is not that Hamlet-the-work is in some distinct realm of ideas, but is really constituted — to varying degrees — by all the pages with Hamlet printed on it, and in the minds of all the people who have read the play and seen it performed. The *Mona Lisa* is the same, though its form (a painting, not language) means that it is less easily distributed (*deterritorialised*, in Deleuze and Guattari's terminology in *A Thousand Plateaus*). Still, people have heard of the *Mona Lisa*, and really discuss it and write about it, and could continue to do so even if the painting were destroyed. Even the *idea* of the work is real.

“abstract machines” (Deleuze & Guattari 1987) that describe how different assemblages are able to interact and either stabilise or give rise to new things. In “The Postulates of Linguistics”, for instance, they describe the relationship between language and the physical: language, far from being an abstract idea, is real, and *affects* the physical by giving order to it. (Deleuze & Guattari 1987) Massumi is particularly eloquent in describing this process with relation to a woodworker carving wood: it is an interaction between two assemblages (the wood, with its history — growth — giving rise to grains and other properties to be interpreted; and the woodworker, an assemblage of person, tools, knowledge and experience). (Massumi 1992) In Chapters Two and Three, I make the same case for textual transmission. It is the interaction between multiple assemblages: documents, an editor — or editors, or editors and computers: an assemblage —, an output, and the user. The challenge, as noted above, is to describe how such dynamic processes give rise to a degree of identity between source and output.

Deleuze’s writings are noted for their difficulty, and *A Thousand Plateaus* is no exception. It is replete with terminology that is either invented or borrowed from a wide number of fields (literature, biology, geology...). For this, they have received their fair share of criticism, not least from Sokal and Bricmont (1999), for *misuse* of scientific and mathematical terminology. (They also, it is fair to say, level similar criticisms at virtually all French structuralist and post-structuralists.) Smith and Protevi, however, defend Deleuze and Guattari from “intentional obscurantism”, instead noting that, in their use of scientific terminology they are “seeking to form ... properly philosophical concept[s] by means of extracting certain forms of thought”. (Smith and Protevi 2017) However, the fact remains that their writings are difficult, and thus critical guides such as that of Adkins (2015), Massumi (1992) and Williams (2008) have been invaluable in aiding my understanding.

This difficulty notwithstanding, the works of Deleuze and Guattari have found increasing application in recent years, and in fields that are at least tangential to textual scholarship. In information science, the concept of the rhizome is often cited. Burnett, and Robinson and Maguire, apply the concept of the rhizome to information organisation, particularly in hyperlinked systems such as the Web. (Burnett 1993; Robinson & Maguire 2010) Sukovic also explores the rhizome as a model for information structures in libraries. (Sukovic 2008) Landow also considers hypertext as a rhizomatic system. (Landow 2006, 60) From those closer to textual scholarship, Drucker's view of textual transmission as 'performance', though not referencing Deleuze, evidently shows the same orientation (Drucker 2007); her article "Distributed and Conditional Documents" (Drucker 2014) also marks a distinctly assemblage-like view of text. From further afield, DeLanda has further developed the idea of *assemblages* and has applied such theories to history, science and social science. (See DeLanda, *Assemblage Theory* 2016) In a somewhat similar vein is Latour's "Actor-Network Theory" approach to social science. (Latour, *Reassembling the Social*, 2005) In the introduction to this work, Latour questions the validity of attributing aspects of a given scenario or formation to a 'social aspect,' as if 'the social' were an external and singular phenomenon; instead, he argues that 'social' elements are inseparable, a fundamental part of a complex network that gives rise to formations. (Latour also suggests that ANT could equally have been named "actant-rhizome ontology", acknowledging the connection of his work to that of Deleuze and Guattari: Latour 1999)

In general, therefore, I think utilising some of the concepts developed by Deleuze and Guattari is justifiable. A perspective developed far from the realms of textual scholarship offers new scope for analysis. That said, by adopting a theoretical platform that is not explicitly informed by textual scholarship, there is a risk of reproducing ideas and theories

that have been articulated elsewhere. Where this occurs, I try to acknowledge the overlaps in perspectives. At the same time, I do not think that this is necessarily a bad thing. A good — and even an original — theory does not override all existing knowledge. If anything, adding to the weight of evidence for an idea by considering it from a different direction strengthens the validity of both perspectives.

*

Finally, I would like to briefly mention the *Letters of 1916* project. The project began in 2013 at Trinity College Dublin, under the directorship of Susan Schreibman (also my PhD supervisor), before moving to Maynooth, along with Professor Schreibman, in 2014. The project aims to collect and digitise letters and other epistolary documents from the period of the Easter Rising (April 1916), using a crowd-sourcing methodology to complement expertise in text digitisation and (digital) history:

Letters of 1916 is Ireland's first crowdsourced public humanities project. It has created an online collection of letters about Ireland written around the time of the Easter Rising. Letters cover all aspects of life, from major political events, the Easter Rising and the Great War, to letters that provide us with a candid and unmediated glimpse into ordinary life, love and work a century ago. There are thousands of letters in the collection connecting hundreds of lives. (Schreibman 2016)

In 2017, the project was awarded a new grant to extend its period of interest, and now covers the Irish 'revolutionary period' from 1916 to 1923. (It should be noted that my research and involvement with the project was carried out before this extension; I consider the *Letters of 1916* in its original incarnation.) The project has been an immense success, collecting several thousand letters, engaging the support of over 1500 volunteers (predominately in the crowd-

sourcing of transcriptions), and producing high-quality TEI encoded letters and a full-text searchable database. I discuss the project more fully in Chapter Five of this thesis.

My work on the project acted as a formative background to my PhD programme, and from the outset I envisaged using *Letters* as a test case. In particular, I sought to answer the question of how *Letters* might be ‘reconciled’ with traditions of scholarly editing.⁴ At first glance, this creates an obvious categorical problem, which may be looked at one of two ways. Firstly, we may simply say that *Letters* is a scholar edition (even though it does not consider itself to be one). Of course, it differs in many ways from what might traditionally be considered a scholarly edition — but, as already suggested, textual scholarship today is a highly pluralistic field, and adding one more ‘category’ (say, the ‘Crowd-sourced Edition’) would not pose any particular problem. But this approach essentially serves to close down enquiry: if *Letters* is a scholarly edition, the ‘reconciliation’ has essentially already been done. As a result, we are reduced to merely enumerating established differences: traditional scholarly editions are more selective, *Letters* more open; *Letters* uses crowdsourcing, while scholarly editing is normally carried out by experts; and so forth. The latter point of distinction, indeed, raised several obvious questions. True, *Letters* crowdsources transcriptions, but it also has those transcriptions checked by experts. Is there, then, a fundamental difference? The answer, from this point of view, is that it hardly matters: it *is* a scholar edition, because that position was taken as axiomatic to the argument. The differences are already built into our definition of ‘Crowd-sourced Edition’.

⁴ This is, indeed, a stated requirement of my funded research project: ‘[e]xplor[ing] methodologies common to digital scholarly editing practice and mass digitisation’ and ‘[i]nvestigat[ing] ways to reconcile the differing methodologies’. (DiXiT Initial Training Network, <http://dixit.uni-koeln.de/fellowships/early-stage-researchers/#esr1>)

What of the other perspective: that *Letters* is not a scholar edition? This, I think, makes the problem even more intractable. For one thing, it mandates a more exclusive view of textual scholarship, denying its obviously pluralistic nature (probably to the exclusion of other kinds of scholarly edition too), and, accordingly, leads us to ask, ‘What is *Letters*, then?’ The truth, I would suggest, is that it again hardly matters. Whatever we decide *Letters* is, we are once more reduced to repeating the same categorical distinctions outlined above, without any way of saying why it matters. But, it must be observed, this approach is not only fruitless but nefarious: it sets up *Letters* as something defined primarily by deficiency: a failure to live up to the standards of scholarly editions.

Finding a way of unpicking this knot has driven the development of the theoretical ideas presented in this thesis: a theory of scholarly editions capable of dealing with *Letters*, one whose first move is not to arbitrarily include or exclude *Letters*. Or, rather, it is not a theory, but a methodology. Its first step, as argued earlier, is not the establishment of categories that embed difference, but the effacement of categories: laying things out on a plain of consistency, on a bedrock of textual transmission. This leads in turn to an analysis of process. How does a scholarly edition ‘transmit’? How does *Letters* transmit? It is only by taking a step back that such questions can really be tackled: expert transcription and crowd-sourced transcription should be considered different if they produce different results. Such an approach is, in one sense, more abstract — it involves an analysis of the abstract mechanisms underpinning the process of transmission — and simultaneously more concrete. Such processes are not simply theoretical; they are real, they must really happen, and they give rise to real differences in transmission.

1. Literature Review

I was astonished at how diverse [definitions of 'life'] were... But when you come to think about it, a collection that includes a fruit-fly, and Richard Dawkins, and the Great Barrier Reef, is an awkward set of objects to try and compare.

— Douglas Adams

1.1. Introduction

In “On the term ‘text’ in the digital humanities”, Caton articulates two possible models for defining ‘text’. The first model imagines ‘text’ as a series of concentric circles, in which there resides a stable core that can be seen as *defining* text. He gives the example of jazz music. Some jazz, like Charlie Parker, is more evidently ‘jazzy’ and resides towards the core, while “other kinds of freely improvisational music” occupy the periphery. (Caton 2013, 209) In such a model, there are what might be described as ‘core tenets’ of jazz, certain properties shared, albeit to a lesser or greater degree, by all exemplars: a “*relative definition*”, as Caton puts it. A second model, Caton suggests, is based on Wittgenstein’s conception of the term ‘game’. While there are local “similarities, affinities”, there is no central core; instead, “a complicated network of similarities overlapping and criss-crossing: similarities in the large and in the small.” (Wittgenstein 2009 [1953], 36) There is much similarity between this description and Deleuze and Guattari’s concept of the rhizome (Deleuze & Guattari 1987).

According to Burnett:

[T]he rhizome is an acentered, nonhierarchical, nonsignifying system without a General and without an organizing memory of central automaton, defined solely by a circulation of states. (Burnett 1993)

In his article, Caton makes a convincing case for ‘text’ being considered a concentric, ‘jazz’ system. What are we to make, by contrast, of scholarly editions and scholarly editing? Defining a scholarly edition or scholarly editing is a particularly difficult task (as if defining ‘text’ were easy!): there is always the risk of presenting a view centred on a given school of editing, to the exclusion of other perspectives. The question is whether there is really a ‘core’: a set of shared characteristics, the possibility of a “relative definition”.

Let us begin therefore by looking at several definitions of a scholarly edition. For Shillingsburg, a scholarly edition “preserve[s] or rescue[s] a work of artistic, social, intellectual, or historical importance as an artifact” (Shillingsburg 1996, 3). Gabler’s definition, by contrast, is more pluralistic, including texts of documents in addition to works: “[A] scholarly edition is the presentation of a text — literary, historical, philosophical, juridical — or of a work (mainly, a work of literature) in its often enough several texts, through the agency of an editor *in lieu* of the author of the text, or work.” (Gabler 2010, 44) Sahle takes things further in this direction: “A scholarly edition is an information resource which offers a critical representation of (normally) historical documents or texts.” As he goes on to argue, a work, from an editorial perspective, is no more than a particular configuration and critical representation of document texts. (Sahle 2014)

Other theorists, perhaps recognising the impossibility of a comprehensive, all-inclusive definition, steer clear of offering one. Greetham’s primer on textual scholarship chooses to articulate many forms of scholarly edition (critical, non-critical, documentary) without

elevating any to the level of a definition. (Greetham 1992, Chapter 9) Pierazzo, likewise, carries out a detailed investigation of traditions and schools within the practice of textual scholarship, without defining a scholarly edition. (Pierazzo 2015) The volume *Scholarly Editing* (Greetham, *ed.* 1995) provides a collection of twenty-four essays relating to the practices of editing various national literatures, without producing an overall definition. Indeed, Tanselle, in the introductory chapter, presents a tree-diagram highlighting connections between various sets of practices (starting with a division into “Historical” and “Nonhistorical”). (Tanselle 1995, 11)

Price in particular notes the use of “‘scholarly’ [as a] descriptive word [...] applied to numerous approaches: authorial or social, critical or documentary, genetic, eclectic, or best text”. (Price 2009) In such a view, ‘scholarly’ as an adjective can be seen as a marker of quality, related to myriad practices, rather than a prescribed form. An earlier essay by Price supports this view: “An edition is scholarly both because of the rigor with which the text is reproduced or altered and because of the expertise brought to bear on the task.” (Price 2007) Kelly, likewise, points to expertise as a principle element (while opting for a pluralistic definition in line with Gabler): “A scholarly edition is a publication that provides an important work of literature or historical document that has been prepared by experts in the field.” (Kelly 2015)

Here we see not so much a definition as an articulation of shared properties: rigour; produced with expertise. Kline’s definition is likewise quality-oriented: “An edition with texts established and verified according to the standards of the academic community.” (Kline 1998, 273) Again, not what *is* a scholarly edition, but what is expected of it. Bowers’ justification for the scholarly establishment of texts echoes just this sentiment: a text riddled with errors is a dubious basis for further scholarship. (Bowers 1959, 2)

If ‘scholarly editions’ have a core, therefore, it would seem to be inhabited primarily by adjectives. This, I would argue, does not amount to a definition, even a relative one. Indeed, I do not think it is unreasonable to suggest that terms such as ‘rigour’ and ‘expertise’ serve as a proxy for ‘scholarly’ anyway. A scholarly edition is an edition that is scholarly, that is produced by acts of scholarship.⁵ Moreover, such a basis cannot operate as a central core in a concentric model of scholarly editions and scholarly editing. For one thing, such notions as ‘being rigorous’ are largely binary. A relative definition, like music with a trumpet and syncopated rhythm being ‘somewhat jazzy’, or a particular text ‘somewhat textual’ in Caton’s argument, does not work in the case of scholarly editing: something is not ‘somewhat a scholarly edition’ for being ‘a little rigorous’. In this light, the “network of overlapping similarities” of Wittgenstein, or (I would contend) even more so the Deleuzian rhizome, present a better model: a dispersed set of practices, with local similarities; and a *connection* to notions of scholarship, expertise and rigour, without elevating them to a defining core. This is the view adopted here. In doing so, the wide variety of things that purport to be scholarly editions can be investigated — and their relations historically situated — without any overbearing normativity. I will not begin by defining what a scholarly edition *should* be.

Having argued the above, it is important to note that this description relates to the field of scholarly editing as it currently exists. This does not mean that, historically, scholarly editing was not more restricted, growing outwards from a central core, rather than being pluralistic from the beginning. Indeed, it would be surprising if the latter were the case, as it would suggest that all the schools and methods and practices sprung up simultaneously and fully-

⁵ Perhaps the best that can be offered (as a here-unnamed textual scholar once wryly suggested to me) is to simply say that we know a scholarly edition when we see one.

formed. Rather, I think a field that once did have a solid centre has become multiple, and fundamentally decentred. I think the reading of the literature surrounding scholarly editing, in the next section, will support this argument. For the present, I would like to suggest several mechanisms that have operated, at least to a degree in concert, to bring about this shift.

The first is the prestige attached to the label *scholarly edition*. When organisations such as the MLA Committee on Scholarly Editions⁶ place their stamp of approval on certain kinds of editions — favouring in all but very recent versions of their guidelines adherence to the Anglo-American, Greg-Bowers line (Sutherland 2013, 51) — it is natural that those working from different theoretical bases, but with no less rigour or scholarship, should also expect recognition. The second mechanism has been the theoretical reconsideration of some of the tenets of ‘core scholarly editing’, finding them less-than-secure principles on which to draw dividing lines. One such distinction is that between *critical* and *non-critical* editing (a perfect choice of terminology to make clear which is considered the superior⁷). As Kline argues, however:

Documentary editing, although noncritical in terms of classical textual scholarship, is hardly an uncritical endeavor. It demands as much intelligence, insight, and hard work as its critical counterpart, combined with a passionate determination to preserve for modern readers the nuances of evidence. (Kline 1998, 2)

When being critical can be shown to apply equally to *critical* and *non-critical* editions, the validity of the divide no longer holds. In this case, documentary editing has now been well-

⁶ Formerly the Centre for Editions of American Authors (CEEA).

⁷ The latter, in Greetham’s words, “[s]imply reproducing a text already in existence”. (Greetham 1992, 347; my emphasis)

established as a methodologically independent branch of textual scholarship. Another such expansion was provoked by the meeting of Anglo-American and German editorial theories (what Eggert calls “the German encounter”: 2009, 203), forcing a mutual recognition of the validity of each approach. A third mechanism is the necessary adaptation of existent editorial practices to new materials and textual situations. Greg’s *copy-text* theory of editing, for example, can be seen as an adaptation of earlier paradigms to the works of Shakespeare. (Gabler 2012, 25) And finally, there are transformations to the field brought from outside, not as an additional paradigm (as in documentary editing), but causing internal reconfigurations. These reconfigurations can be local or global, determined by prevailing philosophical paradigms or technologies. As an example of the former (to grossly, for the moment, oversimplify), how can a critical editor appeal to authorial intent following Barthes’ declaration of the ‘Death of the Author’ (Barthes 1968)? At a more global level is the impact of digital technology, which has profoundly disrupted almost all theoretical paradigms within scholarly editing, as well as (arguably) extending the field with the addition of new paradigms.⁸ In this light, new scholarly editing paradigms can be regarded as only temporary stabilities — coherent sets of practices, producing certain kinds of editions from certain kinds of texts or works or documents — within an otherwise shifting and dynamic field.

However, this does, I think, create something of a problem. With the heterogeneity of the field as currently construed, and the potential for further disruption lurking (if not already apparent), how are we to think of scholarly editing? Where should we draw boundaries, if at

⁸ I say ‘arguably’ as it is not clear to me that a ‘digital scholarly edition’ represents the addition of a new paradigm, rather than an adaptation of existing paradigms to the new medium, a point borne out by Dahlström: “To even talk about *digital* editions as one particular type of edition is ... debatable.” (Dahlström 2004, 20) This, however, does not preclude the creation of new paradigms: *de facto* digital, because the digital is a requirement, but not *because* of the digital; rather, due to subsequent theoretical developments.

all? A theory, Robinson argues, is necessary: “A theory of scholarly editions should offer a set of principles to guide practice. What is a scholarly edition? how should a scholarly edition be made? who should make it?” (Robinson 2013, 105) It is the intention of this chapter to consider such questions, though not (unlike Robinson) because I think scholarly editing needs normative restrictions imposed, but because answering such questions are a fruitful way of moving forward.

1.1.1. Basic terminology: document, work, text

As argued above, various schools and paradigms within scholarly editing comprise temporarily stabilised configurations, producing particular kinds of interplay between myriad aspects: conceptual understandings of key terminology (*work, document, text, author* and *authority; editor* and *editing*; descriptive terms such as *critical*), methodology for determining their representation, the kinds of works or documents being edited, the technological means by which they are represented, and the cultural and economic conditions in which they are produced. Here I shall begin with an outline of terminology, which is, in many cases, the most universal aspect.

The first is the *document*. In *Scholarly Editing in the Computer Age*, Shillingsburg says of the document: it “has physical existence” and “consists of the physical material, paper and ink, bearing the configuration of signs that represents a text” (Shillingsburg 1996, 47). There is little in this definition that is controversial, or really restrictive. Indeed, Shillingsburg himself enumerates other possible “[text] contain[ing]” or “incarnat[ing]” objects: “book, manuscript, phonograph record, computer tape” (Shillingsburg 1996, 174). The characteristic of ‘physical existence’ is important in this regard, particularly when that

physical existence is not as (seemingly) self-evident as in the case of a book or manuscript (or engraved rock). I am thinking particularly of digital documents, in relation to which one might be tempted to use terms such as ‘virtual’. But it is important to recognise that digital documents are just as material as any other kind of document. It may be common to describe the digital as comprising binary ones and zeroes, but this is an abstraction: the ‘binary’ code is encoded physically as electric charges on, for instance, a magnetic disk. (For a more detailed discussion of digital documents, see Kirschenbaum 2013; see also Appendix A to this thesis). Moreover, to draw a distinction between digital and other documents on the basis that the former require technological mediation is, I would suggest, misguided. All documents require a technological mediation — whether by glasses, microscopes, hyperspectral scanners, even *eyes* — and it would be a mistake to regard some as in some way ‘natural’ (i.e. non-technological). This is not, of course, to say that all documents are the same: each material embodiment has its own properties, which affect how a reader may interact with it.⁹

What is more notable about Shillingsburg’s definition — indeed, I would suggest, almost all conceptions of a document in textual scholarship — is its reciprocal dependence on the term ‘text’. As Eggert notes, “The documentary and the textual dimensions are interdependent. They are separable for purposes of discussion, but they are not separate.” (Eggert 2010, 183) This raises the important question of when exactly a ‘material’ attains the status of ‘document’. If we take Caton’s definition of text (alluded to earlier; also below), it is — logically — only after some physically encoded marks on some material have ‘resolved to language’ that we may say there is a text, and, hence, that the material is a document. This is

⁹ If one were perverse enough to print this thesis as a scroll rather than a book, a reader’s engagement with it would be notably different: just accessing this footnote would entail a farcical amount of ‘scrolling’ to reach the bottom of the ‘page’!

a question that will be addressed in more detail later (especially in Chapter 3). For the purposes of this overview, let us continue to the concept of the work.

A *work* is commonly understood as a single artistic or literary endeavour (OED: “A literary or musical composition or other piece of art”) — or even more broadly, an act of communication. According to Shillingsburg’s glossary, a work is, “the message or experience implied by the authoritative versions of literary writing”, of which “variant forms have the same name” (Shillingsburg 1996, 176). This posits the work as an ideal conception, over and above its material embodiment in documents and its linguistic form as text. In this we again see a conceptual interdependence. Eggert exemplifies this most prominently: “[the work] as constantly involved in a negative dialectic of material medium (the documentary dimension) and meaningful experience (the textual).” (Eggert 2009, 237) Treating the work as an ideal conception — or even, as Eggert would seem to suggest, an emergent phenomenon of the document and text — raises the question for editors of where exactly the work stems from, particularly on what grounds it can be determined. Typically, the *author* of the work is taken to be the person responsible for a work.¹⁰ What role as an *authority* over the text — as the “person [...] that has the right to generate or alter *text* for the *work of art*” (Shillingsburg 1996, 173) — the author has is more widely disputed. Gabler, for one, employs the Foucauldian concept of the ‘author-function’, allowing the role of the author in the originator work, but withholding absolute authority over the work. (Gabler 2012) By

¹⁰ In the broadest sense: the author is ‘intellectual originator’ who ‘conceives’ and ‘writes’ the work. Insofar as the term ‘author’ means anything, this would be generally what is understood. Whether the author can be seen as *the* (as in, the sole) creative force behind a work is highly debated, as later discussions (for instance, of McKenzie and McGann) will highlight.

contrast, Anglo-American editorial theory (most obviously that of Bowers and Tanselle) tend to view the author as the absolute authority.

Of the document–work–text triumvirate, *text* is perhaps the most elusive. In attempting to provide a base definition, it is useful to distinguish the document-text as the writing on a document — “The series of words and pauses recorded in a document” (Shillingsburg 1986, 171) — and the work-text, the words that, regardless of where they are written, constitute a work — “a complex of all the versions and deviations belonging to a work” (Martens 1995). Caton’s article (referred to above) provides a more exhaustive survey of usage and further elaborates on the distinctions between document-texts and work-texts. Caton’s formulation of an embodied, material text (that is to say, a document-text) as marks upon a page (or other medium) which “resolve directly to language” (Caton 2013) — as opposed to, say, pictures, which do not — seems adequate as a definition for text as an abstract noun. *The text or a text* — that is, text as a fixed, countable noun — would appear to correlate more closely with the ideal notion, the work-text; as Caton observes, delimiting ‘a text’ is most problematic.

The understanding of these general concepts is, however, affected (and reciprocally determined) by the specifics of a given school or tradition. Most important, perhaps, is the nature of the text or document or work to be edited. Pierazzo notes that “[traditions] are roughly determined by period [of the material to be edited]” (Pierazzo 2015, 17) However, she goes on to note that this “periodization is [...] only theoretical, as it is also dependent on the country and the discipline of the scholar.” (Pierazzo 2015, 17). In this, we see again a reciprocal reinforcement: tradition as stemming from the practice of editing particular kinds of text, but also the effect of tradition in reinforcing practice.

This is, evidently, far from a comprehensive overview of terminology (and, as the above paragraph suggests, hardly static ones). Other terms, more closely linked to specific practices, will be discussed in the coming sections. However, it is necessary to start from some base understanding.

1.2. Scholarly Editing Traditions

1.2.1. Critical editing

Greetham (1992) describes textual criticism as “the most ancient of scholarly activities in the west”, and by the sixth century, scholars had established a text of Homer to remove apparent errors that resulted from oral transmission. The intention behind such activities was, as in the case of Homeric texts mentioned by Greetham, the recovery of lost ‘original’ texts from the error-laden texts produced by scribes when copying manuscripts: “to separate authentic text from textual corruption in the transmission” (Gabler 2012). This trend culminated in Lachmann’s stemmatic approach, which remains the dominant framework for editing classical and medieval texts (Pierazzo 2015, 18). According to Pierazzo, the Lachmannian method, “[A]ims at reconstructing the ‘original’ work using significant errors made by scribes as a guide to reconstructing the genealogical relationships (organized in the *stemma codicum*) among the surviving manuscript sources (the witnesses).” (Pierazzo 2015, 18)

Such an approach is predicated on the state of classical and medieval texts: normally, works of literature whose original is lost, but survive through extant documents or witnesses. “Critically analysed, the results from collating all extant document texts are schematised in a graph, the stemma” (Gabler 2012) based on examinations of errors introduced by scribes in

each generation of copying. From this, a so-called ‘archetype’ could be produced by methodical removal of these errors. Stemmatics thus mandates an idealistic conception of text: the physical embodiment of text in a document may be of use to the editor in, say, dating manuscripts, but these characteristics are far from represented in the final edition. Instead, there is assumed to be an ideal form of the *work* itself, which can be posited and thence reconstructed. This must necessarily be the case, as the process requires the ‘extraction’ of text from its material form as discrete signs in order to compare witness-texts on a word-by-word or character-by-character basis. Such editions, therefore, are much more work-centric than document-centric; the documentary instantiation of the text is viewed as a corruption of the ideal work. The end product of this editing is a single text, with variants perhaps indicated by a critical apparatus (*apparatus criticus*).

Also central, although played down compared to later traditions, is the role of authorship. Despite, in many cases, disputed or unknown authorship, the notion of authorial intention — the text that the author wrote — is clearly implied by the aim of recreating a lost original text: it is, after all, presumably the author’s original that aims to be recreated, or that which, as far as possible, the archetype purports to represent.

The text of such an edition was — and is, in many traditions — considered a ‘critical text’, one “that derives from more than one source text” (Shillingsburg 1996). Indeed, such is the prevalence of, if not the Lachmannian method itself, the objective of this methodology, that it largely defined scholarly editing: to “rescue a work of artistic, social, intellectual, or historical importance” (Shillingsburg 1996, 4). As Shillingsburg earlier notes, the terms ‘scholarly editing and ‘critical editing’ have often been conflated (1996, 2). Accordingly, editions comprising the text of single documents were typically labelled ‘non-critical’, and

were concerned with merely “reproducing” the text of a document rather than “establishing” the text of a work (Greetham 347).

While, as Pierazzo suggests, Lachmannian stemmatics is still the dominant form of editing for classical and medieval texts, later-developed traditions in scholarly editing may be looked upon as an adaptation of stemmatics, or at the least elements of its methodology, aims and outcomes, to differing textual scenarios.

1.2.2. Anglo-American Editing

The ‘Copy-Text’ theory of editing developed by Greg, which forms the basis of much of the Anglo-American school of editing from the 1940s onwards, can be seen as, according to Gabler, “a special modification of the original endeavor of textual criticism as exercised upon classical and medieval texts” to early-modern and modern manuscripts (Gabler 1995, 2). The term *copy-text*, from which this theory takes its name, was coined by McKerrow to “indicate that early text of a work which an editor selected as the basis of his own” (Greg 1950, 19) and was employed by Greg himself to mean the choice of base text that was considered ‘authoritative’ in terms of *accidentals* (such as choice of spelling variants and punctuation usage). The rationale behind such an approach is that a printer or typesetter may adopt some spelling convention, for instance, that differed from the author’s usage — and should thus be corrected. The *substantive* readings, contrastingly, should be modified according to later print editions, which may reflect the author’s revisions. (Greg 1950, 21–2). As Pierazzo succinctly observes: “had authors been able to control the printing process in the same way they had controlled their own handwriting, surely they would have employed their own accidentals.” (Pierazzo 2015, 17)

This theoretical approach is intended to deal with situations in which the author's original manuscript (or manuscripts) exist along with multiple print editions, over some of which the author may have had an influence (Pierazzo 2015; Greetham 1992). As such, there is a stronger interplay and potential overlap between the creative process of authoring a work and the reproductive process of its distribution (Gabler 2012), remnants of which, especially the former, may still be extant (in stark contrast to medieval and classical texts).

Starting from Greg's work editing Shakespeare and other early-modern authors, this method recognises the differing textual condition of early-modern works: firstly, the advent of printing, which produces (ostensibly at least) identical copies, but also adds the involvement of others — typesetters, printers — to the 'creation' of texts; and, secondly, a more 'linear', as opposed to radial distribution of text, with newer print editions being based on previous editions (and potentially incorporating their errors). Perhaps the beginning of this approach, as Greetham suggests, was Greg's questioning of "the rationale that the most authoritative copy-text for an edition should be the last edition published during the author's lifetime" (Greetham 1992, 333), on the basis that an author's ability to maintain 'authority' over his work would wane (and ultimately cease), and errors from previous editions would accumulate. In Greg's approach to editing Shakespeare, he included amendments made in later texts that he judged to be authoritative — and with some justification (McGann 1983, 17). The effect of this is what has been termed the 'eclectic' edition: a text derived from multiple sources without the idea that a single witness ever represented the author's intention.

The copy-text theory was further developed and codified by, most notably, Bowers and Tanselle, to accommodate later modern texts, whose extant remains could comprise not just printed editions (in many variants) but original manuscripts, fair copies and corrected

printers' proofs (varieties of '*avant-texte*'). By the third quarter of the twentieth century, the influence of this theory had become pervasive in Anglo-American schools of editing (Greetham 2012, 334-5).

This approach to scholarly editing gave increased prominence to authorial intention (Pierazzo 2015; Greetham 1992; Gabler 2012), and, in particular, the notion of 'final authorial intention' (Bornstein). In part, such a focus on the author as the sole repository of 'authority' governing the accuracy of a text was due to the increasingly prominent position of authors, and a romantic notion of authorship. (McGann 1983) With medieval or classical texts, though the author was possibly named (or given an identifier: 'The Beowulf Poet'), their importance was largely downplayed (even as editors sought to resurrect the original texts). The editing of texts from the early modern period onwards brought a greater force to the author as a "real, historically situated individual" (Gabler 2012, 20), capable of asserting intellectual ownership of a work. This may be added, as Greetham observes, a "Romantic concept of the solitary author creating a work in an 'originary moment' of composition" (Greetham 1992, 337). Accordingly, "[m]ethod in textual criticism and editing turned from being indigenously based on a critically established validity of text, to being exogenously predicated on (authorial) authority." (Gabler 2012)

In his lecture *The New Textual Criticism of Shakespeare*, Bowers makes clear this endeavour of authenticating a text based on authorial intention: "We must keep one thing firmly in mind. Ultimately, textual enquiry comes to rest on the authentication of every individual word in every Shakespeare play as — on the available evidence — what we may believe to be the very written word of Shakespeare." (Bowers 1966, 66) The foremost problem — the solution to which gives rise to the term *eclectic* editing — is that no single document, even

the author's manuscript, can be seen as an embodiment of the author's intent. Thus, instead of constructing an (albeit lost) archetype, as in the stemmatic approach, the editor constructs 'eclectically' (from a number of documentary sources) a text that represents the author's intent, but which does not, and never had, existed except as an abstract concept: "What became known as the 'text that never was' (but by implication, *ought* to have been, in the best of all possible worlds, since it constructed authorial intention in despite of the testimony of individual documents)" (Greetham 1994, 334). This eclectic editing has been criticised for its ahistorical approach, producing a 'contaminated' text. (Zeller [quoted in Eggert 2009, 209] gives such a view from the German school of historical-critical editing.)

The often-employed formulation '*final* authorial intent' is intended to allow consideration of texts that were revised by an author, and republished, during his or her lifetime. Such revisions clearly suggest that an author's "intention" towards a work changed. Consequently, it is assumed that the later revisions be treated, within the copy-text theory, as the most authoritative — that is, representing that which the author deemed the most complete version. While, according to Bryant, "the critical edition [can] theoretically accommodate a multiplicity of intentionalities: the early, intermediate and late desires of single writers or collaborative groups", "[i]n practice, the Greg-Bowers approach has generally been used to promote the *final* intentions of an *autonomous* author" (2002, 44). Bryant, however, notes that a broader interpretation of intentionality was proposed, particularly by Tanselle, whilst operating still within the overall approach of Greg-Bowers (Bryant 2002, 44).

The theory of copy-text has much in common with stemmatics. With its even greater emphasis placed on the notion of the work, to the downplaying of document-texts, this approach produces a single (critical) text. Indeed, it is possible to suggest that the printed

edition of a single critical text is the ‘natural’ form for such an edition. Producing a printed book was clearly the intention of authors; as such, a scholarly edition in such a form may be seen as replicating the intended form. Thus created is a set of expectations — from the reader’s side, having the text in the form of a book like any other; from the editor’s, the desire to create a single critical text that represents authorial intent — that closely binds the form of the edition (a critical edition) to the editorial activity (critical editing).

1.2.3. ‘Critique...’

Criticism of the copy-text theory of editing — at least from *within* the Anglo-American tradition of critical editing¹¹ — comes essentially from two perspectives: those of a more philosophical nature, from the New Criticism’s ‘Intentional Fallacy’, through Barthes’s ‘Death of the Author’, to poststructuralism and Foucault’s ‘Author-Function’; and, from a more technical perspective, the actual conditions of the preparation and distribution of an author’s work.

The latter perspective is perhaps most directly addressed by McGann in his *Critique of Modern Textual Criticism* and McKenzie’s ‘sociology of text’. McGann addresses, specifically, the application of this theory to modern texts intended for print. McGann’s argument is not only that the text being reconstructed never existed, but rather with Bowers’ (and later Tanselle’s) argument that an author’s manuscript be chosen as copy-text on the basis of its authority. Implicit in this approach is the view that all amendments in the final printed edition to an author’s intentions are treated as corruptions — as they would be if introduced by a medieval scribe copying a manuscript. This is partially true: compositors typesetting

¹¹ The German tradition has, since the 1950s at least, held rather a different view (see earlier reference to Zeller).

make numerous errors, and editors introduce unauthorised changes, either accidentally, or based on personal preference or external pressures (such as censorship). However, as McGann argues, the process of producing a published work is not solely the responsibility of the author — undermining the primacy of authorial intent. Rather it is the “result of interventions by many agents (such as copyists, printers, publishers) and material processes (such as revision, adaptation, publication)” (Siemens *et al.*, 2012). Authors also relied upon editors to clean up their work and were often accepting of changes. As such, an author’s manuscript may not represent their ‘final’ intent: potentially, the author’s ‘intent’ was in fact for an editor to correct errors, or even adopt his own preferences (for example, a house-style). This becomes more problematic when dealing with texts that underwent multiple authorial revisions pre- and post-publication.¹²

McGann (1991) further elaborates McKenzie’s ‘sociology of the text’, arguing that material aspects of print production — such as typography, layout and binding; ‘bibliographic codes’ (Pierazzo 2015) — contribute to the nature of the text as well as the purely linguistic codes, which hitherto, in general, formed the basis of textual criticism. As such, critical editions established by, amongst others, the Greg–Bowers line of editing, will tend to privilege the linguistic codes — those which can be ‘extracted’ and edited in abstraction — over other bibliographic codes. In this perspective may be seen, if not a document-orientation as such, a conceptualisation of the work as intrinsically bound to its documentary form. While accepting the validity of this understanding within the confines of bibliography,

¹² McGann gives the example of several poems by Byron, which were subject to such a process (McGann 1983, 52-4). Over the course of several years, Byron made so many changes (and reversions) that charting exactly what is ‘authorial intent’ (or even *final* authorial intent) is extremely difficult. Zeller argues that though seemingly belonging to the same work, the process of authorial revision has created two distinct versions that cannot be reconciled, and should be treated as separate texts (example given in McGann 1983, 58).

Shillingsburg argues that a 'sociology of text' produces no basis for editing: if a text's meaning is inescapably connected to its material embodiment, the text cannot genuinely exist outside it (Shillingsburg 2011). By extension, as Shillingsburg points out, a scholarly edition — having for itself material form — can represent only the 'sociological' activity of the editor (and the typesetter and binder, etc., of the edition) rather than those of the author and other (material) contributors to the 'original' (Shillingsburg 2011).

For the moment, I shall return to the other criticism of the Greg-Bowers-Tanselle line of copy-text editing previously mentioned, namely the philosophical interrogation of *intention* in a text. From the perspective of New Criticism, Wimsatt and Beardsley challenged the notion of authorial intent, arguing that it was irrelevant to understanding a text; meaning resides purely in the text itself: "it is detached from the author at birth and goes about the world beyond his power to intend about it or control it" (Wimsatt & Beardsley 1946). Barthes's "Death of the Author" (Barthes 1968) similarly attaches production of meaning to the reader rather than the author, removing the idea of a single meaning or interpretation and rendering meaning fundamentally unstable. Foucault's "What Is an Author?" (1969) proposes the concept of the 'author-function', prescribing the various functions attached to production of a text in a social and historical context, but again without designating an individual 'Author-God' (to use Barthes's term) as the sole fount of meaning.

The direct applicability of these arguments to textual scholarship is debatable. Such theoretical perspectives are more frequently applied to literary criticism (as opposed to textual criticism). That is, they apply in a hermeneutic sense, which does not necessarily preclude the more formal construction of the text of a work — distinct from what it might 'mean'. Nevertheless, by forming part of the intellectual backdrop against which textual scholarship developed in the later part of the twentieth century they can be seen, as Pierazzo suggests, as

a motivating force behind theoretical developments such as McGann's, discussed earlier (Pierazzo 2015). If, to paraphrase the argument, the author can no longer be seen as the sole authority governing the text, the validity of editing with a view to reconstructing the author's intent is itself questionable.

In such theoretical positions as those given by, in particular, McGann and McKenzie, we see what I would suggest is a transition in theory of textual scholarship from *replicating* or reproducing a text towards a kind of *modelling*: that it models the forces that shape the meaning of the text — the 'purely textual', the material and the sociological. The term and concept of modelling is commonly used in Digital Humanities (McCarty is particularly noted in this respect: see McCarty 2004; also Pierazzo 2015) — and, though it seldom found use in textual scholarship before — I think McGann's and McKenzie's theoretical positions can be seen as a tentative move in this direction.

As McGann later observes, development of editorial theory in the 1980s (following the publication of his *Critique* and McKenzie's 'sociology of the text') marked a "rethinking a crucial distinction we traditionally make: the distinction between a text and its context" (McGann 2010). Clearly, it is possible to represent (or re-present) the text of a work or document on another document — at least, the linguistic codes. The context — the material and sociological codes that shape meaning in a text — cannot be reproduced. (This is precisely the argument that Shillingsburg puts forward.) Such 'contextual' forces shaping a text can, however, be modelled: that is, represented in some abstract way.

To say that this modelling-bent coincides with the theoretical approaches of McGann and McKenzie is, perhaps, both an overstatement of the case and a failure to recognise that print

scholarly editions have always, in some ways, modelled the text. They have always assembled ‘paratexts’, or “auxiliary sections of materials around the text” (Gabler 2010), which, in effect, ‘model’ those sections of the text which cannot be reproduced. The critical apparatus models textual versions by mapping variants onto a base text (created according to some paradigm). As such, it is a recognition of the basic limits of reproduction. Other traditional components of a scholarly edition, such as annotations and commentary, “[m]ediat[e] the [...] text’s or work’s meaning to readers” (Gabler 2010) — that is to say, descriptively model the contextual aspects of the documents or work, rather than reproduce them.

Such features are, nevertheless, clever ways of working around the constraints placed on modelling by the print edition. Though, for instance, capable of modelling textual variants, print editions are — at least considerably more often than not — dependent on mapping these variants onto a single base text. Counterexamples are infrequent, and severely limited in scope. McGann discusses the ‘multiple version’ edition of *The Prelude*. He notes that the “facing-page edition” (with alternative versions on facing pages), while useful, could not cope with texts existing in more than two versions, for the obvious reason that no more than two pages can ever be facing (McGann 1983, 32). This may be seen as an example of a desire to more closely model the complexity of a textual scenario hampered by the material constraints of print. As Greetham, writing in 1992, observes, “the revisionists [McGann et al.] have not yet produced a critical vehicle for representing [their approach] in a scholarly edition.” (Greetham 1992, 354).

Price, likewise, gives the example of neither *The Collected Writings of Walt Whitman* nor the earlier *Complete Writings* “bother[ing] to include” the first publication of *Leaves of Grass*: “How do we explain this omission? To a large extent, this odd result stems from twentieth-

century editorial practices for establishing authoritative or definitive texts that encouraged the selection of a single text.” (Price 2009)

Given the above, it is possible to suggest that one driving factor may be the constraints of the print edition to provide anything other than a single text — that is, the technical capacity of print to enable modelling, rather than re-production, of a text. At the same time, such a form can be also attributed to the editorial theories being employed (particularly in Anglo-American editing, dominated by Greg-Bowers). Of course, as my attempt to highlight the competing forces that govern the form and production of scholarly editions suggests, it would be wrong — or, at least, illusory — to attach primacy to either of these factors. Rather, the single-text scholarly edition exists as one possible stable solution to a series of reciprocal functions: texts must be single, due to the medium; editors edit single texts according to a theory that endorses single texts, which comes about as an attempt to find theoretical grounding for a single text; audiences anticipate single texts because that is what is typically produced; and, that things ‘intended’ (either from the author or the audience’s perspective) to be ‘books’ should be reproduced in book form.¹³ Approached in this way, we can see the disruption to this stability caused by the theoretical developments outlined above, but no new stability emerging; this, I would suggest, is primarily due to the incapacity of print to adequately model more complex textual theory.

Only with the emergence of the digital medium have new points of stability begun to emerge: as Finneran notes, there is a coincidence between the development of new (digital) technology and “fundamental shift in textual theory, away from the notion of a single-text

¹³ Robinson (2010) highlights this point, albeit with reference to digital editions, by suggesting that by not producing a single, readable text, editors are alienating audiences.

‘definitive edition’’. (Finneran 1996) Before moving to consider the impact of the digital medium on scholarly editing, and particularly its effects on the Anglo-American school, it is necessary first to look at alternative theories and schools of scholarly editing in their pre-digital manifestations.

1.2.4. Documentary editing

The Anglo-American school of editing typically distinguishes critical editing, which I have thus far described, and so-called non-critical editing. While critical editing “establishes” a text, non-critical editing “seeks to present a faithful version of a single document, with only minimal textual involvement by the editor” (Greetham 1992, 349). Non-critical editing, then, can be considered editing without critically establishing a single text from a number of versions, or, as Greetham puts it, “simply reproduce[s] a text already in existence” (Greetham 1992, 347). Of course, Greetham’s “minimal textual involvement” needs qualifying. As Dahlström argues, the idea that an edition can totally reproduce a text or document is a form of “mimetic fallacy” (Dahlström 2004, 8): that any text ‘reproduced’ from a document requires some ‘ideal’ conception of that text.

According to Greetham, “since by definition [non-critical editions] do not involve any criticism of the text, their [their] characteristics have not been in the forefront of textual theory, as have the features of critical editions” (Greetham 1992, 351) Certainly, their characteristics not being “at the forefront of textual theory” is true. However, as suggested above, “not involving criticism of the text” is clearly spoken with a particular notion of “criticism” in mind. As Kline argues, non-critical editions are hardly “uncritical”. (Kline 1995, 2) (Dahlström’s *mimetic fallacy* demonstrates just this point; since a text cannot be reproduced wholesale, any interaction must be a critical engagement.) As Pierazzo notes, the

description of such editions as ‘non-critical’ may be seen as deliberately downplaying act of editing — i.e. the editor plays no ‘critical’ part — and thus be regarded as pejorative. (Pierazzo 2014)

Consequently, both Pierazzo and Kline employ the term *documentary editing* — an editing with an orientation towards the document as a historical artefact, rather than the material support of an abstract text. Cappon (1966) instead uses the term ‘historical editing’. *Diplomatic edition* is also (more traditionally) employed, to designate an edition that reproduces the document (the “diploma”) (Pierazzo 2011). Such a term also fits with Tanselle’s distinction, presented in *A Rationale of Textual Criticism*, that critical editing edits the text of a work, whereas non-critical (or documentary) editing edits the text of a document (Tanselle 1989). I shall refer hereafter to documentary editing, as it is more broadly employed outside the Anglo-American sphere (where ‘critical’ — and, by extension, ‘non-critical’ — carry a more specific interpretation; see Gabler 1995; Sahle 2014).

Pierazzo distinguishes between several types of documentary editions in pre-digital editions. Diplomatic editions, particularly, in addition to reproducing the text in an abstract form, attempt to recreate some of the graphical features of the text-as-embodied, including “line breaks, page breaks, abbreviations, and differentiated letter forms” (Pierazzo 2011). Further detail may be included in so-called ‘super diplomatic’ or ‘type-facsimile’ editions, which aim to even further reproduce the physical features of the document (including “functional signs or marks”), as far as these may be reproduced in print (Pierazzo 2011). This particular regard highlights the representational constraints of print — at least typically — to that which can be typeset. Diplomatic, and especially super-diplomatic, editions can clearly be seen as an imperfect response to the theoretical desire to encode material aspects of a document, in spite

of limits of the edition's medium. These editions, with their habitual typographical complexity and disruption of the flow of text, are both difficult to read and can only give a limited sense of the text-as-situated. Tanselle, it may be noted, argues that photographic reproductions of documents should be included, though acknowledges that this may be impossible, not least due to cost (Tanselle 1989). Pierazzo also argues that transcription, the act of re-inscribing the text of a document, should not be conflated with a documentary edition: the transcription is an editorial work-in-progress, while the edition itself is a preparation of said text for public consumption (Pierazzo 2011).

Documentary editing can be seen as editorial practice rooted in historical, rather than literary, studies, to provide a text for historians to use as evidence. (Kline 1995, 11) Typical examples of such texts these include private correspondence and diaries, which would typically not be typeset or reproduced. Tanselle, likewise, suggests that the main focus of documentary edition is private text, that is, texts not intended for publication (Tanselle 1989). In this way, documentary editing is associated more with historical scholars rather than with literary (Kline 1998). According to Shillingsburg, the “documentary (or historical) orientation” is “founded on a sense of the textual integrity of historical moments and physical forms” (Shillingsburg 1996, 17). Its value is more in “stressing the evidentiary value of those documents”, whereas public writing “would seem to demand editions that attempt, through editorial emendation, to offer the texts of works, the works that are but imperfectly represented by the texts of the documents” (Tanselle 1989, 64). Kline, likewise, emphasises the importance of the document itself as historical evidence, thus “demand[ing] that their editors provide editorial texts that themselves will communicate as much of the sources’ evidentiary value as possible.” (2014). Harvey, in *Editing Historical Records*, is less prescriptive as to the act of editing, insisting only on transparency of the edition regarding normalisation

of the text (of his three ‘maxims’, the second, “Say what you are going to do and do it”, best reflects this view) (Harvey 2001, 13).

Further motivations include the utilisation of documentary editions “as substitutes for the original, as long as one understands their limitations.” (Tanselle 1989, 58). Another is that it allows the reader to “take advantage of the editor’s expertise”, particularly regarding difficulties such as handwriting (Tanselle 1989, 59). Greetham, again pointing to the addition of editorial expertise, notes that such versions may also serve as a basis for critical commentary (1992, 347).

In the United States (the perspective from which Kline writes), documentary editing developed in parallel to the more dominant strains of critical editing. Indeed, as Kline notes, the term was seldom used before the 1970s, even though before this a distinctive methodology had begun to evolve — as early as the nineteenth century, in fact, when editions of statemen’s papers and correspondence were produced. (Kline 1995) Kline identifies the criticism of Tanselle (himself a literature-oriented critical editor, and proponent of the Greg-Bowers line of critical editing) of the practices of documentary editing as a key moment — his principal point of attack being historical editors’ tendency to not sufficiently respect the document as historical source, by, among other things, silently amending the document text. (Kline 1995; Simon 1981) According to Tanselle, a documentary edition should, unlike a critical edition (which seeks to correct error-laden texts) replicate the text of the document as it stands. The rationale for such an approach is that, as given above, a documentary edition should replicate the “imperfectly represented” text (Tanselle 1989). Thus an editor should not attempt to “smooth” the text by correcting or modernising spelling or punctuation; to do otherwise is to misrepresent the document text (Tanselle 1989). Such criticism provoked

something of a reappraisal, but also a further acceptance of documentary editing, including by the MLA Committee on Scholarly Editing.

Thus far, we have focused on one of two possible motivations for taking a documentary approach to editing: namely, because a critical approach is impossible — either because there is literally only one version, or (more typically) because a historically-inclined editor views the document as inherently unique, and thus singular. However, there may be an alternative motivation for producing a document-centric edition: because critical editing it is undesirable. This applies especially to editing the manuscripts of a literary work, even when multiple versions (including printed copies) exist. A foremost rationale, as Pierazzo argues, is a rejection of ‘critical’ editing of the Greg-Bowers school. (Pierazzo 2014) The theory of ‘material philology’, likewise, posits the inseparability of the literary work and its material embodiment, and would seem to endorse this approach. (The same, indeed, might be argued of McGann’s and McKenzie’s perspective.)

Such views represent the meeting — full-circle, as it were — of the literary and historical conceptions within Anglo-American editing. In other schools and editorial traditions, however, the division between historical and literary editing were never as clear-cut.

1.2.5. French and German perspectives

The above sections have dealt primarily with traditions of editing as they have developed in the anglophone sphere. That they appear ‘first’ in this thesis — and that the perspectives presented here are treated more summarily — is a reflection of the nature, predominately, of the literature available in English (and, significantly, the relative unawareness, in Anglo-

American schools of editing of other traditions).¹⁴ I do not wish to imply any pre-eminence of Anglo-American schools.

The differences between the Anglo-American and French and German schools of editing (themselves hardly reducible to a single ‘school’) can be seen as illustrating the ‘rhizomatic’ nature of editorial schools. It can hardly be supposed that there is something peculiar about a text written in French or German that mandates a particular form of editing, or that precludes their editing along the lines of the Greg-Bowers *copy-text* school. Likewise, all these schools have their origins in the earlier history of textual scholarship (e.g. Lachmann). Rather, other factors have come into play: the development and refinement of earlier methods, and, importantly, institutionalisation of different editorial methods have entrenched their pre-eminence in particular countries. The (almost exclusive) preference given to Greg-Bowers by the MLA (see above) is indicative of this. Another factor is separation — afforded by, amongst other things, the language used for academic literature. Where two schools are accessible to each other, cross-fertilisation is more likely to occur; when they are not, speciation arises.¹⁵ It remains to be seen whether increased awareness, and the sweeping along of all previous schools by the advent of the digital, will further efface these distinctions. As with all things rhizomatic, we are more likely to see the development of new and interesting hybrids. For now, let us turn to types of edition that are particular to various countries.

¹⁴ That Eggert (2009) refers to the (belated) awareness of Anglo-American editors with, in particular, German editorial methods as ‘The German Encounter’ illustrates the nature of the paucity of literature in English related to ‘nationally-determined’ schools of editing.

¹⁵ It may also be observed that degrees of ‘isolation’ and ‘cross-fertilisation’ may differ on each side: Eggert notes the relative familiarity of German editors with the editorial methods of the anglosphere, while the reverse was much less prevalent (hence, ‘Encounter’...)

The *édition génétique* (genetic edition) — and genetic criticism (*la critique génétique*) — arose during the 1970s to occupy “une place nouvelle dans la recherche littéraire française”¹⁶ (Gresillon 1994). It can perhaps be best characterised as sitting at the intersection of textual scholarship and literary studies. Its aim — quite in contrast to, in particular, Anglo-American editing practices, but also to a lesser degree the German school — is the reconstruction of the *genesis* of a literary work. It focuses on the work¹⁷ (as a process) of the author (“le travail de l’écrivain”: de Biasi 1996, 161), opposing itself as Grésillon states to the idea of the work as ‘closed’. (Gresillon 1994) In this view, versions of a text comprise what Deppman describes as “meaningful variant” — in contrast to “errors or corruptions” (Deppman et al. 2004, 2), that is, aspects to be amended in the pursuit of a ‘correct’ (or final) text.

Accordingly, it can be seen as taking a document-orientation towards literary manuscripts (to a very specific degree). A genetic edition “présente ... l’ensemble des documents génétiques conservés d’un oeuvre ou d’un projet.”¹⁸ (Grésillon 1994, 243) Its principal focus is on authorial manuscripts and other notes (so-called *avant-texte* material). It aims to not only represent the text of documents but to present an analysis of the artistic creation of the work through studying documentary features such as the author’s revisions and amendments as encoded in the material evidence. It is thus, rather than a representation of the document’s text spatially-situated (as a type-facsimile may be), a temporal representation of the act of creation of a work.

¹⁶ “a new place in French literary research” [my translation]

¹⁷ Work in the sense of the French *travail*, rather than *oeuvre*: the actual work that is done, not the literary work as an entity. (As Eggert, 2009, observes, English is one of the few languages that does not make this distinction.)

¹⁸ “which presents ... all of the preserved genetic documents of a work or a project” [my translation]

As such, despite the main focus on the document, it is also necessarily concerned with the work whose genesis it purports to represent. As Bowman suggests, such a view may imply a kind of teleological bent (Bowman 1990, 628). This is not to suggest — quite the contrary — that genetic criticism sees a trajectory towards a single, ‘definitive’ work; however, it is also undeniable that it maps a trajectory (indeed, multiple trajectories) towards the creation of an entity (the work, even in its multiple states). It is not a model of pure textual fluidity.¹⁹

The chief difficulty for genetic editions is the impossibility of (literally) representing time as inscription on a two-dimensional surface. As such, the *édition génétique* exposes quite clearly the duality of the scholarly edition: the text of the edition — that is to say, the writing itself — is a representation of the text of the document, while the edition simultaneously *models* the creation of a work. Indeed, as Pierazzo observes, the modelling aspect of genetic editions is sufficient that the text of the edition itself is difficult to read. The chief cause of this may be seen as the constraint of print to convey the three dimensional page and the fourth, temporal dimension on what is effectively a two dimensional surface;²⁰ instead, abstract representations of time — through marking phases of writing with special symbols or diacritics — must be employed instead.

In the Germanic tradition of editing, the *Historisch-kritische Ausgabe* (historical-critical edition) tends towards the presentation of the textual history of a work, including all its

¹⁹ To even say, ‘we are creating a genetic edition *of* something’ suggests one or more objects of study, to which we may tentatively assign the label ‘work’. It is possible that genetic criticism (especially in the digital realm) may come close, in certain textual conditions, to kind of pure ‘cartography’ (as Deleuze and Guattari put it: Deleuze and Guattari 1987, 17) — that is, texts used and reused in a circulating system with no overall direction at all.

²⁰ A book is, of course, three dimensional; but the third dimension — carved into or protruding from the page — cannot feasibly be used to encode information.

variants, in place of establishing one single text. According to Gabler, the aim is not to “critically shape” a text, but rather to “provide a segment of the text’s history” with the history of the text provided in the form of apparatus (Gabler 1995). Though, as Gabler points out, both the Anglo-American and German schools of editing may be regarded as both *historical* and *critical*, the Anglo-American school emphasises the critical aspect, whereas German tradition tends to focus on the historic. According to Zeller:

Only one specific version, one specific witness document which, as the carrier of a particular state of the work, represents a unique historical moment or a discrete stage in the development of the work and the author, forms the basis of the edited text. (Zeller 1995, 27)

This marks a great difference from the intention-based, eclectic editing of, especially, Greg-Bowers. As Gabler observes, though both Anglo-American and German editors edit *texts*, the difference is between editing texts in their historical manifestation — that is to say, the materiality of their embodiment — as opposed to an abstract, ideal nature of authorial intent (Gabler 1995). Zeller notes that spelling and punctuation (*accidentals*, in Greg’s theory) should not be modified as they represent a text in its given material condition; to alter the text of a particular document with reference to another document (regardless of whether one would deem a particular version to be authoritative) results in a “contaminated text” (Zeller 1995). (Thus, to invert the earlier reference to Greetham: the text that never was *really* never was, and thus really *should not be*.) Though not explicitly document-orientated in the same way as diplomatic editions, with their focus on the material embodiment, the German *Historisch-kritische Ausgabe* takes a much less idealistic and abstract notion of the text and the work. It can, perhaps, be said to be version-orientated: the work is not an ideal conceived only in the author’s mind, but never realised, but is represented in its development (from notes and manuscripts to edited proofs and multiple printed editions) by the totality of the

documents in which it is embodied. Changes made by an author post-publication, likewise, constitute a different historical text, a version to be represented as one among many, rather than an altering of authorial intent; Gabler succinctly characterises this as an endeavour to record “what the author *did*” as opposed to “what the author intended” (Gabler 1995).

The editor’s objective, then, is not to correct a text, with emendation strictly limited (Gabler 1995). Instead, a particular version may be selected as the “orientation-text”, a version that represents the structure of the work, around which it is possible to organise the variants embodied in the documentary history of the text (Gabler 1995). Even this, as Gabler observes, is seen as an editorial choice, and one that is often eschewed to avoid giving precedence to the version chosen; the edition may be “apparatus-only” (Gabler 1995). Zeller, also, notes the de-centring of an edition, away from any base-text, instead consisting purely of textual variant (recorded as apparatus):

In several editions the apparatus has developed into an autonomous part of the edition, equal in standing to the edited text, and at times indeed incorporating it. (Zeller 1995, 96)

Such a theoretical basis can be seen as presaging the criticisms of intentionalist editing by McGann and others; Gabler, indeed, notes that, of such views, the “salient points have already been anticipated by German textual scholarship” (Gabler 1995). Such a time-lag may be explained by a number of factors. As Gabler notes, citing Shillingsburg, the dominance of the Greg-Bowers school in Anglo-American editing provided such a critical consensus that any departure was, at the least, difficult (Gabler 1995; Shillingsburg 1986, 1996). Foremost also must be the relative inaccessibility to an English-speaking audience of German editorial theory. In his otherwise comprehensive chapter on ‘Textual Criticism’ in *Textual Scholarship* (1994), Greetham does not refer to any German theoretical position. Indeed, the first real

attempt to introduce an English-speaking audience to such theories was the collection, *Contemporary German Editorial*, edited by Gabler and published in 1995 (from which the majority of the above references are drawn).

A perhaps greater argument still is the apparent change in the function of an edition that adopting such editorial procedures might entail. As was earlier noted, producing an edition containing multiple concurrent versions of a text in print was extremely challenging. (Greetham noted that a “critical vehicle” had not yet been invented to contain such views; Zeller’s ‘facing-page’ edition could not, logically, represent more than two versions of a text on facing pages.) Undoubtedly such representation is possible in print, though not without resorting to complex apparatus; consequently, as Gabler observes, “[i]t is aimed less at the reader than at the user of the edition” (Gabler 1995). Anglo-American editing can be seen as an extension of Bowers’ aim of providing an accurate and authoritative text for literary criticism (Bowers 1959); even when critical editions laden with apparatus were published, the possibility (even the necessity) of, at the same time, producing a reading edition with a single text representing the work, demanded a single critical text. To say that German historical-critical editions, by contrast, demand an understanding of *textual* scholarship might well be true, but this misses the essential point: that the use-case for such an edition is in understanding the textual history of a work. As such, what appears to be, from an Anglocentric perspective, a necessity that the text of the work be readable *as a text* does not apply: the societal (and hence economic) pressure that a scholarly edition present, as its central focus, a single version of the work is not a force that, in this case, determines the editorial model.

This distinction between presenting a work for the same function as the work itself (i.e. for the text to be read) with editorial activity subservient, and modelling the activity of textual scholarship as a fundamental part of the edition presages many of the theoretical debates and tensions with regard to digital scholarly editions. Essentially, is the purpose of a digital scholarly edition to be read (or to facilitate the study of its textual content alone) or to open up the mechanics underlying the creation and genesis of the text — or can it do both, without the context drowning out the text?

1.3. Digital Scholarly Editions

1.3.1. The digital turn

It is difficult to determine the extent to which the advent of the digital as a medium for scholarly editions has destabilised previously stable sets of practices. From one side, undoubtedly this shift in medium has led to the reappraisal of theories of textuality, arising (primarily) from the heightened representational capacity of the digital. However, the dominant theoretical tendencies in textual scholarship (especially the Greg-Bowers school in Anglo-American editing) had already been destabilised to a large degree. Accordingly, Finneran's observation that the advent of a new medium coincided with a "fundamental shift in textual theory, away from the notion of a single-text 'definitive edition'" (Finneran 1996) seems most apposite. The digital turn afforded new potentials for representing the results of new editorial theories that has previously only been articulated as a critique of existing practices.²¹

²¹ That McGann's *Critique* is explicitly a critique, and not a manifesto for a new school of editing, is, I think, illustrative of this point. Indeed, Greetham, writing in 1992, observed

That said, digital scholarly editing cannot be seen as having crystallised into a single coherent set of practices. As Pierazzo (2015) observes, attempts to theorise digital scholarly editions and digital scholarly editing have been notably lacking. Likewise Robinson, who argues that developments in digital scholarly editions since the early 1990s have been practice-led: an exploration of the potentials (and constraints) of the digital medium. (Robinson 2013) Sahle’s notion of a digital scholarly edition — an edition whose functionality is reliant on the digital medium, not merely a *digitised* edition (Sahle 2014) — may be seen not so much as theory of digital scholarly editions, but a definition (and one ‘defined’ implicitly against print editions). As such, it is not surprising that Pierazzo (writing as recently as 2015) is still able to pose the question: “Is digital scholarly editing the same as scholarly editing but in a new medium, a new methodology or a new discipline altogether?” (Pierazzo 2015, 9) This question, as well as recent theoretical developments, will be discussed in the next section. This section will concentrate on development of the digital as a medium for scholarly editions.

Identifying a starting point for digital editions is, in itself, difficult. With a due degree of hindsight, it seems reasonable to argue that the pervasiveness of Internet connections, the advent of the World Wide Web as a viable delivery system, and HTML rendered by a web browser as a mode of representation, marks a point from which it is pragmatic to talk of digital editions.²² Moreover, when we talk today of digital scholarly editions, what is almost

that, for all the validity of such a critique, it lacked a “critical vehicle” for its instantiation (Greetham 1992, 354) — that is, there was a theoretical destabilisation without the possibility of restabilisation.

²² This is not to discount, for instance, Robinson’s *Wife of Bath’s Prologue on CD-ROM* (1996) or other CD-ROM-based scholarly editions; however, it cannot be denied that the distribution of scholarly editions on physical disks was — again with hindsight — something

invariably meant is a web-based edition (this is not, of course, to suggest that a digital edition must be web-based; rather, it is a *de facto* reflection of reality). However, the viability of the digital as a medium for scholarly editions lagged several decades behind the use of the digital as a tool for textual research, including in the preparation of print editions. (Vanhoutte 2013)

Father Roberto Busa's *Index Thomisticus* is often cited as the first Digital Humanities project. (See Hockey 2004) Using IBM computers, Busa constructed a lemmatised concordance of the complete writings of Thomas Aquinas. Though much of the work was automated by computer processing, the end result — the actual concordances — was published originally in print form, and only much later on CD-ROM and then via the web. (Hockey 2004) Busa's work in this regard can be seen as the beginning of the digital medium as a tool for processing and carrying out statistical analysis of text — in other words, for studying text, rather than for its representation and dissemination. For scholarly editors, computers were also used as a “research assistant” (Shillingsburg 1980, 31). Tools were developed to assist in, and in part automate, “concordancing, collation, analysis of variants, stemma determination, reconstruction, building and photo composition.” (Vanhoutte 2013) Faulhaber gave a similar view, citing the computer as a “facilitator” of textual criticism. (1991, 123) Gabler's use of the *TuStep* collation tool (in preparing his edition of *Ulysses*) and Robinson's *Collate* software stand as prominent examples of digital tools employed in the preparation of print

of a stop-gap. At the same time, it should be observed that these editions were *truly* digital editions: paralleling Sahle's definition, Robinson describes “electronic editions” as “[E]ditions conceived and executed exclusively for electronic publication, and impossible in any other form.” (Robinson 2009, 51) This stands in marked contrast to other simultaneous uses of CD-ROMs as accompanying print editions with supplementary materials (see Kline 1995, 78). Accounting for the ‘demise’ of CD-ROMs is beyond the scope of this thesis — though the platform-independence of the web and (comparative) cost of the disks suggest themselves as potential reasons (even before laptops stopped coming with CD or DVD drives).

editions (Hockey 2000, 126). From a more organisational perspective, Kline describes the utility of “computerized databases” for maintaining “control files”, a catalogue and index of documents being selected for an edition. (Kline 1995, 63) And, of course, there was the use of more generic digital tools, such as word-processors, by textual scholars.²³

In these cases, the use of the digital as a medium for presenting and reading texts was far from a consideration: access to computers was limited, due to their number, relative complexity and cost, and their employment for computational (rather than representational) work. This is implicit in Faulhaber’s declaration that, in carrying out computational work on texts, such as the preparation of concordances, “[a]s a byproduct, but only as a byproduct, the computer produces an electronic version of the text.” (Faulhaber 1991, 123) At the same time, he sees dissemination as “the next logical step”; “to deliver [the] electronic version as the end product itself.” (Faulhaber 1991, 123) From the 1980s onwards, much formative thought was given to storing texts in database form, and have them accessible from a computer terminal connected to a network. Ross considered the textual database as “the ‘front-end’ for a document retrieval system ... the kind of information in the database could be displayed to a scholar working at a terminal, where passages from all sources could be called up.” (Ross 1981, 161) While seeming primitive by today’s standards, and operating primarily on local networks, such a system mirrors a Web-based edition in all essential functionality.²⁴

At the same time, scholars were realising the potential of the digital medium for more dynamic representations of text. Laufer in 1989 proposed an ‘*édition-diffusion électronique*’, using “multiple frames on the screen as a manipulation tool for the dynamic reading or

²³ Shillingsburg (2013) also claims to have regularly employed the Donald Knuth’s *T_EX* system for the typesetting of his editions.

²⁴ Indeed, the ‘primitiveness’ is really only in the reach of the network.

consultation of parallelized documents (version, variants, facsimiles, annotations...)” (Cited in Vanhoutte 2013) Logan *et al.* argue for the superiority of such systems over print editions for accessing and representing complex textual scenarios: “[T]he windows of the electronic edition can replace not one but several place-holding fingers [...] they allow appropriate temporary rearrangements of the pages of the edition.” (Logan *et al.* 1986, 322) According to Lebrave, a database containing a number of versions of a text, presented digitally, would “allow any reader to construct his own reading according to the hypothesis he wants to build up.” (Lebrave 1987, 42) In these early conceptions of digital editions, we see a tendency towards, if not developing an explicit editorial theory, then exposing the medium-bound exigencies of extant editorial schools. As McGann argues, older theories were, in part, founded — or, rather, stabilised around — the constraints of the medium rather than a purely textual rationale:

The exigencies of the book form forces editorial scholars to develop fixed points of relation — the “definitive text”, “copy-text”, “ideal text”, “Ur text”, “standard text”, and so forth — in order to conduct a book-bound navigation (by coded forms) through large bodies of documentary materials. (McGann 1991)

In this, we see the “coincidence” between digital developments and editorial thought alluded to by Finneran (see above). It is, perhaps, most interesting that these theoretical positions began to be articulated before it became practical (and, certainly, before it became commonplace) to produce digital editions. In this regard, the potential itself of the digital medium, quite apart from any actual employment of it, can be seen as a further destabilising force on top of theoretical critiques — precisely, it should be noted, because it does not propose any restabilisation in actuality.

It is also true that such developments were taking place against a background not only of the rise of poststructuralist literary theory, but of the supposed technological instantiation of such theories. Landow's *Hypertext* (1992) is, perhaps, the most obvious case-in-point (witness its subtitle: *The Convergence of Contemporary Critical Theory and Technology*.) From the perspective of scholarly editing, however, it is not so much about the hypertext systems' "promises to embody", as Landow (1992, 3) puts it, a particular theory. Indeed, as Vanhoutte argues, such discussions of hypertext within scholarly editing are prone to hyperbole: many such editions employ hyperlinks, but only within their own local, closed system. (Vanhoutte 2013) The reason for this, I would suggest, is the fact that what is being represented is not, in itself, 'hypertext' — at least in the same way as Aarseth's notion of *ergodic literature* (Aarseth 1997) — but a digital representation of sources that are evidently non-hypertextual. As Schreibman notes, "the pattern of centres and margins at any point within the hypertext may be in flux, but the possibility of any one node becoming a centre is established by the writer or editor." (Schreibman 2002, 78) The digital medium allows a plurality of texts to be conveniently co-presented (the convenience of not requiring Logan *et al.*'s "several place-holding fingers"), and allow much closer integration of apparatus. But this is, I would suggest, a condition of a digital hypertext *system*, rather than an editorial theory that argues that historical text(s) can be treated *as* hypertext in Landow's conception. As such, Siemens' view seems more pragmatic and appropriate: "the computer facilitates a non-linear interaction with the text [...] facilitat[ing] a reader's interaction with the apparatus." (Siemens 2012, 447) — 'facilitates' being the important word here (in its proper etymological sense: to make *easier*, not to *enable*). Not, of course, that the effects of this facility should be underestimated.

This representation of text via a digital system (in particular the Web) was, nevertheless, only really idealised as a possibility: as a “byproduct,” according to Faulhaber, of computer processing of texts and text databases. Vanhoutte, in particular, observes a particular friction between the computer-readable and –processable (the database) and the presentational (in the form of hypertext) in many of these early discussions. Thaller, in particular, distinguished between text databases, which are “able to support historical *research*”, and hypertext editions, which “administer, in a convenient way, *results* of historical research.” (Thaller 1996) Burnard, likewise, argues that the hypertext representation of a text gives a “view of information as an emergent property resulting from connection between one piece of discourse and another”, whereas a database system requires a greater “formalization” of the text (1992, 17). Vanhoutte has suggested that the “object”-like representation of text-as-hypertext tallied more closely with an editor’s notion of an edition. (Vanhoutte 2010) An HTML file, for instance, may be seen as a (relatively stable) representation of a document (or text), while a database involves the breaking-apart and arbitrary formalisation of texts, and arbitrary querying and recombination in a way that has little to do with the original text.²⁵ However, the arbitrariness of this distinction was beginning to become apparent, Robinson for instance noting that a collection of hypertext documents with cross-referencing links was essentially a database. (Robinson 2004)

It is difficult, for all the above, to determine whether the digital medium caused any particular paradigm shift in editorial theory. It is certainly true that the digital did not supplant — nor

²⁵ Such a view is later echoed in articles such as that by Bradley, which discusses the utilisation of relational databases for extracting and querying prosopographical information from texts. (Bradley 2005) Evidently, even with such a database made available online, the representation of text is placed second (actually, it is impossible) to the study of particular aspects of those texts.

has it fully supplanted — print as a medium for scholarly editions. As noted earlier, I think this stage of development can best be seen as adding fresh impetus to a theoretical rupture already taking place. That said, several early digital editions can be seen as attempts to stabilise such new paradigms, combining the affordances of the new medium with theoretical ideas within textual scholarship. According to Vanhoutte, McGann's *Rossetti Hypermedia Archive* was “created ... to demonstrate the practical feasibility of his social theory of the text that was at the heart of *A Critique of Modern Textual Criticism*.” (Vanhoutte 2013) Shillingsburg's later notion of a digital edition as a “knowledge site” (Shillingsburg 2006), and Gabler's (similar) conception of the “work site” (Gabler 2010), may be seen in a similar vein. However, in the main, early digital editions may be seen as representing what (at least with hindsight) might be seen as under-theorisation of the medium.

1.3.2. Scale and heightened representation of the digital

Aside from the nature of hypertext discussed above, two characteristics of the digital medium (and in particular Web-based editions) led to what may with hindsight be seen as a degree of hyperbole. These are the multimedia capacity of the digital — in particular, the ability to cheaply include page facsimiles (Schreibman 2013) and its apparent boundlessness in terms of scale. These elements — especially taken together — appeared to enable the lessening of the subjective interpretation of the editor in favour of greater objectivity (giving rise to so-called ‘un-editing’); or the partial transferral of critical activity to the reader of the edition. In terms of scale, Price (2004) notes that “[o]ne distinguishing feature of electronic editions is their capaciousness: scholars are no longer limited by what they can fit on a page.” (His own work on *The Walt Whitman Archive* is an example of this positive accumulation of material.) McGann, likewise, emphasised the possibility of increased scale:

The computerized edition can store vastly greater quantities of documentary materials, and it can be built to organize, access and analyse those materials not only more quickly and easily, but at depths no paper-based edition could hope to achieve. (McGann 2007)

This accumulation of materials may be seen, in part, as the rejection of author-focused critical editing, which has already been discussed. The ability to include and represent multiple versions of a text may be looked upon as the antidote to the requirement for a single, critical text (Finneran 1996, makes this point). Robinson, however, writing in 2013, highlights the somewhat hubristic nature of earlier digital scholarship:

We could include everything; we could enable every way of using everything we included. There were no limitations beyond our imaginations: resources alone bounded what we could and could not do. There was no need for theory. Practical possibility alone was our guide. (Robinson 2013, 106)

It is, I think, no coincidence that such early editions adopted the ‘digital archive’ moniker; they were seen as the construction of digital databases of material, rather than the explicit editing of text (especially from within the Anglo-American tradition of critical editing). The surface effect of objectivity in a large, multi-version archive may be the deferral of criticism to the reader, and the end of the editor-as-authority. Especially when the edition also includes page facsimiles, the reader has access to the same material (in every sense except the physical) as a critical editor. Shillingsburg, using Eggert’s term, characterises this as the ‘archival impulse’ (versus the ‘critical impulse’) within scholarly editing (Shillingsburg 2013, 3).

In a similar way, the possibility of including digital facsimiles of documents led to similar notions of a new ‘objectivity’. Inclusion of page images in print editions was, of course,

already possible. Indeed, Tanselle advocated the inclusion of “photographic reproductions” where possible in document-oriented editions. (Tanselle 1992, 60) That this could be done in virtually all digital editions, especially without the constraint of printing costs, can be regarded as a significant shift. Moreover, the dynamic nature of the medium enables a reader to zoom into a particular region of a document, heightening its use — in some ways even beyond that of the original document. Some scholars have taken this possibility to suggest a greater degree of objectivity is possible thanks to the inclusion of facsimiles. Kiernan’s argument that a page facsimile “subsume[s] the diplomatic edition” (2006, 266) appears to fall into this category. Other theorists, however, have questioned this apparent objectivity. Dahlström, for instance, observes that a digital facsimile of a document is itself a subjective representation. (Dahlström 2004)

The result of these possibilities led to a greater focus on the document within digital editions — albeit an under-theorised focus. There existed a broad notion that providing, on one hand, all the textual evidence, and, on the other, providing it in an ‘unadulterated’ fashion, could replace the work of textual scholars. Schreibman perhaps best outlines this tendency:

These first-generation digital editions did not appear, at least on the surface, as scholarly or critical in the ways that the textual editing community had come to understand the terms for print publication. Frequently these early editions did not include the apparatus, critical notes, and contextual material that, by the late 1990s, had come to signal the apotheosis of scholarly editing. (Schreibman 2013)

More recent theoretical considerations have (as with Dahlström and Pierazzo and Sutherland above) brought a greater understanding of the limits of the objectivity. Price notes that the

technical capacity to do away with a single critical text does not usher in a new objectivity, but an alternative subjectivity:

Even an inclusive digital archive is both an amassing of material and a shaping of it. That shaping takes place in a variety of places, including in the annotations, introductions, interface, and navigation. Editors [...] are and should be, after all, something more than blind, values-neutral deliverers of goods. (Price 2004)

Taking this further, it may be argued that editors cannot be “values-neutral deliverers of goods”, even when they may wish to be. Dahlström makes this point most forcefully, asking (as per the title of his article), *How Reproductive is a Scholarly Edition?* He suggests that any remediating act — including photographing or transcribing (as opposed to editing in some way) — is a critical act that fundamentally alters the objects being studied. (Dahlström 2004) Pierazzo argues — against suggestions by Huitfeldt and Robinson to the contrary (Huitfeldt 2006; Robinson 2009) — that transcription can never be an objective act. (Pierazzo 2011) As such, any such attempt to push critical activity towards the user is not value-neutral; subsequent criticism carried out by the user is a criticism-upon-a-criticism. The danger, as Price observes (above), is that we pretend it is otherwise. In deferring criticism (in the sense of providing versions, rather than a single, critical text) to the reader, we are actually deferring possible criticisms — determined by the editor — and going some way to shaping them.

1.3.3. The digital encoding of texts

As has already been seen, the earlier stages of utilising the digital medium for text analysis and representation were marked by a division between databases — particular the relational database paradigm (see Ramsay 2004) — and text encoded *as* text for representation. (Knuth’s *TeX* system was designed for the typesetting of documents; HTML as a mark-up

language for rendering by a web browser. Both of these examples treat mark-up as presentational.²⁶) This ‘encoding for analysis’ and ‘encoding for presentation’ divide can, I think, be seen as intrinsically linked to the presentational orientation of the above (and other) text encoding systems. Scholarly work on a text requires a degree of machine-readability — thus ‘descriptive’ encoding. Hockey (2004) notes resistance in the scholarly editing community to HTML, suggesting lack of expressiveness for description. According to the project history of the *Walt Whitman Archive*, the original intention was to encode text as HTML, though XML was later adopted due to just this lack of expressiveness.

The Text Encoding Initiative was originally established in 1987 to provide a standardised descriptive encoding format for texts, for re-use, interchangeability and platform independence. It initially used SGML (Standard General Markup Language) before migrating to XML. (Hockey 2004, Schreibman 2013) Such technologies afford the possibility of encoding text *as text*, rather than broken up for computational work in a database, and in a descriptive manner (what a thing *is*, rather than how it should be presented). One of the most prominent theories (which at a syntactic level, XML imposes²⁷) is that text is an ordered hierarchy of content objects (the so-called OHCO thesis: De Rose

²⁶ This is, perhaps, not true of HTML in its most pristine state; rather, it can be regarded as a kind of descriptive mark-up for a rather constrained document model, transformed into a generic presentational form by the browser. Later abuses (browser vendors allowing and — save us! — <blink> and <marquee> tags) shifted this balance towards the presentational. While not descriptive mark-up in the same way as TEI, I would suggest that it *is* descriptive nevertheless; rather, the distinction lies in the imposed constraints of the document model.

²⁷ Actually, it does not. As Cummings argues (in almost all his introductions to XML) XML is an abstract data model, and thus has no such restrictions. It is in its serialisation as *inline* mark-up that the problem of overlapping hierarchies appears. XML is, in theory, capable of representing any structure — as seen by such techniques as pointers and ‘stand-off markup’. GraphML is an XML schema for encoding graphs as lists of nodes and edges: and if you can encode a graph, you are really not bound by any hierarchy. That said, it is also reasonable to argue that inline markup is a great strength of XML from the encoder’s perspective.

et al. 1990). This view of text can be seen as deriving predominately from print culture: a book contains chapters, which contain paragraphs, which contain text. However, when desiring to encode features of a text not bound to such an obvious hierarchical disposition, the problem of overlapping hierarchies arises. (Barnard *et al.* 1988; Huitfeldt 1994; Pierazzo and Stokes 2010) The case of poetic lines, words and enjambment is one often-cited example (Schreibman 2013). From a documentary perspective, likewise: paragraphs, for instance, overlap page boundaries. Renear’s attempts to ‘rescue’ the OHCO thesis by, for instance, allowing that a text contains multiple “perspectives”, which can be “decomposed into OHCOs”, were not (as he admits) successful. (Renear *et al.* 1993)

Such problems aside, the TEI has become a kind of *de facto* standard for encoding texts in the humanities — though it frames itself more as a set of guidelines and a standardised (though extremely wide-ranging) vocabulary. The TEI provides facilities for encoding critical texts and apparatus as different readings, as well as a vocabulary for marking up features of a text (names of persons and places, for instance). Its most recent iteration (TEI P5: 2007) added additional support for manuscript encoding, taking pages, regions, zones and lines as base units (rather than more ‘abstract’ textual codes, such as paragraphs), facilitating a document-oriented approach. It also affords a standardised, yet highly adaptable, approach to encoding metadata (the *teiHeader* element).

The success of XML and TEI can also be attributed to the wide number of tools for manipulating structured text. XSLT (combined with XPath; but also XQuery) in particular facilitated the use of TEI for digital editions by allowing the use of descriptive mark-up in XML, before transforming it into a presentational format (HTML) for display on the web. Most significantly, this decoupling of data and presentation allows the creation of multiple

views of the text, generated from the same source file. For instance, with a critical text marked up with all readings, the versions that would, in print, be relegated to the critical apparatus, could be presented in full. Further developments, particularly XML-native databases such as BaseX and eXist-DB, have, to an extent taken us full-circle: preserving the capabilities of descriptive mark-up, while also bringing the power of complex database queries, including full-text search.

Development in this area, particularly with regards to XML, is far from static. Pierazzo's keynote presentation at the TEI Consortium 2015 meeting argued the necessity of separating TEI as a model for text encoding from its dependence on (and shaping by) XML: "the TEI-as-model conflates two different types of models: a conceptual model of text and textuality, and also its implementation in the shape of an XML schema." (Pierazzo 2015a) In her proposal, XML would be one of several possible languages in which to 'use' the TEI. In the development of other languages, one of the principle aims has been avoiding the problem of overlapping hierarchies. LMNL is one such approach; syntactically similar to XML, it does not require the proper nesting of elements, being non-hierarchical: "[It] does not necessarily impose structure over the text; instead, it is construed as indicating arbitrary regions or ranges." (LMNL 2012) Broughton, while using XML, has attempted to rigorously separate and decouple 'layers' of encoding (document structures, text structures, and higher levels of interpretation). (Broughton 2016) Treating text as a graph has also been suggested as an approach. Andrews and Macé have explored this idea in relation to modelling textual variants. (Andrews and Macé 2013) Tummarello explored the potential of RDF graphs for text modelling. (Tummarello *et al.* 2005) A more recent approach by Dekker and Birnbaum is the TAG (Text as Graph) model, which considers text as a hypergraph. In this regard, it avoids the problem of overlapping hierarchies, while also enabling discontinuous mark-up. At present, it lacks a serialisation format, though it has its own query language (TAGQL),

similar in form to SQL. These developments should be regarded seriously, as they offer the opportunity to separate text from arbitrarily-imposed bounds, instead considering it as what Witmore has termed a “massively addressable object”. (Witmore 2010)

1.3.4. Digital documentary editing

The above sections have aimed to outline the development of the digital medium as a viable vehicle for scholarly editions, and a number of tendencies arising from initial experimentation with the medium. This does not, however, answer the question posed by Pierazzo as to whether scholarly editing in the digital medium is editing as before, only transferred to a new medium, or representative of new theoretical paradigms. As the introduction to this chapter suggests, schools within textual scholarship may be seen as the confluence of particular kinds of text, theoretical positions towards that text, and the requirements of producing an output, in terms of medium, audience, cost, and myriad other factors, which, taken together, achieve a kind of (temporary) stability. If anything, the digital medium has brought previously marginalised considerations — not least, as McGann argues (see above), the constraints of print — to the fore. Whether any new stabilities have arisen is, however, an open question.

Having said the above, it is clear that the digital medium has favoured particular kinds of editions, and therefore conceptions of text, over others: most notably, a renewed focus on the document (as opposed to the work), and (almost as a corollary) documentary editions. As Pierazzo says, “It is a truth universally acknowledged that documentary editions have found a very welcoming home in cyberspace.” (Pierazzo 2014) For the documentary edition, the digital offers many advantages over print, not least the ability to include high-quality

facsimiles, and the possibility of encoding (that is to say, formalising) physical aspects of the document. These affordances have further shifted the remit of documentary editing from its traditional historian-orientation to encompass literary works that would, more conventionally, be represented in critical editions (or editions with a greater emphasis on their status as literary works).

The *Austen Manuscripts* (Sutherland, *ed.*), for instance, presents an edition that is explicitly document-oriented, even to the extent of ignoring ‘immaterial’ structures in the text. There is, for instance, no delimitation of individual stories or pieces, except insofar as they coincide with physical page boundaries.²⁸ Moreover, many of the visual features of the document-text, particularly deletions, are represented as such (a more work-oriented approach would, typically, remove the deletions on the grounds that the author evidently intended them to be deleted). That these kinds of features — to which must be also added complex typographical rendering of, for example, different sizes of Austen’s handwriting — can be represented easily in the edited text is an illustration of the strength of the digital medium. Moreover, the collocation of the edited text with page facsimiles allows a much easier contextualisation of the (nevertheless necessary) formalisations required even with such detailed descriptive encoding.

In this respect, a digital documentary edition may be seen as the digital successor to various forms of diplomatic (diplomatic, super-diplomatic, type-facsimile) editions in print, only more so. According to Pierazzo, who has been perhaps the most consistent in arguing the

²⁸ Robinson (2013, 125-7) is critical of such an approach, highlighting the difficulties of locating a given story; this eschewing of the work-text, in favour of the document-text, is, he contends, “distancing” for users of the edition.

merits of documentary editions (see Pierazzo 2011, 2014; Pierazzo and Stokes 2010), the digital documentary edition enables:

The recording of as many features of the original document as are considered meaningful by the editor, displayed in all the ways the editors consider useful for the readers, including all the tools necessary to achieve such a purpose. (Pierazzo 2014)

Indeed, it is easy to see the groundwork laid for a more robust engagement with documentary editing in earlier editions, such as those discussed above. Both the *The Walt Whitman Archive* and *The William Blake Archive* took a documentary approach — at least insofar as they deliberately eschewed the establishment of works as a single, critical text (the degree to which they set out to record the physical characteristics of the actual document varies). Indeed, such ‘archive’-style editions may be conceived as a collection of documentary editions. Moreover, the proliferation of digital facsimiles online point to a natural starting point for the development of documentary editions: “[t]he availability of digital facsimiles and their ease of publication on the web have already been cited as one of the main reasons for the success of DDEs.” (Pierazzo 2014) From this starting point, the furnishing of a transcription — especially to enable text-searching — can be seen as a natural succession. This suggests, in fact, that apparently documentary editions (even when not specifically theorised as such) arise as a *de facto* outcome of digitisation and publishing on the Web.

This ability to combine — or rather, present in parallel — a photographic facsimile and a transcribed text is particularly advantageous (indeed, this side-by-side layout is virtually ubiquitous in most scholarly editions, even those that do not purport to be documentary editions). Such separation means that the sole weight of rendering the physical details of the document does not reside with the ‘text’, as it might in, especially, a super-diplomatic or

type-facsimile edition. The various levels of textual ‘meaning’ comprising the document are spread across two, both in their own way more particularly adapted, forms. The purely visual, such as intricacies of layout and signs or symbols difficult to convey in typed glyphs, may be represented *in situ* by the facsimile (and its presence indicated in the transcription), whilst the pressure on the transcription to indicate these features is reduced, allowing the linguistic features to be modelled in a way that is easier to read (and search) as text. From this perspective, it is possible to see the flaw in Kiernan’s assertion that a pure facsimile edition “subsumed the diplomatic edition” (Kiernan 2006, 266). Rather, for Sutherland and Pierazzo, the digital documentary edition suggests a new relationship between the transcribed text of a document and page-images: whereas, in print diplomatic editions, the transcription was seen as a surrogate for the manuscript, and thus judged on “the implied precision of the editor’s translation” (Sutherland & Pierazzo 2012), in digital editions that are image- and transcription-based, the transcription (formerly the basis of the edition) now stands as an “interpretative tool” for the reader of the manuscript (Pierazzo 2011). As such, the facsimile is more than merely an “illustration”, as it might have appeared in print. (Gabler 2007, 2)

As Pierazzo points out, the creation of this kind of edition is, regardless of the detail that it may be possible to encode with TEI (indeed, with any formal system), still the result of an editorial act, that is to say, a ‘critical’ interference with the text. (Pierazzo 2011) While, for example, the manuscript description set of elements provided by TEI P5 may not ‘interpretative’ at the level of the text itself, even such ‘pure’ description should be considered interpretative. Citing Sperberg-McQueen’s argument that a document contains an infinite number of facts, of which only a limited proportion may ever be represented, Pierazzo argues that it is useless to consider a documentary edition anything but subjective (Pierazzo 2011, 4–7). Both *what* to encode and *how* to encode it are subjective decisions. Digital

documentary editing in this respect mirrors at a lower level of granularity the same arguments concerning objectivity through the accumulation of material seen earlier with digital archives: the accumulation and formalisation of physical detail of a manuscript does not represent objectivity either, but instead marks a deferral of critical activity to the reader, and, moreover, a critical activity constrained nevertheless by prior editorial activity.

Such a perspective is echoed by Dahlström (2004). While recognising the potential for such a deferral of criticism — via a digital encoding of texts to later serve as the basis for critical activity — Dahlström views the production of a digital version of a text to be an act of “complex translation rather than simply one of transmission” (2004, 1), an act that is not neutral. As such, its capacity as a reproductive device is constrained. As Dahlström suggests:

Any close reading of the tools as texts, however, reveals their situatedness: their dependency on particular historical media settings, their socio-cultural roles and functions, or their argumentative, rhetorical dimension. (Dahlström 2004, 19)

With this in mind, Pierazzo’s statement that “the provision of digital documentary editions *per se* does not represent an editorial theory” (2014) may appear somewhat odd. From one perspective, the provision of document-texts (rather than works) is not new: the digital component does not, despite heightened representational capabilities, necessarily imply an particular conception of the text . In this regard, it is a continuation of a pre-digital approach in a new medium. However, as argued earlier, any kind of editorial practice entails an interplay of theory and (amongst other things) representational technologies. As such, the

provision of digital documentary editions, or at the very least, the turn towards a more documentary focus in the digital, *is* a theoretical–pragmatic position.²⁹

For Gabler, a documentary focus is clearly a theoretical position (see Gabler 2007, 2012). True, this is not an explicitly digitally-bound theoretical position, and, indeed, bears many of the hallmarks of earlier editorial theories, particularly within the German school of editing. Indeed, his desire to put “the horse of the document properly before the cart of its eventually emerging text” (2007, 201) may be seen in the light of the general move away from a focus on author-centrality (and in particular authorial intention) in scholarly editing. However, its theoretical position is tightly bound with what he views as the potential of the digital medium to instantiate editions on such lines. His conception of a *work site* — in which the myriad documentary evidence comprising the work is arranged, with the work conceptualised as an emergent phenomenon from this “web of discourses” (Gabler 2010) — is explicitly conceived as digitally-based: “digitally explorable knowledge sites”. (Gabler 2010, 43) In such a conception, we see echoes of both McGann’s *Rosetti Archive* and Shillingsburg’s “knowledge sites”. However, Gabler goes further in envisaging a new relationship between the ‘edition text’ and apparatus, commentaries and other traditional components of print scholarly editions — not with the later hierarchically subordinated, but placed at the same critical level as the former. (Gabler 2010, 44-6) In some ways, this can be viewed as simply the outcome of the facilitation of access to apparatus seen by Siemens (see above), in that they need not be relegated to footnotes or endnotes. But, I think, it points to a wider theoretical position: stressing to a greater degree the role of the editor as mediator, rather

²⁹ To argue otherwise, I think, is to suggest that print-based scholarly editions were manifestations of purely textual theory, with no consideration of how the medium affected such theory.

than adopting a self-effacing posture — a “hand-maidenly bashfulness,” as Gabler puts it. (Gabler 2010, 45)

1.3.5. Digital critical editing

This tendency towards documentary editions, and, more generally, a shift in textual focus from the text of the work to the text of the document, has not been without its critics. In particular, this is due to a (perceived at least) marginalisation of traditional textual critical activity. Fischer, for one, has observed the lack of “truly digital and truly critical” editions. (Fischer 2012) Robinson is perhaps the most outspoken critic of this trend (see Robinson 2009, 2013). For Robinson, the prevalence of document-oriented editions in practice has come at the expense of theoretical concerns:

The dominance of the document model of textual editing in the digital realm suggests that a theory of digital editions is emerging, based on page-by-page transcription of individual documents, which asserts that a digital edition should concentrate on the documents alone. (Robinson 2013)

In other words, the *de facto* tendency towards a document-focus has begun to be encapsulated as a theoretical paradigm. This is, in part, a welcome corrective to the “the millennia-long concentration of scholarly editors on the work rather than the document.” (Robinson 2013, 127) However, Robinson views this as potentially problematic. His critique rests on two positions, one pragmatic and the other theoretical. The pragmatic argument has already been alluded to: that, especially in the case of literary works, an explicit document-centric approach

undermines, and even masks from the user, the text of the work.³⁰ Such a view — and in particular its contrast to Gabler’s outlined above — can be seen as striking at the heart of the function of scholarly editors. Robinson’s perspective might well be seen as containing echoes of Bowers:³¹ a scholarly editor’s role appears to be the provision of texts (reliable texts) to allow further study of them *as* works of literature. In other words, that there is a clear separation between textual and literary criticism, the former serving as the foundation for the latter. From this position, features of scholarly editions such as a critical apparatus are important only insofar as they validate the editor’s choice of reading (unlike Gabler, for whom presentation of the *textual tradition* of a work is the function of a scholarly edition³²). This is not, then, a theoretical criticism of documentary editing *per se*, but rather the identification of an apparent mismatch between the results of such an undertaking and the requirements of (at least a specific set of) users: Why, Robinson essentially argues, should a user have to wade through myriad versions, littered with corrections, to access the *work* that is ostensibly the object of study?

Robinson’s second criticism of an explicit documentary focus is more theoretically grounded, framed as a response to Gabler’s argument that the editor should concentrate on what is “endogenous” to the document (as opposed to exogenous: the author, wider notions of the work or text). Robinson suggests that not only is impossible to draw such a sharp distinction,

³⁰ In other words, as a *single* text that (however this might be defined) comprises the work, and may be read without the encumbrances of its actual existence in any document.

³¹ Bowers’ contestation that most literature scholars are more worried about the “pedigree of their dogs” than the authority of the texts they use in their scholarly work marks, perhaps, the apogee of this perspective; accordingly, the duty of scholarly editors was to ‘authorise’ (in the absence of the author) a particular version of a work, providing a stable foundation for further work.

³² In this divergence of views, we might well see in microcosm a repetition of the divide between the Anglo-American and German editorial traditions discussed earlier in this chapter.

but that any engagement with a document requires consideration of exogenous factors. (Robinson 2013, 116) He cites Bordalejo's example of the word rendered as "“d u r a-with-an-underdot o” in a manuscript of Dante's *Inferno* (Bordalejo 2010): the “a-with-an-underdot” should be interpreted as a deletion of the “a” (rather than as a diacritic) — but knowledge of Italian (that “*durao*” is not a word) and of scribal practice (the dot designating a deletion) are essential to interpreting this as such. These factors are exogenous to the document. Indeed, I would suggest, this argument may be extended to cover any interpretation of any inscribed text: a shared awareness of the forms used to encode text in the document is essential for any interpretation. While it also seems to be the case that the kind of ‘exogenous’ factors Gabler has in mind — particularly authorial intention — operate on a higher level than these minutiae,³³ Robinson's argument does illustrate the impossibility of drawing too sharp a distinction between the endogenous and exogenous aspects of a document. However, it is nevertheless not clear that such an argument stands as a critique, *per se*, of a documentary focus, or of the production of digital documentary editions; in one way, this may simply be saying that no interpretation of a document is objective (see Pierazzo's point earlier).

Fischer's conception of a ‘truly digital’ and ‘truly critical’ edition (Fischer 2012) differs somewhat, positioning the critical text not as a representation of the ‘work’ *per se* — at least, not as the totality of the work in the kind of platonic conception given by Robinson — but as an organisational and interpretative tool. A critical text, included as part of an edition comprising documentary evidence, provides “guidance” by offering “suggested reading based on expert analysis” (Fischer 2012). A critical text, in this understanding, may be regarded as

³³ One could consider the author's use of specific forms as ‘non-intentional’ — being part of a wider convention, separate from ‘the work’ *per se*.

a suitable way for an editor to offer his or her own theory or argument about a text,³⁴ whilst still, positioned as it is within a web of documentary evidence, allow the user to explore (and perhaps question) the editorial choices made.³⁵ In this respect, we see a position closer to Gabler's notion of a "work site" (and a similar foundation in the German editorial school): that is, the practice of textual criticism without the constraints of a print edition's requirement for a single, 'critical' text (the *utility* of a critical text can be seen, I think, as distinct from its *necessity*).

The so-called 'hyperstack' edition of St Patrick's *Confessio* (in whose development Fischer was closely involved) exhibits such an approach, including not only multiple manuscript versions, but the texts of existing critical editions. Apollon considers the possibility of critical editing in the digital realm in a similar way: "[C]ritical editing' would mean not only to collect, to select, to reconstruct, and to establish but also to systematize a work's contextualization in a vaster environment (e.g., by creating hyperstructured multimedia collections)." (Apollon 2014)

From the above arguments, it is possible to see an emerging set of paradigms in digital textual scholarship — specifically, an emergence of a critical approach that is, or can be, either

³⁴ The notion of an edited text representing the editor's 'argument about a text' has been made in numerous quarters, including Cerquiglini (1999) and Eggert (2009). In all theories that eschew an essentialist nature of text, any edited text represents an 'argument' about that text, rather than the 'text itself' (even editorial traditions inclined to the reconstruction of *the* text — normally according to the reconstruction of intention — do so via argumentation, and acknowledge that other arguments may be made).

³⁵ Fischer's argument presented in this paper is made, perhaps, less contentious by its specific application to medieval texts edited using a stemmatic method. As has already been suggested, the removal of textual contamination through a study of propagated errors is generally regarded as less controversial than constructing a previously non-existent text based on authorial intent.

explicit or latent: that a work can be determined either via a single critical text or be present as an emergent phenomenon arising from the collocation of documentary evidence; or, indeed, both. Andrews, most notably, has called this “third way philology” (Andrews 2013): the utilisation of the digital medium to cleave a path between the “old philology” based on rescuing texts from corruption and “new philology” which seeks to represent texts in its historically contingent states. For all this, it is difficult to see that digital scholarly editing has crystallised into a distinctive ‘school’, or even multiple schools. The theoretical divisions highlighted in previous sections seem to be continuations of previous, pre-digital debates, even as certain positions are given fresh impetus via the pragmatic possibilities of the digital. What distinguishes certain kinds of digital edition from others is, primarily, the type of text being edited, and the domain in which, and for which, it is being edited: the single greatest divide is still between editing of literary works and the editing of historical documents.³⁶ That said, the digital turn may be seen as effacing somewhat older definitions, even as debate rages: ‘critical’ editing has been widened from its earlier narrow scope of ‘creating a critical edition’, and theorists are much less likely to assert the ‘non-criticality’ of documentary editions.

It may even be argued that the digital medium, while not eschewing theory, does not lend itself to the establishment of ‘schools’ of editing. This, I would suggest, is a result of the malleability of the medium, and the rapid development of what can be represented by it and how. In particular, I am thinking of user-interfaces. If print-based scholarly editions were apt to fall back on conventions, this can be seen as reflection of the medium’s constraints: only a limited number of possibilities were available, and these tended to be codified and

³⁶ Despite the historicising of literary works through a renewed focus on their material embodiments, and the facility offered by digital editions to represent this, in the background there always exists, even implicitly, the conceptual work. (If Gabler insists on editing documents, it is no accident that he also refers to the digital accumulation and structuring of these documents as a “work site”.)

crystallised in practice. (If, for instance, a reader had to learn to read a text through myriad, interfering typographical signs, it made sense to adopt a system with which users were already familiar.) The digital medium allows, however, a much closer alignment of the interface with the specific documents or works or texts being presented. Each edition can (for want of a better description) establish its own ‘best practice’.

1.3.6. Digitisation for other purposes

This chapter (indeed this thesis) takes its starting point from textual scholarship as traditionally practiced, and latterly transferred to the digital medium. However, it is, I think, mistaken to treat digital scholarly editing in isolation, separate from general milieu fostered by the same developments (technological, and especially, though not restrictively, within academia) that enabled this digital turn in textual scholarship. The advent of the digital medium — and in particular the Web — as a tool for myriad scholarly activities has led to the digitisation, in one form or another, of text-bearing objects by people who are not (ostensibly) textual scholars, for reasons that are not (explicitly) ‘textual scholarship’. In other words, lots of people, even within academia, put historical documents, texts and works on the Web, without considering themselves textual scholars — and, indeed, approaching their work from rather different perspectives.

Coyle’s article, ‘One Word: Digital’, is a brief but thorough outline of differing motivations for digitising text. Of these, “Digitisation for Preservation”, “Digitization for Discovery”, “Digitization for Delivery”, “Digitization for Reading”, and “Digitization for Research” (Coyle 2006) are perhaps the most pertinent to textual scholarship. That is to say, they overlap with aspects of scholarly editing, as transferred to the digital medium, in a way that is difficult to precisely delineate. Indeed, the roots of digital scholarly editing have already

been seen as an extension of digitisation-for-research (the digital text-as-text being a “byproduct” of the digital text for computational research). In digitisation-for-preservation, again, we see echoes of Shillingsburg’s description of a scholarly edition as an effort to “preserve or rescue a work of artistic, social, intellectual, or historical importance as an artifact.” (Shillingsburg 1986, 4). Digitisation for “discovery”, “delivery” and “reading” can be seen, moreover, as *de facto* properties of digital scholarly editing (if no user can find the edition, access it, nor read it digitally, the edition might as well not exist).

As this overlap does not, however, necessarily require that any one endeavour, carried out for its own end, fall into the domain of scholarly editing, the question arises: to what extent does an activity, for some other purpose (especially a scholarly purpose) automatically become an issue of textual scholarship? Such a question, I think, strikes at the heart of the problem: that there is no overarching conception of what a scholarly edition is.

From one perspective, this shift can be seen as a ‘widening’ of the scope of textual scholarship from within, particularly with regard to what might be considered ‘critical’ activity. The argument that creating an edition of a single document’s text is a critical activity on par with ‘critically’ editing a single text from multiple documents exemplifies this tendency. By the same token, however, it also lessens the grounds for excluding other kinds of editions — in particular, facsimile editions — on the basis that they, likewise, involve the critical application of expertise. Gabler, for one, highlights this tendency, suggesting that facsimile editions are only “marginally [...] considered *editions*” as a result of the neglect of the material document in favour of the text. (Gabler 2007, 198) True, such editions might mandate a “bolster[ing] by paratexts” (Gabler 2007, 198) as a sign of their intellectual validity, but it is nevertheless an indication that the text of a document might be represented without any

explicit ‘editing’ of the text itself. In this respect, we see what might be regarded as the widening of the field of textual scholarship from within. To this, I think, it is appropriate to add the role of textual scholars as foundational to the development of the Digital Humanities as a field.

From an external perspective, however, it is also possible to see an ‘encroachment’ of other fields around the margins of scholarly editing, enabled by the digital medium. In particular, archives and libraries assume new functions in the digital realm. These fields have their own traditions and theoretical positions; moreover, both have traditionally played a role in textual scholarship, especially in activities preliminary to the creation of scholarly editions (bibliographical work being the most obvious). Archiving and librarianship both play a different role in the digital realm, in that they entail a *remediative* aspect quite apart from their roles in preserving and organising physical, text-bearing artefacts. Archival digitisation,³⁷ for example, preserves artefacts by producing digital surrogates, especially photographically. This stands in marked contrast to both traditional archival practice, which preserved the original documents; and ‘traditional’ textual scholarship that preserved the linguistic text by extracting it from its degradable material instantiation. However, as Van Hulle notes, digital tools enable a crossing of this division:

Usually, digital ‘archives’ are distinguished from ‘editions’ because the latter offer a critical apparatus. An interoperable tool such as CollateX can enable any user - not necessarily an editor – to transform a digital archive into an electronic edition. As a consequence of these developments in scholarly editing, the strict boundary between digital

³⁷ *Not* a ‘digital archive’ in Price’s sense. (Price 2009)

archives and electronic editions is becoming increasingly permeable, resulting in a continuum rather than a dichotomy. (Van Hulle and Nixon, 2011, 118)

As someone whose work straddles the divide between (digital) textual scholarship and other forms of digitisation, the theoretical perspectives offered by Dahlström perhaps best exemplify this movement from an archival, library and museum perspective (see Dahlström 2004, 2010, 2012; Dahlström and Hansson 2008). Firstly, he sites digitisation within museums and archives as a critical, signifying practice (Dahlström 2012), both in terms of remediation of actual documents and their organisation (the overlap with textual scholarship on these grounds is evident — see the reference to Pierazzo 2011 above). Secondly, and most pertinently, Dahlström explicitly draws a line between the practice of digitising text *as text* carried out by libraries and archives and the creation of scholarly editions. (Dahlström 2010) While recognising the traditional absence of an overlap — chiefly on the grounds that library digitisation does not aim to ‘establish’ a text from multiple documents — Dahlström notes that this distinction fails to take into account tendencies towards a more document-oriented mode of editing described above. Moreover, the practice that he frames as “critical digitisation” (Dahlström 2010, 2012) can be seen as establishing, if not an “ideal text, then at least [it entails a] kind of document criticism that seeks to establish an ideal document.” (Dahlström 2010, 95) The distinction between *image digitisation* and *text digitisation* (Dahlström 2010, 88) also holds little water: not only does image-digitisation digitise the text of the document (if only by proxy), facsimile-only editions are, as noted above, increasingly seen as valid. In addition, he observes the overlap in practice: in terms of method and digital tools, and in collaboration between library digitisation and the creation of digital scholarly editions. “So scholarly editing and library digitisation are perhaps approaching a

point where the two not only meet but merge, at least on the project level.” (Dahlström 2010, 96)

Such a conception can, perhaps, best be regarded as a cross-fertilisation: in carrying out ‘critical digitisation’, librarians and archivists are borrowing the methodologies from textual scholarship; and, in creating scholarly editions that explicitly aim to reproduce the text of documents — rather than critically establishing the text of works — documentary editors are carrying out an operation that cannot be realistically distinguished from digital librarianship or digital archiving.

Given this, it becomes increasingly difficult to establish a general case for scholarly editing to be considered as distinct from any kind of text-based digitisation — including at the periphery. By this ‘periphery’, I mean particularly so-called mass-digitisation projects, of which Google Books is the most prominent. In comparing ‘critical digitisation’ to mass-digitisation, Dahlström goes some way to articulating the differences between the two: their orientation towards scale or depth and the amount of critical work (over automation) are two most pertinent characteristics. (Dahlström 2010, 91) Many critics of mass-digitisation echo these criteria in their criticism. Coyle (2006) observes the high level of errors in Google Books — pointing to a lack of rigour that is clearly unacceptable in a scholarly edition of any kind. Nunberg’s famous and often-cited blog post, *Google Books: A Metadata Train Wreck*, is even more scathing, pointing to the misclassification of books (a bibliographic failure *par excellence*). Moreover, as Dahlström suggests, perhaps the greatest failure of mass digitisation is its conceptual treatment of all instances of a book as identical. (Dahlström 2012)

To my mind, however, it is difficult to exclude such mass-digitisation practices from the domain of scholarly editing on theoretical grounds. In terms of quality, for instance, it is not difficult to envisage more sophisticated optical character recognition algorithms producing fewer — if, indeed, any — errors (or to be at least as good at transcription as human readers). ‘Quality’, after all, can only really be described as sufficient for a given purpose, and the same applies to scholarly editions.³⁸ Likewise, scale itself cannot be seen as grounds for any distinction (there are very large scholarly editing projects). And the most obvious dividing line — textual scholarship’s aim of establishing a ‘lost’ text — has been relegated to one among many paradigms and thus no longer holds true. In one sense, therefore, textual scholarship’s embrace of a pluralistic understanding of ‘critical’ activity has opened the door to myriad paradigms that can be framed as critical; and it is a door that is difficult to shut again.

As a final consideration, we may consider the form taken by the digital scholarly edition. At the forefront of such arguments is the question of whether a scholarly edition need actually represent the text of documents or works *as text*. In this we see, perhaps, the encroachment of more broadly placed digital humanities work, whose results are made available as digital resources. Indeed, this may be regarded as a return to digitisation for computational study — the relegation of the digital text to a by-product. Again, the recognition that a scholarly edition is only a partial representation of a text rears its head: if any representation is incomplete, on what grounds can a scholarly edition of a text be distinguished from some other representation of the text (assuming of course that it is carried out with scholarly rigour)? The output of myriad computational tools — concordances, key-words-in-context,

³⁸ The tendency of scholarly editors to label their editions as “definitive” are long gone. (See Kline 1995)

n-gram searches, topic models — all represent the text to some degree, and may be considered ‘scholarly’. This is not, of course, to suggest that anyone *would* consider these scholarly editions, but that finding a coherent rationale for excluding them (and, more important, *not* in one stroke excluding things that *are* considered scholarly editions) is fraught with difficulty.

A recent symposium at Maynooth University on ‘Virtual Worlds as Digital Scholarly Editions’ marks perhaps the apogee of this tendency. “Immersive environments provide a venue to raise issues of how the phenomenology of place and space can be used to design a new language of scholarly editions.” (Schreibman 2013) Schreibman has suggested that her *Battle of Mount Street Bridge* (Schreibman 2015) project can be regarded as a scholarly edition — even a critical edition. Though it presents a three-dimensional reconstruction of the battle, it takes as its basis myriad, often conflicting, documentary accounts of events; the visually-re-enacted battle can therefore be regarded as a ‘critical’ representation — a weighing up of the validity of textual evidence — of its textual underpinnings. Again, this conception of a scholarly edition is far from universal, but at the same time, and especially without attaching some normative understanding of the kinds of textual activity that comprise scholarly editing, it is difficult to formulate a coherent counterargument.

1.4. Conclusion

This chapter began by proposing a model of scholarly editing as a network of partially overlapping practices, with local similarities and shared concerns, but no overall defining characteristic beyond a universally shared sense of rigour. Moreover, as the discussion in this chapter I think has shown, there is no unified sense of what this rigour entails in practice; rather, whatever is done should be done rigorously, judged only by a localised set of practices.

These are in turn determined by the types of text dealt with; theoretical conception of texts, documents and works; as well as the technological restrictions, academic landscape and cultural and economic factors that stabilised them as schools in the first place. This points to a Deleuzian understanding of things as inherently dynamic systems, always in a state of “becoming”, rather than a fixed set of identities or ‘definitions’. (Adkins 2015, 7–9) A school or editorial tradition, in this understanding, is a slowing-down of ‘becoming’, a temporary coagulation of theory and practice; it has a beginning — the breaking apart or extension of previous schools — and may well have an end, or it may change of all recognition.

DeLanda describes Deleuze’s replacement of the Aristotelian notions of the ‘general’ and the ‘specific’ with universal-singulars and individual-singulars. (DeLanda 2006) However, as DeLanda suggests, many things that Aristotle would consider ‘general’ are, in fact, individual-singulars — that is, they can be shown to be historically contingent, not universal. Take the example of species and organisms: a species is not a static, transcendent ideal, of which organisms are instantiations; rather, organisms are individual-singulars, and the species too an individual-singular — “larger spatio-temporal individuals ... [resulting from] causal interactions of a population of smaller scale individuals.” (DeLanda 2006). Accordingly, “a given species is as singular, as unique, as historically contingent as the organisms that belong to it.” (DeLanda 2006) The same may be applied to schools and individual editions. A school is, in this sense, the editions that it produces: a population (At the same time, this points equally to individual editions being singular and unique, an instantiation of a highly specific ‘paradigm’ based on the precise documents or texts at hand.)

In answer to the question, what causes stabilities, therefore, we may say: practice. Theoretical and technological advances (‘shifts’ may be a better, less teleological, term) cause ruptures, or at least increase the speed of change, but it is practice that stabilises. These shifts may be small

and localised: a tentative evolution here or there. Eggert gives a nice example of Thorpe and Gaskell's critique of the Greg-Bowers approach.³⁹ (Eggert 2009, 174) Doubtless it caused some shift — influencing their own editions at the very least — but as a modification was firmly within the general 'copy-text' approach.

Other shifts — particularly those caused by external influences rather than internal critiques — have more global, and more pronounced effects: an asteroid collision, in evolutionary terms. Its effects are not only substantial, but have effected almost all schools of scholarly editing — and not simply by forcing each school to re-examine its own methods, but by effacing the boundaries between some schools (one can, for instance, make an edition containing all the versions of a work *and* a critical text), and shifting the general balance of power (most notably from critical to documentary editing). Likewise, the general shift away from the author as the origin, guardian and determiner of meaning provoked by the impact of critical theory. A shift, therefore, towards “becoming”, towards fluidity.

It is this state of fluidity that has also contributed to the expansion of the field from outside. Once internal structures, such as schools and traditions, have been broken down, the external boundaries of the field (which were only ever *de facto* boundaries) also disappear. Practices that were, traditionally, external to textual scholarship bring themselves into the mix — archival digitisation, for instance — because there are no (even temporary) rigid boundaries to stop them.

³⁹ That a first edition should be chosen as copy-text, rather than the author's manuscript, as it was a better reflection of authorial intent (i.e. that the author *intended* production to be collaborative).

In this chapter, we have focused primarily on the causative mechanisms underlying the emergence, stabilisation and (at times) formalisation of ‘schools’ of scholarly editing. It should, as a parting shot, be recognised that other accounts of textual scholarship that are distinctly ‘anti-school’ have been elaborated. In particular, Van Hulle and Shillingsburg — in *Orientations to Text, Revisited* (2015) — offer a very different perspective. In this account, they present seven ‘orientations’ to text that an editor may take. These orientations, born of a kind of abstract survey of textuality and interpretations, are distinguished from each other by a degree of incompatibility — that is to say, a lack of rationality arises from any *ad hoc* mixture of these orientations. This approach has a particular affinity with that presented here. In particular, it is explicitly non-normative; indeed, if one might for a moment read it in a Deleuzian light, it can be seen as a kind of ‘survey’ of editions that exist (an engagement with the real, though done fairly abstractly) and the extraction of conceptual orientations that exist. Though these can in certain cases (particularly genetic editing) be mapped onto schools, they do not claim to ‘explain’ schools. While this is hardly a weakness of their analysis, I would contend that an analysis of schools is necessary: they are real, and they have effects on editors and editorial practice in ways that do not derive explicitly from an orientation to the text (though they may well imply one). (I will return to this theoretical approach later in Chapter 3, as I think it overlaps neatly in certain regards with the Deleuzian analysis of textual scholarship I present.)

To conclude, finally, we may say that to ask, *What is a scholarly edition?* in today’s climate is really to ask, *Have things settled down yet?* That Pierazzo (2015) poses it as an open question, and Robinson (2013) argues that it is about time it should, suggests not. This, as I suggested in the introduction, poses a problem if we want to analyse digital scholarly editions and other kinds of digitisation — particularly, in the case of this thesis, with a view to reconciling such

practices. The open-ended view would suggest, in fact, that there is no problem: scholarly editing and textual digitisation are well on the way to being reconciled — if only by a lack of boundaries. Alternatively, we might try (as Robinson suggests) to theoretically erect boundaries.⁴⁰ In this, however, there is a risk: that, in erecting some boundary or other and elevating it to the position of a definition, all that can be studied is, somewhat tautologically, that boundary. In the following chapter, I propose a third approach.

⁴⁰ His (somewhat caustic, somewhat tongue-in-cheek) call for “Digital humanists [to] get out of textual scholarship ... and if they will not, textual scholars should throw them out” (Robinson 2013) can be interpreted as such an attempt to construct a boundary.

2. Textual Transmission and Validity

Reason is always a region carved out of the irrational—not sheltered from the irrational at all, but traversed by it and only defined by a particular kind of relationship among irrational factors. Underneath all reason lies delirium, and drift.

— Gilles Deleuze, Desert Islands

2.1. Introduction

This chapter intends to consider scholarly editing and the digitisation of text-bearing objects as instances of *reproduction*; or of *remediation*; or of *transmission*. These three terms are, at least in the realm of textual scholarship, all connected. A digitisation of an object, or the scholarly editing of a text, *reproduces* the object or text (or work or document). It does so by *remediating*, transferring the text from one ‘medium’ — a material instantiation — to another medium or material instantiation. In doing so, it *transmits* the original via the reproduction. Of these, *transmission* is the essential characteristic; reproduction and remediation are simply facts of the matter (scholarly editions do not transmit by, for instance, sending out the original manuscript).

Digitisation and scholarly editing, of course, have their own stated rationales or intentions — none of which are invalid on their own terms — but these must be considered supplementary: without the basic fact of their creating a reproduction, via a remediation, with the purpose of transmitting the original, other rationales fail. “Preservation, access and dissemination” (Buzetti & McGann 2007) are three such common rationales; as is “rescue a work of artistic, social, intellectual, or historical importance.” (Shillingsburg 1996, 3) All perfectly valid — but a digitisation or scholarly edition does not preserve or rescue a text or

work or document if it does not transmit it. It does not disseminate or enable access to a work if it does not reproduce the work. “Editing itself is nothing if not a form of textual transmission” (Pierazzo 2015, 75).

The contention in this chapter is therefore that all kinds of scholarly edition, digital scholarly edition, and digitisation project can be considered, at a base level, as a *transmission*. A transmission, quite simply put, must happen. What, then is transmission? From the point of view of the recipient, the ‘transmission’ is what is received. But, as Dahlström argues, a scholarly edition is “not a simple transmission” but a “complex translation” (Dahlström 2004, 17). In this, we see the other side of transmission: its actually being carried out, by a complex process. The challenge is, therefore, to account for both sides. How can complex processes of translation produce, as an end product, from the point of view of the recipient, what appears to be — or better, what functions as — a transmission? An account of transmission is a study of, above all, its mechanisms.

The first step — with which this chapter is primarily concerned — is the development of an abstract conception of the mechanisms of transmission. This is not to say that such a conception is not real. By investigating abstraction we are not, as Deleuze and Guattari contend, engaging with some ‘ideal’ notion. Abstract mechanisms are as real as anything else: they are the real functions carried out in an actual process, regardless of what, in any specific process, carries out such a function.⁴¹

⁴¹ Integers, for instance, are abstractions — which can be specifically embodied in apples as well as oranges, as well as anything else that can be treated as a unitary whole and counted. The mechanism of addition is unconcerned with the specifics.

Approaching the question in this way allows the sidestepping of any normative conception or defining scholarly editions (especially) according to some kind of essential characteristic. While different kinds of transmission are undoubtedly different, starting from a point of similarity allows us to analyse these differences, and particularly how they arise from the kinds of transmission processes that are employed. In some ways, this is not without problems; but I would suggest it is a necessary first step in this analysis. It avoids, for one thing, any arbitrary distinction between print and digital scholarly editions; and if digitisation is predicated on the digital medium, this is — at least for the purposes of this analysis — beside the point. As Mak argues in her analysis of *Early English Books Online*, “Digitization [should not be] conceived primarily as technological, or, at least, no more or less technological than medieval illuminated manuscripts or printed newspapers.” (Mak 2013, 4)

If I am focused here on scholarly editions and digitisation, it is because of the increasingly close overlap between the two. Dahlström (2010), in particular, coherently outlines the basis of such an understanding, observing the drawing together of scholarly editing and ‘scholarly’ digitisation, as well as exposing the artificiality of possible distinctions — especially those founded on textual intervention *versus* objective reproduction. (Dahlström 2010, 81) Both exhibit “*a general trait of textual transmission*”. (Dahlström 2010, 87, original italics) Briefly, then, it can be argued that, as scholarly editing can now operate in the digital realm, it is digitisation. And digitisation within cultural institutions, as Dahlström notes (2010), increasingly replicates the operations of scholarly editions (as a dual effect of a medium-enabled movement towards ‘scholarly edition-hood’ within domains such as archives, and the general adoption of a more pluralistic understanding of textual scholarship).

This initial position can therefore be seen as taking as its philosophical basis what Adkins sees as the key aspect of Deleuze’s metaphysics (indeed, his whole philosophy): “continuity”.

(Adkins 2015, 1)⁴² Placing things on a continuum can be seen as a rejection of ‘arbitrary’ categorical distinctions, which, as Adkins observes, can only reproduce themselves as “good sense” and “common sense” — in other words, preventing us saying anything new. (Adkins 2015, 23) Pierazzo, for one, has taken a similar strategy — that of the continuum — to highlight the artificiality of rigid boundaries between various kinds of scholarly edition. (Pierazzo 2015, 60)⁴³ My own continuum might be described as a continuum of *criticality*. At one end we might site *mass digitisation* (the ostensibly uncritical, technologically-based application of a generic reproduction) and at the other end, the *scholarly edition* — with a healthy overlap between scholarly editing and digitisation, centred around (so-called) “critical digitisation” (Dahlström 2010; Dahlström, Hansson & Kjellman 2012). Beginning with this continuum, therefore, enables analysis without the restriction of pre-emptive distinctions (being scholarly or not; being digital or in print; critically editing a text or ‘reproducing’ it; Eggert’s ‘editorial’ *versus* ‘archival’ as categories⁴⁴).

This continuum, in and of itself, does not (as argued in the Introduction to this thesis) say much: it would be to simply extract an arbitrary feature (criticality) and to assign a particular value to every transmission. By studying the processes of transmission, however, I will show that a number of critical thresholds arise. These represent ‘turning points’, beyond which significant differences occur. In this regard, it follows DeLanda’s description of thresholds in water temperature: of course, water can have any temperature, but moving below 0° C (one

⁴² See the introduction to this thesis for a fuller discussion of what might be described as Deleuze’s ‘ontological project’: the rejection of any ontological separation of the intelligible and the sensible, in favour of a ‘univocity of being’ (Adkins 2015, 6).

⁴³ Her example places varying degrees of ‘materiality’ on a continuum: at one end, the facsimile edition (the most materially-bound conception of text), and at the other, “unclaimed texts on the web” (“those texts which are presented without any reference to the material editions from which they are presumably drawn”). (Pierazzo 2015, 60)

⁴⁴ In his keynote speech at ESTS 2016, Eggert presented the ‘archival’ and ‘editorial’ tendencies as poles of a continuum.

threshold) causes a state change (to ice), while moving above 1000° C causes another (to a gas); by contrast, moving from 51°C to 52°C has no particular effect (see DeLanda 2015).

I shall refer to the ‘turning points’ on this continuum as *validity thresholds*: that is, the degree to which representation may be validated. One threshold is between *any kind of document* and a *document that transmits another document*. Another is between what I shall call *naïve transmission* and *articulated transmission*. Finally, a third distinction is drawn, around what shall be called *reflexive transmission*. Though this latter falls, broadly speaking, as a divide between scholarly editions and digitisation, it should not be taken as a definition of scholarly editions: indeed, it would seem to exclude some things traditionally considered as scholarly editions, and includes some things that are not. Rather, it is a threshold beyond which a reproduction might be considered ‘scholarly’ — because the reproductions beyond this threshold function as scholarly editions.

These thresholds can be seen as implicit in Buzetti and McGann’s description of scholarly editing: “Scholarly editing is the source and end and test of every type of investigative and interpretational activity that critical minds may choose to undertake.” (Buzetti & McGann 2007) From the above outline of transmission and reproduction, we clearly have a source and an end: the original and the reproduction. What determines these various thresholds is the capacity of the transmission to pass a certain level of “test”.

What gives rise to these thresholds? In this chapter I will suggest that it is process: the process of transmission. Certain types of process may be mapped onto certain characteristics (such as ‘scholarly’ and ‘rigorous’), but it is not the characteristics themselves that give rise to a certain level of validity; it is the process. If a scholarly edition has some unique status within this wider realm of textual transmission, it is a status that must be produced — that is, the

result of a specific kind of transmission process. At the same time, other forms of digitisation themselves, by virtue of different transmission processes, have a different status, and produce different kinds of representation.

Likewise, we must ask, how is this validity measured? What is the test? Here Deleuze's concept of *affect* is useful. This, in essence, represents the reconsideration of properties of an entity not as something intrinsic, but as a capacity to do something. Take, as DeLanda does in several examples, a knife. 'Sharpness' is a property of a knife. As an affect, however, it can be considered as the ability to cut. The important distinction is that an affect is always determined in relation to something else: a knife cuts cheese, but not iron bars. As such, according to DeLanda the term affect is essentially a shorthand for, and should be read as, *the capacity to affect and be affected by*. (DeLanda 2006) What produces a level of validity is therefore capacity of the transmission to produce a given result (i.e. that which it transmits) in the hands of a user. There is a strong overlap here with Drucker's view of text as performative. For instance, a text may, in one sense, be viewed as a property of a document; however, for Drucker, the text is *performed* by someone in interaction with that document. "Objects don't represent, they perform." (Drucker 2013) In other words, the text is an affect. Reproduction, then, is "all about turning this performativity into a conduit." (Drucker 2013)⁴⁵

⁴⁵ Thanks especially to my colleague Vinayak Das Gupta for pointing out a seeming paradox: a perfectly good scholarly edition that happens to be locked in a box that cannot be opened *is not a scholarly edition*. By way of reply (somewhat after our debate!) I would suggest that whether it *is* (or is not) a scholarly edition is not the question; *functionally* it is not, because if no one can read it, it cannot transmit.

With this in mind, in order to explore these validation thresholds, it is first necessary to consider the underlying mechanics of reproduction/remediation as a transmission of the source. The approach to analysing this transmission is heavily based on the philosophy of Deleuze, particularly the concepts drawn by Deleuze and Guattari in *A Thousand Plateaus*.⁴⁶ As such, it considers all text-bearing documents and works as an assemblage. It models transmission as a series of interactions and material translations. This comprises the first half of this chapter. The second half will focus more concretely on what it means for a transmission to be ‘validated’ in such a model, again using concepts derived from Deleuze and Guattari.

Though this model is, I think, an original treatment of scholarly editing and digitisation as the transmission of text-bearing material, it inevitably overlaps with previously established ideas and reproduces many established conclusions along the way. It is certainly not original to regard textual artefacts as assemblages — indeed, Deleuze and Guattari themselves use books to illustrate their conception of assemblages. (Deleuze & Guattari 1987, 5; see also Savat & Harper 2016, and Drucker 2014.) McKenzie’s *Bibliography and sociology of texts* (1986), in its consideration of bibliographic codes and material forms as intrinsic to the meaning of text, can be read as treating the book as an assemblage. Likewise McGann (in particular *Radiant Textuality*, 2001) and his conception and theorisation of the ‘social text’. (Also McGann 2003, 2010)⁴⁷ There is also a strong overlap with Bryant’s conception of the “fluid text” (Bryant 2002) in the initial conception of the source material presented here.

⁴⁶ Deleuze (alone and in collaboration with Guattari) often shifted terminology between books — either assigning new terms to previously established concepts, or modified and/or extended.

⁴⁷ Interestingly, in the former, he describes “textual fields” (McGann 2003) rather than “the social text” (McGann 2010). I prefer “textual field” — as I will argue — as a base-condition for all text-artefacts.

In terms of textual transmission, there may also be seen echoes of Shillingsburg's conceptions of "Reception Performance" and "Reception Text" (Shillingsburg 1997, 54), as distinct from material or ideal texts, and particularly implicated in the activity of transmission. From a perspective of modelling, Pierazzo's two models of the relationships between material, text, work and actors in scholarly editing (see Pierazzo 2015, 53–6) are exhaustive in their detail, and particularly relevant for this argument (being modelled in such a way also makes it easy to identify where our approaches diverge).

Of particular importance establishing a starting point for this discussion is Dahlström's *How reproductive is a scholarly edition?* which pre-empts many of the issues raised here: the iconic status of a scholarly edition, undermined by the 'mimetic fallacy' and 'complete encoding fallacy' (Dahlström 2004), are especially relevant — and are essentially reproduced as arguments by my own model. They point, especially, to the requirement of idealism in reproduction (where "contents are disembodied, separable from their physical document carriers and hence transportable ... to other carriers": Dahlström 2004, 24), at the same time as (in my reading at least) cogently expressing a deep scepticism for the validity of relying on such idealism as an explanation for transmission.

Indeed, beginning from a Deleuzian position necessarily involves a deep suspicion, if not outright rejection, of appealing to the ideal as a mechanism that accounts for transmission. For Deleuze difference necessarily pre-empts identity: "In a Deleuzian ontology resemblance and identity must be treated not as fundamental but as derivative concepts." (DeLanda 2006) This implies the "extreme position" outlined by Dahlström: that "texts are not only media typical but even exclusive to a particular material media". (Dahlström 2004, 24) This is not to say that we cannot consider two materially-embodied versions of the same text as 'the same' somehow, but rather that they cannot be considered 'the same' by default. In this

analysis, I insist on the contrary: by default, they are different; ‘sameness’ is derivative, dependent on interpretation ‘as the same’.

Consequently, my approach here aims to look at the mechanisms underlying transmission; to ask why it is that we consider two inherently different materially instantiated texts as ‘versions’ of each other in the first place — how the mechanisms of interpretation (for it must be carried out somehow) give rise to ‘versions’ out of difference. What I think this model affords is the possibility of seeing the ideal text not as an explanation, but as a rhetorical device at the point of its emergence.

My default position of the text is thus that described by Buzetti:

From a semiotic point of view the text is intrinsically and primarily an indeterminate system. To put it briefly, there are many ways of expressing the same content just as there are many ways of assigning content to the same expression. (Buzetti 2009, 50)

If a text is intrinsically indeterminate, to talk of identity between two ‘versions’ — or even of two versions at all — is impossible. Identity might be asserted, but only as a result of interpretation (a particular determination of an indeterminate system). In this, we see echoed McGann’s third criterion of the ‘social text’: “Textual fields arise co-dependently with interpretative action.” (McGann 2010, 38) Indeed, the position I shall outline does strongly mirror McGann’s theory of the social text.⁴⁸ The logic behind this is to suggest that text-as-semblage is the default state, or put another way, that all text is always, prior to interpretative action, a ‘social text’. In doing so, my objective is to site other theories of text

⁴⁸ Interpreting McGann in the terms of Deleuze and Guattari’s theories of assemblages — an undeniably more abstract position — allows for its extension, as well as affording a raft of critical tools.

as emerging from interpretative activity, rather than pre-empting it (this, I think, is the great strength of McGann’s understanding: if it can be considered a ‘theory’ at all, it is one that recognises other editorial theories as the result of a subjective interpretation).

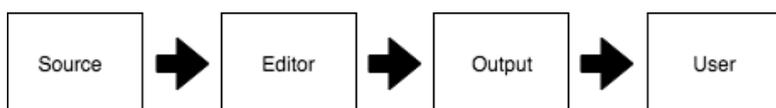
The question arising from taking this position is how to account for transmission from a position where difference (indeterminacy) takes precedence over identity (determinacy).

2.2. Transmission of textual material by reproduction

To begin, we must outline a basic understanding of reproduction that accounts, in the most general sense, for scholarly editing and digitisation activities:

The function of a reproduction is to transmit some source material — literary or historical, physical or ideal, document or work — to the user of the reproduction.

This suggests a path of transmission illustrated below:



*Figure 1: the path of transmission: information flows from the source material, through the ‘editor’ into the output, and thence to the ‘user’.*⁴⁹

⁴⁹ It should be noted that the term ‘editor’ is employed here for brevity (versus alternatives such as *remediating agent[s]*); it should thus be taken in the broadest and most neutral sense. An ‘editor’ might be one person (who may or may not consider themselves ‘editors’ as such), or an editorial team; it might be a scanner or camera connected to OCR software; or any combination of these. Internally, the ‘editor’ might be a complex configuration of human and other actors (really, it is an assemblage): from an external, and simplified, perspective, it can be modelled as a single stage in the process of transmission.

The justification for this transmission might be varied, and may typically be seen as the bridging of some kind of impediment to the source material being otherwise accessible. At the physical level: the document is too fragile to touch; the document is in an archive while the user is on a different continent; the document (necessarily) has only one incarnation, while the users are multiple. As an ideal form (a literary work, for instance): it may be corrupted by damage to its physical manifestation, or by errors introduced in copying, whether by scribes or typesetters or interfering editors; the work may be spread across multiple documents, and thus inaccessible as a work; most fundamentally, because a work is always inaccessible without a physical medium (one cannot ‘transmit’ pure thought). Or the source may be inaccessible, by its nature, to all but an expert, at the level of the inscribed text or linguistically. At a further level, the work or document is necessarily divorced from its historical context in time and space, again requiring expert knowledge to bridge the contextual divide. Most prosaically, the potential user is impeded by not knowing the document or work exists in the first place (a scholarly edition or digitisation conveys the existence of its source).⁵⁰

It should be clear that both scholarly editing and digitisation ‘bridge’ some or all of these impediments: some of these might require more expertise and thus might be more naturally the preserve of scholarly editing; but it cannot be denied that a digital reproduction of a document, for instance, *rescues* degraded material as much as a scholarly edition does. Without drawing unstable categorical distinctions between the kind of impediment bridged, the case for considering both as transmission stands.

⁵⁰ Dahlström says one of the functions of a scholarly edition is as a “bibliographic tool” (2004, 19).

There is a further evident commonality: that transmission in this sense occurs by the creation of a new material document. As Dahlström says, “[r]epresentation is in this case an instance of the activity of copying, reproduction, or ... *media translation*.” (Dahlström 2004, 22; original italics) Transmission is therefore not a constant, linear flow of information. Instead, the transmission occurs via the ‘packaging up’ of information into the output, and the wholesale forwarding of the output to the user. This is necessary, as there is evidently a separation in time and space between the production of the transmission-output and its reception. As such, the output must be considered a discrete, closed-off entity. Moreover, while, for instance, a digital scholarly edition (or digital collection of any kind) might be retroactively changed or added-to by the editor after publication, at the point at which information is transmitted from the output to the user, it is always closed. The output can only transmit to the user the information that it contains at the point of this stage of transmission. There are, accordingly, two stages to transmission — the ‘editing’ stage and the ‘reception’ stage — separated in time and space (see Figure 2). The two stages can be broadly described thus: in the ‘editing’ stage, the editor reads the source material, and then writes the output; the output is then sent out into the world, and is taken up by the user in the ‘reception’ stage, where the user reads the output.

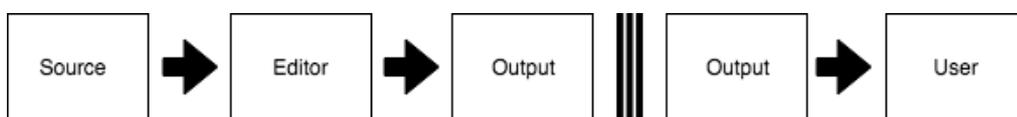


Figure 2: Transmission of source over separation in time and space.

Accordingly, no knowledge of the source (or, indeed, the editor) is transmitted to the user, except that which is contained within the output. For the transmission flow to still function as a result of this separation, therefore, there must be identity between the source and the output. This does not imply that the source and output must be the same — evidently they are different material documents — but that there must be some functional equivalence

between the two. In other words, the user must be able to read in the output that which they could (assuming sufficient expertise in certain cases) read in the source if they had access to it.

An explanation that would, of course, account for this equivalence is that the source and output are analogous to each other (insofar as the same information may be read in both). Indeed, they must be ‘analogous’ (in a manner of speaking) in order for one to transmit the other. But there is nothing in the transmission model presented above to account for this analogy. Without taking material reproduction of any kind to automatically lead to identity — without, in other words, assuming an analogous relation as a given — there is an intractable problem. This is the point to which adopting Deleuze’s argument that difference precedes identity inevitably leads.⁵¹

While couched in different terminology, this is hardly an original observation: many critics essentially say as much from different perspectives. Pierazzo emphasises that all transcription (a material reproduction of text) is subjective (Pierazzo 2011, 466) — thus any intrinsic analogous relation between source and output cannot be founded on objectivity. Sperberg-McQueen observes that all documents comprise infinite facts, only a finite number of which can be represented in a reproduction (Sperberg-McQueen 2009, 31). Logically, therefore, the output re-presents a finite — and therefore infinitely small — number of the source facts (and as such can hardly be taken to be analogous-by-default: quite the reverse).

⁵¹ It should be noted that the analogous nature of the source and output could be determined by *simultaneous* reference to the source and the output: it would be possible in this instance for a user to determine exactly *how* the output is analogous. In this way, we see the establishment of identity in its proper place: not as a given, but as a function of difference between source and output. But this is anathema to a transmitting reproduction, which, quite simply, cannot function as such if it is reliant on access to the source.

The same logic underpins Shillingsburg's critique of McKenzie's 'sociology of text' as a basis for editing: that it necessarily produces an entirely new socially-determined and situated document that cannot represent the original: "The scholarly edition, regardless of its editorial principles, indexes only the forces that produced it, not those that produced the source texts upon which it is based." (Shillingsburg 2011, 17) (Though he talks explicitly of scholarly editions rather than other remediation, including digitisation — "*any* reprint, even an accurate one with no emendations", as he puts it (16) —, the point still stands.) Representation, is, therefore, a determination or interpretation, not a replication.

Dahlström again most directly addresses the issue: it is, he says, "vital to recognise that the target document is always *derivative* to the departure document" (Dahlström 2004, 21) — not a representation. From Buzetti's description of a text as an "indeterminate system" (see above), transmission must be taken, firstly, as a determination of that system: therefore the output, while another indeterminate system, cannot be taken to be the same indeterminate system. The same point is made by McGann's first principle of the social (from which, he notes, all the other principles essentially flow) that a textual field is a Bakhtinian space or "heteroglossia". (McGann 2003, 1)

The problem that arises from this situation, however it might be framed, is therefore an impossibility of representation — against which, seemingly paradoxically, we must set the fact that representation is essential if scholarly editions or digitisations (the output) can function as a transmission of the source.

In describing the situation in this way, we have reached, I would suggest, the logical endpoint of (at least) forty years of debate within textual scholarship. From the 'high point' of

objective, empirical editing, embodied in the Anglo-America sphere by, especially, Bowers (see Bowers 1959, 1964) and Tanselle, there has been a gradual erosion of certainty, to be replaced with doubt and unease. Against the backdrop of structuralism and poststructuralism, we can see the emergence of subjectivity, accompanied by, at each turn, an attempt to peg objectivity at a lower level. The stone cast by, for instance, McGann's *Critique of modern textual criticism*, undermined the authority of the author (and thereby the assumed authority of the editor to edit on the basis of authorial intention). One response to this is to reproduce the texts of documents: if editors could no longer validate their production of a single, critical text, they could at least objectively transcribe the textual content of all the documentary evidence (in this, the digital medium has undoubtedly been a driving force). But here, in turn, is only established a new subjectivity, one providing more information but still not objective. As Buzetti and McGann note:

Facsimile editing is sometimes imagined as a relative straightforward and even simple scholarly project, but in fact the facsimile editor's task is every bit as complex and demanding as the critical editor's. In certain respects it can be more difficult precisely because of the illusions that come with the presence of a single documentary witness, which can appear as a simple, self-transparent, and self-identical 'object'. (Buzetti & McGann 2007)

If the material evidence itself cannot be objectively interpreted, there seems to be nowhere left to go. What 'objectivity' here points to is grounds on which we might ignore the problem of representation — or, rather, on which the problem does not exist at all. But remediation is always subjective, and therefore representation cannot be viewed as a given.

The challenge, therefore, is not to dismiss scholarly editing — and, indeed, digitisation — as fundamentally flawed activities, but to account for their functioning.

2.3. Assemblages, intensities and extensities

To begin approaching the question above, it is necessary to return to and elaborate the functional transmission model articulated in the previous section. Making any further headway requires, in the first place, a much greater specification of each of the terms used: what exactly is meant by the source (when it must, in theory, encompass documents and works in order to account for scholarly editing of all stripes, as well as digitisation practices); the ‘editor’ — which can comprise any number of people and remediating technologies; and the output, which again might encompass any number of forms and technologies. To take the position offered by Deleuze and Guattari, these ‘actors’ in the process of transmission may be considered as *assemblages*.

The study of assemblages forms one of the central principles behind *A Thousand Plateaus* (Adkins 2015; Nail 2017). According to Adkins, “[w]e can think of ‘assemblage’ as an answer to the venerable philosophical question, What is a thing?” (Adkins 2015, 10) An assemblage is used to describe a dynamic system of heterogeneous components (physical systems, living organisms, linguistic utterances, concepts, events). As both Nail (2017) and Panagia (2017) observe, the English term ‘assemblage’ is itself problematic as a translation of Deleuze and Guattari’s term *agencement*: the French term denotes “a construction, an arrangement, or a layout”, while ‘assemblage’ only describes “a bringing or coming together.” (Nail 2017, 22) Accordingly, an ‘assemblage’ in this context should be read not merely as an accumulation, but a functioning system of ‘assembled’ components.

For Deleuze and Guattari, considering things as assemblages (capable of interaction with other assemblages in a heterogeneous system) allows the overcoming of the ontological discontinuity between the ideal and the real, by allowing them to be considered in the same

heterogeneous system. (Adkins 2015, 11)⁵² Assemblages can be characterised by their simultaneous tendencies towards stasis and change. (Adkins 2015, 13) According to Nail, the “theory of assemblages provides an alternative to the logic of essences [...] An assemblage does not have an essence because it has no eternally necessary defining features, only contingent and singular features.” (Nail 2017, 24) An assemblage is, therefore, distinguished by the *exteriority* of relations (connections between molecular components) rather than *interiority* (in which such connections are hierarchically subsumed as *properties* of the extensive whole).

While it is possible to consider an assemblage from the perspective of its unity — as an *extensive*, molar entity — it may also be considered as an *intensity* of molecular components. If assemblages can only be considered from the side of their unity (“[the side that] faces the strata”: Deleuze & Guattari 1987, 4), the ‘naturalness’ of the unity is always a presupposition. In a similar vein, Foucault questions such apparent “Unities of Discourse”,⁵³ arguing that such unities are not naturally occurring but are always constituted by discourse. (Foucault 1972, 22). As such, we are not bound to consider things as ‘already established’ wholes, but see the whole for the contingencies that shape it (via interaction with other assemblages in a wider system):

For example, we cannot extract the being of a book from the vast
historical conditions of the invention of an alphabetic language,

⁵² “The establishment of discontinuity between the sensible and intelligible [...] is the inaugural gesture of Western philosophy and has been the predominant way of dealing with the paradox of the ‘thing’ [its identity as permanent, but also its capacity for change]. The stability of a thing is attributed to its intelligible nature, while the thing’s ability to undergo change is attributed to its sensible nature [...] the properties related to the thing’s intelligible nature are its essence, while the properties related to its sensible nature are its accidents.” (Adkins 2015, 11)

⁵³ The title of the first chapter of *The Archaeology of Knowledge* (Foucault 1972, 22).

distribution of paper, the printing press, literacy, and all the social contingencies that made possible the specific book of our inquiry, with all of its singular features (color, lighting, time of day, and so on) and all the conditions under which someone is inquiring into the book. (Nail 2017, 24)

What Deleuze and Guattari provide, therefore, is an approach that not only throws apparent unities into perpetual doubt, but allows both sides of the equation to be considered. Considering assemblages from the perspective of a molecular intensity allows the questioning of unities: how they arise, under the exertion of which external forces, and, moreover, the mechanics of their dismantling. At the same time, unities as molar extensities are equally real, not merely illusions resulting from an insufficiently deep analysis.

To talk more concretely in terms of textual material, let us take the example of a literary work as an assemblage. On one side, as an extensity, we may say that a work of literature has a natural unity. Moreover, we can give an account of the interactions by which we consider it a unity: it was (possibly) conceived as a unity by the author; it was published as a single entity; it has a physical form with clearly delineated boundaries; there is the name of the work printed on the book. From this perspective, it appears to be a unitary whole. In the words of Buzetti and McGann, therefore:

We name, for example, D. G. Rossetti's 'The Blessed Damozel' with that one name, as if it were a single self-identical thing (which — in that special perspective — it in fact is; which, that is to say, it 'is taken to be'). (Buzetti & McGann 2007)

This 'special perspective' can be read as 'the work-as-unity', in the light of the myriad factors that cause it to be viewed as a unity. To this may be added the printing of lots of copies of

the work, which are taken (as a result of real interactions) to be the same as a function of the difference between them (in this case a relative lack of it). This, it should be noted, is fundamentally an inversion of the traditional conception of an ideal work that is instantiated: there is a real and actual process which creates the work, in its physical form as well as virtually in thought.

At the same time, as Foucault (1972, 23) notes, various works are more or less unitary: a collection of poems has a single form (book), but is more obviously fragmented than a novel; a work may be constituted by a series of books (*A la recherche du temps perdu*; which, of course, might also be considered seven works: they all have different titles and were published at different times); works that form connections to other works (Foucault cites the example of Joyce's *Ulysses* and its relation to the *Odyssey*⁵⁴). What, then, does it mean to edit a work? What exactly is being edited? Does a work pre-exist as a result of the social processes mentioned above, or does an editor, as Pierazzo suggests, "assert" the work? (2014, 56) The answer is both, to a lesser or greater degree, depending on the apparent unity of the work in question. Some works exist in far-flung shreds of documentary evidence that have been forgotten; others are well-established in physical as well as conceptual unity. The question is not whether a work is a unity or not, but how easily it can be delineated as a work (which is really only a question of whether someone — or some process — has already delimited it, and how).

Likewise, text (*the* text or *a* text) may be considered an assemblage. To take up McKenzie's assessment that "the ostensible unity of any contained text ... is an illusion" (McKenzie 1986,

⁵⁴ A perfect example, an assemblage of assemblages: *Ulysses* in interaction with the *Odyssey*, the translation of 'structural' forms, the overwriting of linguistic forms; a complex and dynamic system, which draws in other assemblages (e.g. the physical dimension of Dublin, reterritorialised and recoded in prose).

50), it is really impossible to consider text in any other way (though perhaps “illusion” is too strong a word: the unity of a text at any given point is real and not merely illusory; but it is only a fleeting unity, determined by interactions at a given point. This is, I think, what McKenzie is suggesting in his wider arguments). Sahle’s pluralistic model of text (the “Wheel of Text”) can similarly be seen as suggestive of such a perspective, identifying six ‘dimensions’ of text: Idea, Work, Linguistic Code, Document, and Sign (Sahle, as reproduced in Fischer 2013, Pierazzo 2014.)⁵⁵

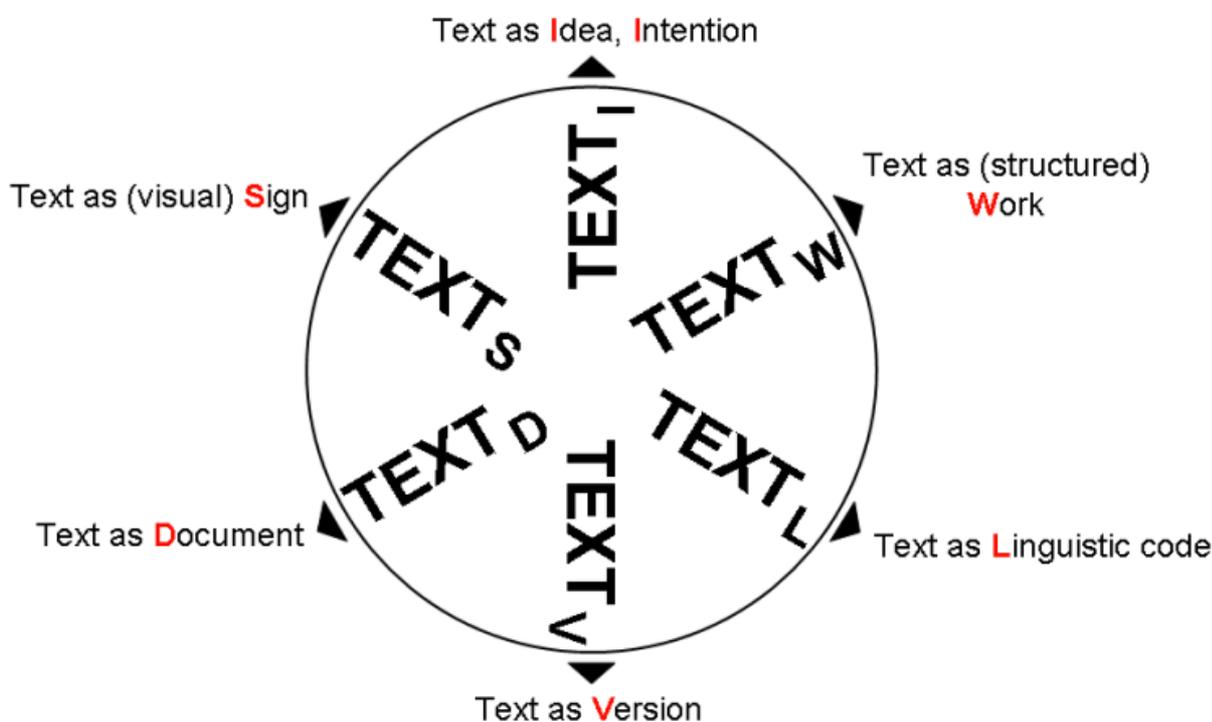


Figure 3: Sahle’s ‘wheel of text’, reproduced from Fischer (2013)

These seem to lie on a continuum from abstract concept to physical instantiation (as such, it echoes the hierarchical FRBR model).⁵⁶ But, one assumes, in Sahle’s model they are placed

⁵⁵ As noted (including by Pierazzo), Sahle’s ‘wheel of text’ — or, at least, the articulation of its theoretical basis — is published only in German; like all other unfortunate Anglophones (or, rather, non-Germanophones), I must cite second-hand accounts.

⁵⁶ Functional Requirements for Bibliographic Records, a recommendation of the International Federation of Library Associations and Institutions (1998), intended for use in library cataloguing. It comprises, in a hierarchical system: work, expression, manifestation, and item (the abstract idea of the work ‘containing’ all expressions — the

in a non-hierarchical relationship (as a wheel) to each other to avoid the suggestion that one takes precedence over any other — they may ‘participate’ to lesser or greater degrees. The wheel of text model clearly identifies ‘text’ as a kind of assemblage. However, by placing ‘text’ itself in the centre, it appears as an extensity: text *consisting of* multiple dimensions. In this regard, it is a discrete formation possessing a delimiting boundary and an exterior, interiorising phenomena as properties of text. This it undoubtedly is, from one perspective. However, the consideration of text as an assemblage allows it to be simultaneously considered as intensity: the coincidence and organisation of forces acting in different dimensions. Text-as-intention, for instance, is an extensive view; intensively perceived, however, it is a connection between the document and linguistic code and the author (or some notion of authorship, perhaps the Foucauldian author-function, depending on the properties attributed to it).

From this perspective, there is no ‘exterior’ to text, and no boundary which encloses the dimensions as properties of the text; rather a system of externally-connected and organised components. From such a view, text is what emerges from the interaction of entities (physical or conception) in these various dimensions, rather than a starting point.

Indeed, one may be given to ask whether the ‘dimensions’ of text that Sahle suggests necessarily comprise unities themselves (to ask, can we treat text as multiplicity rather than plurality⁵⁷). The ‘document’ dimension, for instance, may be more or less fragmentary: a text

“intellectual or artistic realization” —, the manifestation — “the physical embodiment of an expression of a work”, as in a single published edition of a text —, and the item, a “single exemplar of a manifestation”, a particular physical copy of an edition). (<http://www.oclc.org/research/activities/frbr.html>)

⁵⁷ Indeed, I think Sahle’s model already makes this argument, albeit implicitly. The spokes of the wheel are not considered discrete dimensions as such; their organisation *as a wheel* suggests that an infinite number of positions may exist between the identified positions.

on two documents bound together with string *versus* a text inscribed on two documents in separate archives half the world apart. Or put another way, the dimension of ‘documentariness’ is one dimension to be considered; but does not comprise a unity as a starting point, instead requiring a series of presuppositions (which may be justifiable, but are not ‘natural’, pre-existing or self-evident).

However, to start by treating dimensions as extensive means that dimensions are already separate and should be considered as such; but this constrains the ability to analyse the formation of unities that arise due to the interaction of multiple dimensions. For instance, the linguistic dimension of a text is profoundly altered by material degradation: a cigar burn in the middle of a word does not alter the ‘fact’ that the word was once there, but it does alter how that word may be read, and hence the linguistic dimension of the text under consideration. Such an argument is equally applicable to the other dimensions of text. This approach is nothing more than a shift in perspective, from saying a text consists of dimensions x , y , and z , to saying text emerges from the coincidence and particular organisation of dimensions x , y , and z .

Not only is it useful to view text (or works or documents) in this way, it may be argued that scholarly editing and digitisation — any kind of remediation — is necessarily predicated on text-as-intensity. For a scholarly editor, if a text can only be seen as extensive, it is already fixed in all its dimensions: it has reached a state of unalterable, essential perfection. This echoes Barthes in *From Work to Text*, where a work is seen as a text that is already fixed. (Barthes 1984) Confronted with a text in such a state, why edit it in the first place? But this is only a specific manifestation of a wider problem: how can anything be extracted at all from an extensive text? There is nothing that can be extracted. As Nail argues, “unities do not allow for the possible emancipation or recombination of their parts without destroying

themselves in the process.” (Nail 2017, 23) As such, without allowing that texts are intensities of molecular components, they cannot be remediated.⁵⁸ At the same time, extracting molecular components must be seen as a destruction of the (ostensible) unity: “On the other hand, when component parts subsist independently from their internal relations within a unity, they cease to be unities”. (Nail 2017, 23) This, in succinct microcosm, is the (seemingly) intractable problem of reproduction argued in the previous section.

This conception of a text as an assemblage clearly mirrors McGann’s conception of the ‘social text’, particularly his employment of Bakhtin’s term ‘heteroglossia’ (see McGann 2010), which could quite easily be substituted for ‘assemblage’. Indeed, the connection between Bakhtin and Deleuze and Guattari has been well established; Deleuze and Guattari often cite Bakhtin as a primary source (for a comparative reading of the two see Evans 2009). According to Mishra, it was “Bakhtin’s signal achievement to identify the novel as a heteroglot assemblage”. (Mishra 2007) That said, I think Deleuze and Guattari move beyond the idea of the novel as a heteroglossia, and provide a more general account of phenomena as assemblages, and the role of interaction in stabilising and changing them. It is these general principles that will be explored and applied to textual transmission and remediation in the next section.

2.4. Towards stability: the creation of extensities

Deleuze and Guattari call the process of creating unities (or, more precisely, the tendency towards being discrete or extensive) *stratification*. “Strata are Layers, Belts. They consist of giving form to matters, of imprisoning intensities or locking singularities into systems.”

⁵⁸ An ‘extensive text’ would be bound to, for instance, its material embodiment and could not be ‘interpreted’ in any sense. ‘Extensive text’ in this sense is nigh-on oxymoronic.

(Deleuze & Guattari 1987, 45). As Adkins observes, *strata* are the answer to the question: “How do assemblages acquire tendencies toward stasis? Or, better, are there different ways in which intensities are captured and regulated?” (Adkins 2015, 42) Such a mechanism is required as there can be no “universal, transcendent forms” (an ideal form) to organise matter (Adkins 2015, 43). Instead, there is an “encounter between two form/substance complexes” (Massumi 1992, 12). Two real things — assemblages — interact with each other. To borrow Massumi’s example (Massumi 1992, 10-1), making a table is an encounter between wood and woodworker. In this scenario, the wood is the *content* and the resulting table the *expression*. It is not that one (content) is material and the other (expression) is ideal: content and expression are of the same order. Rather, there is an “encounter between expression and content, in which each receives a determination in its struggle with the other.” (Massumi 1992, 13) The distinction between content and expression is, in principle, reversible; that the terms are applied one way round is the result of one ‘overpowering’ the other (the woodworker appears to shape the wood, and not the other way round).⁵⁹

Deleuze and Guattari refer to this mutual encounter as a “double articulation” (1987, 45). Borrowing from the linguist Hjelmslev’s sign model, Deleuze and Guattari divide both the content and expression sides of the double articulation into *form* and *substance* (thus, form of content, substance of content; form of expression, substance of expression):

The first articulation [content] chooses or deducts, from the unstable particle-flows, metastable molecular or quasi-molecular units (substances) upon which it imposes a statistical order of connections and successions (forms). The second articulation [expression] establishes

⁵⁹ It is difficult to see wood shaping the woodworker, though it must, however weakly. Instead imagine one person punching another in the face: on one side, a broken nose, and on the other, a bruised fist. Or, as in the old joke: "I didn't punch him; he ran into my fist!"

function, compact, stable structures (forms), and constructs the molar compounds in which these structures are simultaneously actualized (substances). (Deleuze and Guattari 1987, 46)

They go on to give the example of the formation of sedimentary rock: on the side of the content, there is a selection of *substance* — sediment — whose *form* is a statistical aggregate (particles floating in water organising by size); the processes (an application of a force) that compresses the sediment give rise to an expression, having a *form* — stable layers of sediment — and comprising a *substance*, rock. There is no external ideal dictating the ‘form’ of rock; rather it is the tendency of sediment in water to organise into layers, and then subject to an interaction with a compressing force, that produces the form of rock. In this example, we therefore see a movement from *molecular intensity*, the accumulation of sediment in one place (characterised by having no real ‘exterior’), to *molar extensity*, rock (which clearly has an exterior).

Such a process may further be separated into processes of deterritorialisation and reterritorialisation, and decoding and coding: de/reterritorialisation relates to the “selection or extraction of some set of intensities in order to compose them or place them in a different relation to one another,” (Deleuze and Guattari 1987, *Geology of Morals*) operating on the substance; de/coding relates rather to the altering of form.⁶⁰ Precisely how this happens on any given stratum is dictated by what Deleuze and Guattari term an *abstract machine*.⁶¹

⁶⁰ “Thus wind, water, and tectonic upheaval might both deterritorialize and decode the sandstone. Or, a mason might deterritorialize but not decode the sandstone.” (Adkins 2015, 51)

⁶¹ “An abstract machine is nothing other than the process of double articulation”; “[it] is constituted by the stratum on which it is lodged” (Adkins 2015, 47).

This distinction (de/coding, de/reterritorialisation) expressed in terms of sediment and rock seems rather arbitrary. However, as Deleuze and Guattari observe, different strata (in Chapter 3 of *A Thousand Plateaus* they identify three: the physical, the organic and the linguistic) each have a tendency towards different rates of de/reterritorialisation and de/coding. (“Strata can be distinguished between each other because their abstract machines differ.” Adkins 2015, 47) The linguistic stratum is most interesting, having the possibility to reach beyond itself onto other strata and ‘overcode’ them, to “reorganize other strata on its own terms” (Adkins 2015, 58).⁶² Thus language has the “remarkable ability to organize and stabilise” (Adkins 2015, 81).

[...] the ability of language, with its own givens on its own stratum, to represent all the other strata and thus achieve a scientific conception of the word. The scientific world [...] is the translation of all the flows, particles, codes, and territorialities of the other strata into a sufficiently deterritorialized system of signs, in other words, into an overcoding specific to language. (Deleuze and Guattari 1987, 62)

Accordingly, are we able to organise the physical systems into ‘sediment floating in water’ and ‘sandstone’ (in a physical system, sedimentary particles are only more-or-less solidified). Yet words are not a representation of things (Adkins 2015, 75), as in the Saussurean semiotic distinction between signifier and signified, but an interaction between words and things as an assemblage, a “reciprocal presupposition”: “an assemblage of enunciation does not speak ‘of’ things: it speaks *on the same level* as states of things and states of content.” (Deleuze and

⁶² Deleuze and Guattari begin their ‘construction’ of the linguistic stratum and its mechanisms — in short, how language works — by considering it not as a ‘communicative tool’, but rather in language’s capacity to affect other strata. They do this through the identification of “order words,” considered not simply as imperative verbs but anything capable of, in contact with a certain assemblage, functionally altering system. Adkins gives the example of saying “I do” at a wedding (2015, 71) — which, unlike saying “I do” in normal creation, fundamentally alters the status of two people and creates a marriage.

Guattari 1987, 87) Thus the strata have a reciprocal relationship, rather than a hierarchical one: “[t]he result is a flat ontology.” (Adkins 2015, 58).

Taking this as a starting point, we can now look again at the previous models of transmission by remediation. Instead of transmission, though this is still the overall effect, we can now consider each stage as the interaction between two assemblages, resulting in a translation of a form from one form/substance complex to another.⁶³ This is illustrated by Figure 4.

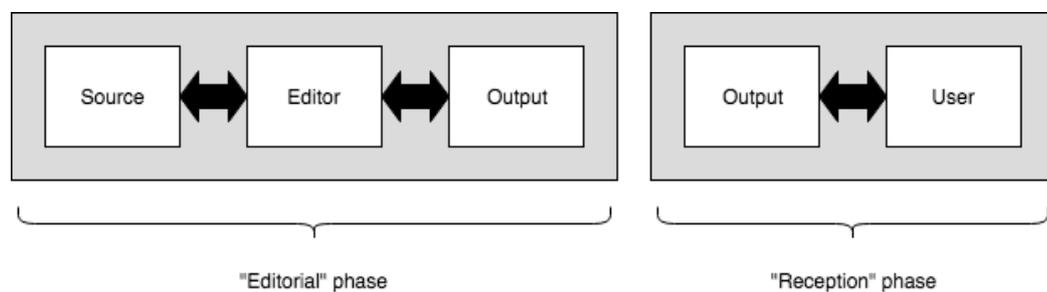


Figure 4: Transmission modelled as a series of interactions.

From the perspective of Deleuze and Guattari set out above, it may be said that the editor is an assemblage,⁶⁴ who interacts on one side with the source material and on the other with an output, with the aim of ‘translating’ the former to the latter. Before continuing, however, a few words must be said about the source. Considered as an assemblage, it can be seen as a

⁶³ *Translation* in this instance is a deterritorialisation (extraction) of a form (‘text’ – though it cannot exist without a substance) from one form/substance complex (the document: text-as-material), and its reterritorialisation on another, the ‘text’ (form) onto the substance of, say, the editor’s brain (neurons ‘encoding’ language).

Translation in this sense is, for Deleuze and Guattari, not a metaphor: it is a translation in precisely the same way as a translation from English to French is. In this case, the *form* is (loosely speaking) the ‘sense’ to be translated and the *substance* is the English language. The form is deterritorialised from English-*substance* and reterritorialised on French-*substance*. Of course, the form and substance of a sentence in English cannot be ‘truly’ separated (they are a singular entity). *Interpretation* — what the English ‘means’; really, what it *affects* (its capacity to affect) — is that which is *translated*. ‘Pass me the salt’ (an example Massumi, 1995, often uses) gets you the salt, as does ‘*Passez-moi le sel*’. (And yet, ‘*passez*’ carries the formality of the *vous* form that was not in the English...)

⁶⁴ This is, indeed, implicit in the conception of the editor made earlier in this chapter: as an assemblage of people, technologies, institutions. Even a single expert working alone is an assemblage of a person and pens and paper and domain knowledge.

complex system of components: the material surface; inscription as a series of marks made in ink (or stone and strategic absence thereof in the case of an engraving); the inscription as linguistic text; myriad meanings; intention; heterogeneous connections to author(s) and editors; bibliographic codes; connections to other material documents; and further connections to modes of production and distribution: publishers, bookshops; as well as readers, forming part of the social, cultural and linguistic background of its (multiple) reception(s). In this description can clearly be seen strong echoes of McKenzie and McGann, and also Sahle's model considered as intensity ('intention', for instance, not as a property of text, but as a connection to authorship).

From this, several points may be noted. Firstly, such an assemblage (a macro-assemblage of assemblages) has no clear boundary, and can be extended indefinitely: a printed document can be connected up eventually to Gutenberg and the invention of the printing press; paper to a particular tree, vellum to a particular goat farmer; the alphabet to the Phoenicians (everything, in principle, to the beginning of the universe). Evidently such considerations are absurd in practice, but at the same time, there is no self-evident line at which to stop: the editor must, entirely arbitrarily, call a halt somewhere. This, indeed, might be seen as the first step: to determine the boundaries of the source-as-assemblage, internalising connected components and transforming them into properties of the source, making the source extensive.

Putting it this way, however, is to mistake the reality of the situation. This is the second point: that editorial interaction is really with the source as material document (or set of material documents). This is the primary encounter: editor and document. Hence Gabler's insistence that, "it is documents that we have, and documents only." (Gabler 2007, 199) And, also, Sahle's description of scholarly editing as, "the critical representation of historical

documents.” (Sahle 2016, 23: my emphasis) The reason for this is, I think, obvious: the heterogeneous connections between the material document and other components that, together, have given rise to (and hence define, in some way⁶⁵) the document are historical connections that have been broken. Thus, while a document might contain traces of its history, we do not know about these without interpretation: “We know nothing about a body until we know what it can do, what its affects are”. (DeLanda 2006)

In this state — fully disconnected — the material document is simply material, without any intrinsic meaning at all.⁶⁶ Thus it may be argued, as Bordalejo (2010) and Robinson (2013) do to a limited degree, that the document alone cannot be the basis of reproduction — even when we ostensibly reproduce the document, or the text of the document. Any interpretative act, of which reproduction is one, involves the plugging-back-in of the document to a system. Any system. Some of the connections made can be at the most generic level, e.g. writing, a system of marks on a page that resolve to a particular language. But also semi-specific systems, such as general expertise. And also highly specific systems: a conceptual model of some of the historical contingencies that gave rise to the material document. This plugging-in is therefore an interpretative action; in fact, how it is plugged in, and to what systems, *is* interpretation. And, most importantly, it is always a subjective act: there is no pre-emptive, universally valid way of plugging in the document. As such, the determination of the boundaries of the source-as-semblage is really the effect of what decisions the editor takes in (re)assembling the

⁶⁵ Deleuze and Guattari’s aphorism that, “[t]here is no difference between what a book talks about and how it is made” (Deleuze and Guattari 1987, 4) pithily illustrates this (and, indeed, their entire perspective).

⁶⁶ As an illustration, take a vacuum cleaner that is disconnected from the electricity. In this state, it is merely a few pieces of plastic and metal joined together (functionally, in other words, not a vacuum cleaner). That we recognised it as a vacuum cleaner — and not plastic and metal — is due to the fact that it *could* be plugged in again to produce an effect (or, more precisely in this instance, had been plugged in in the past).

assemblage. It is this reassembling that gives rise to the kinds of stratification of the material document arising from its encounter with the editor.

2.5. Decoding and deterritorialisation

As a preliminary example, let us take a document inscribed with text in English that an editor wishes to transcribe. Sahle describes a simplistic understanding of transcription as, “reading written down”. (Sahle 2012)⁶⁷ Such a conception points to a very basic model of transcription: there is some text on the page; the transcriber reads it as language; and writes down the language again. As Sahle notes, this notion of transcription is inherently flawed; I agree, and will here give my own reading of the situation (to begin with, leaving aside the re-inscription of text on a new document, the output, and focusing on the first interaction: the interpretation of the document). The obvious problem with this description from the perspective presented here is that it takes *text* as a given. A document alone is a material system, a surface and some marks on it; it does not, intrinsically, contain text. If we take Caton’s definition of text as marks on a surface that “resolve to language in our heads” (Caton 2013, 212) we can see the marks as non-self-evident. Rather, text emerges out of an interaction between the material document and an interpreter (the editor). Without the interpretative action of an editor (or any reader), there are only marks on a page. It is the effect produced on the side of the reader (the ‘resolving to language’) that means we see the marks as text. McGann makes essentially the same point, noting that “[t]extual fields arise codependently with interpretative action” (McGann 2003, 3; 2010, 2)

⁶⁷ From the talk, “Modeling Transcription” at the *Knowledge Organization and Data Modeling in the Humanities* symposium, Brown University, 2012. (Transcript of the talk available online at: <https://datasymposium.wordpress.com/sahle/>) To be clear, Sahle presents this view of transcription as an overly simplistic starting point, and goes on to dissect its inadequacies and underlying assumptions.

We may consider this as a process of de/recoding and de/reterritorialisation. On the side of the document (the content), we have substance (ink) and form (particular shapes); in the interaction with the editor, the form is decoded and deterritorialized from the page; and is recoded as language and reterritorialized in the brain of the editor. In other words, a stratification. While it may seem like a one-way-street — a transmission from document to editor — according to Deleuze and Guattari this transmission is an effect of the process of ‘double-articulation’: not the identification of content, and extraction and reproduction of that content as an expression, but an encounter that produces an effect on both sides. On the side of the source material (the content), a simultaneous reterritorialization and recoding occurs — essentially, an overcoding of the physical marks by language. The marks on the page ‘become’ text out of the encounter with a reader. Remembering that content and expression are reversible, we see that it is only as a result of the overpowering force of the editor that we consider the source material to be the *content* and the language-in-the-brain-of-the-editor to be the *expression*. As such, the editor interacting with graphemes on a page is really overwriting (overcoding) them with language, setting up (however fleetingly) the correspondence between the two on the basis of interaction: ‘These marks on the page *really* represents a given word in English’.

To enable this interpretation, it is necessary that the document be ‘plugged in’ to some system, in this case, the English language. Returning to Massumi’s example of wood carving, we can see the ‘interpretation’ of the wood in interaction with a chisel. If the wood-carver brought along a sledgehammer instead, the wood would be ‘interpreted’ in a quite different way: as kindling. So too the above example: it is as if, metaphorically speaking, the editor has brought along an ‘English language chisel’ to the encounter, which stratifies — separates specific marks on the page from the remainder — on the basis that it can be interpreted as

English. (Therefore if, for instance, the page also contains a number of Chinese characters, they cannot be de/recoded as language, as the only interpretative tool is English.)

That said, even if a text is (ostensibly) written in English, it is not necessary to interpret it as English in order for this process to occur. A non-English-speaking editor whose native language used the Latin alphabet would be able to stratify the document without explicit recourse to interpreting the marks on the page as language: the document would be ‘plugged into’ the Latin alphabet.⁶⁸ The same logic may even be said of a photocopier (in its first, ‘photographing’ stage): it stratifies the marks on the page, extracting black beyond a certain threshold of ‘blackness’ and treating the rest as white (in other words, plugging the document into a simple system of *black relevant, white irrelevant*). As such, I would argue that there is no intrinsic difference between human interpretation of a document and ‘machinic’ interpretation: only the relative nuance of the stratifying forces they are able to bring to bear.

In these examples, we see the process of stratification. As Deleuze and Guattari observe, it always produces a pair of strata, distinguishing a signifying form from the ‘background’. It is tempting to view this as selection based on an intrinsic property — black text from white background — but this is to take an already-determined view of how things work. Rather, viewing things as an assemblage is to take what appears to be an intrinsic property and see it instead as an exterior connection to something else: either conceptual (‘blackness’, the Latin alphabet, the English language) or physical (for instance, people: the author, the editor, someone scribbling marginalia on the page). Thus what determines interpretation are the types of connection that the editor makes with external components; delimiting the text as assemblage *is* what causes it to be interpreted in a given way.

⁶⁸ Indeed, even when marks on a page are stratified as language, there is the potential to fall back to interpreting as alphabetic glyphs, for instance in the case of an unfamiliar acronym.

This process of stratification can therefore be seen (simplistically) as creating an equation, in which the whole is the sum of the two strata:

Material Page + **BLACK MARKS** = THE DOCUMENT

(where bold type indicates the extracted stratum and italics the discarded)

This is not to suggest that a document comprises a material page and black marks; rather, that a certain interpretation allows such a division to be made.

It is also the case that, given the objective of representation and transmission, that an editor necessarily connects up the material document to systems that would ‘interpret’ its components as *more extensive* (or at least equally extensive). This is important as it not necessarily the outcome for all connections that could be made: some forces can *destratify* an assemblage even further. Interpret a document with a cigarette lighter and see the result. We do not even have to resort to wanton physical destruction. ‘Interpreting’ a document by looking at it through a blurring lens will result in a less-extensive interpretation. Instead of discrete letter forms, the interpreter will see only indistinct smudges of black. As such, interpretation in a way that can be transmitted requires that it be transformed into a more, and not less, extensive set of components. This tallies with Buzetti and McGann’s view that the “basic critical method [of editing is] formalization.” (Buzetti & McGann 2007) As such, the suggestion in Pierazzo’s model of editing — that the editor “asserts” facts that the document is “made up” of (Pierazzo 2015, 53)⁶⁹ — may be reinterpreted. The ‘facts’ of the

⁶⁹ The quotations here are the types of the relations between entities in the Unified Modelling Language diagram presented by Pierazzo.

document do not pre-emptively exist as facts, but are created by the editor — not out of nothing, but out of a potential to be interpreted in a given way by a given interpreter.⁷⁰

To give some more examples of systems that the editor can plug in (or encompass within the assemblage being edited), we may take a ‘poetry’ system — which would attach particular significance to line-breaks, versus a prose system which would not. Or, more prominently in the field of textual scholarship, authorial intention may be seen as a system. This, in fact, can be interpreted in many ways, ranging from Foucault’s *author-function* (as suggested by Gabler, 2012) to the full-blown *final authorial intention* of the Bowers-Tanselle line.⁷¹ This suggests that ‘authorship’ must be itself interpreted as an assemblage, which the editor ‘divides’ by some kind of selection, and connects up certain component parts of authorship and not others. To illustrate this, take a word on a page that has been deleted by being struck through. For a photocopier — uninterested in authorship or even text as such — it is merely a set of black marks on the page, not distinct in any way from any other black marks. For an editor reproducing the text of the document without appealing to authorial intention, it is a word with a line through it. To plug in a degree of authorial action (perhaps closest to the *author-function*) it is a word that the author meant to write and then decided to change. In this sense, a quasi-neutral *author-function* Gabler suggests is really an appeal to — or connection to — wider conventions (in which striking through a text ‘means’ deletion⁷²). To

⁷⁰ Deleuze and Guattari describe this as *affect* — which, as DeLanda observes, is really a shorthand for ‘the capacity to affect and be affected by’.

⁷¹ It should be noted that it is not necessary to *presume* authorship (explicitly) in order to stratify, even though common sense would suggest that someone (or something) must have written it (a photocopier, for instance, knows nothing of authorship).

⁷² Or, to take Bordalejo’s example (2010, cited in Robinson 2013, and referenced in the Literature Review) from a manuscript of Dante’s *Inferno* — the word visually appearing as ‘d – u – r – a-with-underdot – o’ — that convention at the time marked deletion with an underdot (rather than it being a curious diacritic mark). This again illustrates the nature of the author as an assemblage, ‘plugged in’, as it were, to wider conventions. It is never possible, therefore, to purely separate out these components. Indeed, as Robinson notes, even if one did not know about the underdot-means-deletion convention, simultaneous connection to

take a final authorial intention perspective, the struck-through text is interpreted as not being there at all. In this respect, therefore, interpretation is really an accumulation or synthesis, of different connections — or, put another way, connections are between things that are themselves assemblages, with their own connections between them (e.g. between authorial intent and writing conventions).

The above examples are focused on interpretation at a level normally considered within the realm of textual scholarship (especially authorship), but the same ‘plugging in’ of documents is required even in the case of practices that are more obviously ‘digitisation’ (including image-based reproduction): interpretation-as-connection operates down to the material document itself. Even the physical boundaries of the document cannot, ultimately, be seen as self-evident, only sufficiently ubiquitous as to appear as such (any connection to a ‘seeing assemblage’ or ‘feeling assemblage’ can interpret the edge of a document). But this does not stop here, and many other aspects of the material document are even less self-evident. For instance, which way up is a document? In theory (let us use the example of a piece of A4-sized paper), knowledge of paper production — and its habitual use — would suggest that a straight edge would be at the top (i.e. either portrait or landscape, *not* rotated 217° clockwise). If the document could be connected up to a language-system or even grapheme-system, the ‘correct’ orientation could be determined. In the case of cross-hatched writing, there are two distinct orientations (with portrait-text and landscape-text); Japanese newspapers often contain articles written both horizontally and vertically (in this case, the orientation of the characters tells us nothing; only interpretation as language determines the orientation of the page). In these examples, we see the mutual implication of page and text

the Italian language would cause alarm bells to ring, as ‘duraio’ is not a word in Italian. But, if one were to stratify purely on the basis of identifiable graphemes, one would produce ‘duraio’ — and it would be quite legitimate.

(and convention) in determining the orientation of the page; the physical page does not have a right way up intrinsically, but only insofar as the whole material document itself is an assemblage. Thus the orientation(s) of the page are the result of being able to stratify significant orientations from an infinite number of possibilities. Moreover, it was necessary to be able to (at least partially) interpret the text to do so. As such, even ostensibly objective reproduction such as photographing a page, is dependent on a fleeting interpretation of the text-as-language (or –as-graphemes) to orientate the page.⁷³

It should be noted, moreover, that in the above analysis we are dealing with a single, physically extensive document. This is seldom the case in any editing or digitisation activity, where a text may be spread over multiple pages of only marginal physical extensity. To take one example, we may have a page that has been ripped in half. Interpretation of the two halves as a single page involves the connection of one to the other (and a plugging-in of the two to an external mechanism: a conceptual ripping-assemblage). Documents bound together depend on an interpretation of the pages through connection to a book-binder (a binding-assemblage), or possibly some kind of linguistic-assemblage (or even authorial assemblage) that could determine or account for the consistency of the text across the pages.

At the most complex level, two documents plugged into each other on the basis of textual similarity (both, as typically conceived, instantiating the same work). This could be interpreted in multiple ways. A purely textual stratification would create two documents containing similar texts, with no way of distinguishing between them (the kind of rationale

⁷³ In a similar vein, we may consider the stratification of a document by a hyperspectral scanner. In theory, this generates an infinite number of layers (the electromagnetic spectrum is, after all, a spectrum, not a series of fixed points). But that a layer be considered as such is the result of a particular wavelength *and* amenability to stratify from the layer revealed by that wavelength some text or other sign.

that produces a variorum edition). Conversely, connections to a fully-fledged final intention-wielding conception of the author would stratify not just the text from each document, but one reading of the text from two possible readings, based on whether it was the author who was responsible for a particular reading, and then whether that reading was the final one intended by the author.

The complexity of this final operation would appear to elevate it far above the more mundane activities of transcribing graphemes, or sticking a document in a photocopier. Indeed it is, if only because it requires myriad connections — not just to the author, but to language, handwriting, the complex interplay of compositional activities, editing, revising, and printing — to determine or stratify the text in such a manner. But, for all these difficulties, I would suggest, it is the operation of the same abstract mechanism: it is the formalisation and extraction of codes from material documents, based on some criteria, and their recoding in some appropriate form in the mind of the editor.⁷⁴

A relationship between the source and the editorial interpretation is therefore established: one of backwards-projection (or overcoding, to use Deleuze and Guattari's term) of a set of determined forms onto an indeterminate system in which the determined forms were only a possibility.

Most importantly, this formalisation is entirely dependent on where the editor chooses to place the boundaries of the text-as-assemblage; on what interpretative systems the editor

⁷⁴ Or, to be clear, from the mind of the editor to an 'intermediate' piece of paper; or from camera-sensor to computer memory, or any other intermediate medium. The editor, as suggested earlier in this chapter, is an idealised description of a complex assemblage. In this analysis, the internal mechanisms of the 'editor' are, I admit, glossed over. This is not to say that they are irrelevant, but that from the mechanistic conception described here, the editor can be considered as a 'black box', taking the source as input and producing an output.

chooses to plug the material document into. And no given system is self-evidently the ‘right one’. An editor may experiment iteratively; the only test of validity (and then, only a weak validity in a global sense) is whether a system is able to stratify. Just as a woodcutter might find that a blade made out of custard makes no ‘interpretative’ cut into the wood, so too an editor attempting to plug an English text into a French-language system: it produces no effect; it is unable to formalise. But to carry out this stage of transmission, some stratification must be made: even if it is, in the case of a photograph of the page, between the physical piece of paper and the photons bouncing off of it, captured and recoded.

In the above discussion, we have covered the first interaction in the process of transmission presented in Figure 4 (see page 125). The second interaction concerns the recoding and reterritorialisation of the forms extracted onto a new document.

2.6. Recoding and reterritorialisation in the output

The second stage of transmission comprises the interaction between the editor and the output. This is, again, an encounter between assemblages. As such, it is subject to the same logic of double-articulation as the source–editor encounter, though with a different power relation. To briefly characterise this stage in terms of Sahle’s description of transcription as “reading written down”, it is the ‘writing down’. Though it is situated in the transmission diagram above (Figure 4; see Page 125) as following the first interaction, the source–editor encounter, it is only conceptually (and pragmatically to a degree) second; the two may occur largely concurrently.⁷⁵

⁷⁵ As suggested earlier, this is an ideal representation of the interactions. I am not saying that an editor needs to take in and memorise every detail of the source material *before* recording the whole in the output — though it is possible, in the case of more mechanical reproduction (e.g. a photocopier). But, though these two interactions may operate iteratively and

In the terms of this analysis, the editor recodes and reterritorialises the abstract codes (now floating round in the brain of the editor, or on intermediary pieces of paper or computer memory) onto the output. This creates an output — a derivative document, in Dahlström’s words — that is stratified *in the same way* that the editor had stratified the source material. The output is thus an encoding of the editor’s determination of the source’s inherently indeterminate system.

This does not mean that there is a natural or self-evident way of recoding the abstracted codes of the source; rather, it is a recoding in what Sahle refers to as the “target system” (Sahle 2012) — the formal representative capability of the output. This, evidently, varies with the medium. To take a simple case, the editor who re-codes and reterritorialises the textual signs in his head with only a typewriter must rely on the twenty-six letters of the alphabet (upper and lower case), ten digits, a limited array of punctuation and skilful manipulation of the carriage return bar. A printed page that is typeset has similar (albeit fewer) restrictions; a printed photographic reproduction again fewer restrictions. Digital representation can make use of full-colour, highly detailed images; or use complex technical systems of encoding — such as TEI-XML — to descriptively mark up stratified features of the text. Evidently, the digital medium places fewer restrictions, but they still exist: no (sane) reproductive medium can currently re-code the texture of the original document, or its smell.⁷⁶

concurrently, at some point the editor must start with the first interaction (what would they be writing down otherwise?)

⁷⁶ Or rather, it is possible, but only with recourse to language — a non-formal system. The informality of the description, however, makes it difficult to connect in a meaningful way to, say, the formally encoded text: “This bit of page smells of coffee” — a useful fact, if the editor deems it — does not alter the text that does smell of coffee. What we see here is a *coffee-smell—text* assemblage that stratification into *coffee-smell* and *text* only destroys.

The ability to formally re-code the extracted strata in the output is (as will be seen shortly) a necessity. That it is a requirement, though, points to a necessary connection to the source material that the previous section omitted: the material source document must be plugged into the formal system of encoding available in the output. Depending on the relative similarities of the source system and the output system, this might be a trivial matter. It is easy to see how the textual codes of a document written on a typewriter might be recoded on another typewriter. However, the typewriter-as-target-system cannot encode other aspects of a document, for example illustrations or large coffee stains. Therefore, a connection to the output-system is required, and its effect summed up in a general rule: the editor cannot stratify the source material in a way that the extracted stratum cannot be encoded in the output-system. Or, rather, the editor might stratify the source in such a way, but it must eventually be discarded (or described in prose: see Footnote 76) for want of a way to actually encode it in the output. One way of looking at this is to say that certain aspects of the source material cannot be represented; but to talk of ‘aspects’ is to imagine an already-determined system. Rather, the source system itself can only be determined in such a way that allows its representation in the output.

Up to this point (and taking into account the user-reception stage described in the next section), the representation of the source by the output can account for a baseline level of transmission — what might be best described as *naïve transmission*. It cannot be denied that, at a very prosaic level, something has been transmitted to the output (which the user, in turn, can extract from the output). That is to say, some aspects of the source have been extracted, and they have been recoded and reterritorialised (somehow) in the output, from whence the user might in turn extract them. Something, in other words, has been copied. To say that

this is not transmission is therefore too strong a statement. Rather, it is a transmission whose reliability is open to all sorts of questions. Hence the description as *naïve*.⁷⁷

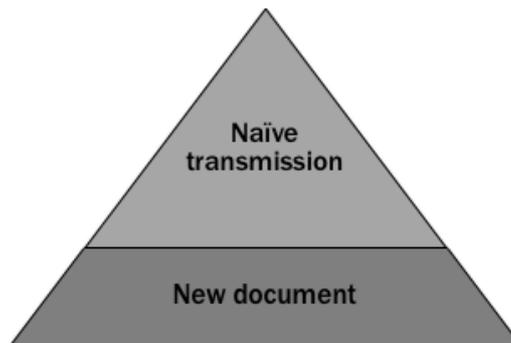


Figure 5: *Naïve transmission* — predicated simply on having been copied — distinguishes the result from an entirely new document.

The reason for regarding this in such a way is that, to borrow terms from information studies, there is no way to distinguish the *signal* (in this case, the elements of the source that the editor wishes to transmit) from *noise* (everything else about the output). This indistinguishability arises because the output is also full in all its dimensions — it, too, is a material document, which the user must stratify in some way in order to extract the transmission. (This third interaction will be dealt with more fully in the next section.) Who, therefore, is to say — given only a ‘naïve transmission’ — which aspects are those originating in the source material, and which are those native to the output? It is true that much of what might be seen as naivety is actually a ‘non-naïve’ appeal to convention: the user of, for instance, a print scholarly edition knows perfectly well that the binding and the paper are not the originals, and accordingly reads the text as signal and ignores the texture of the pages as noise.

⁷⁷ In fact, it may be argued that this alone does not amount even to naïve transmission, on the basis that, confronted with the output alone, the user knows nothing about its having been copied from a source: it is a new document, and thus a completely ‘original composition’ for all anyone can tell. But this can be solved by appending the statement, “This is a copy of *Document X*” (or words to that effect). And, of course, a more honest description would be, “I copied something from *Document X* onto this new document, but I’m not going to tell you what” — hence the naivety of taking it as a transmission.

But this rather points to a general idea of what it takes to make transmission ‘non-naïve’ (what I will call an ‘articulated transmission’): knowledge on the part of the user of what to consider signal and what to consider noise. This knowledge may be in the form of wider convention, or it may be explicitly communicated as an adjunct to the output. From the perspective set out here, what this knowledge amounts to is a description of the editor’s encounter with the source material: i.e., if the user knew — because it was convention, or because they were told — what it was an editor did to stratify the source material (which includes, necessarily, how it was to be recoded in the output; see above), the user would be able to work backwards and determine that part of the source material which was transmitted. In the previous section, we described a particular stratification (of writing from the background page) in the form of an equation:

Material Page + **BLACK MARKS** = THE DOCUMENT

(where bold type indicates the extracted stratum and italics the discarded)

The output thus creates the same equation, which can be ‘solved’ by the user by subtracting the material page. In such a description may be seen echoes of Eggert’s attempt to theorise the literary work by recourse to Adorno’s negative dialectic: “[T]he work’s documentary and textual dimensions dynamically interrelate [...] They are, in this sense, one another’s negative constituting principle.” (Eggert 2009, 234-5) But this can more generally applied to any system that can be stratified.⁷⁸ In Deleuze and Guattari’s conception of stratification — where

⁷⁸ While I am, generally, inclined to agree with Eggert’s view, I think one criticism can be made. Specifically, there is necessarily a leap between saying that the documentary and textual dimensions of a work *interrelate* (undoubtedly true) and saying that they are “one another’s negative constituting principle”. The latter requires the presupposition of a particular dialectic — in other words, an interpretation. Eggert’s combination of this view with a Peircean semiotic (Eggert 2009, 234-6) suggests just the requirement of an interpretant to

strata always come in pairs, a division of a system into two ‘layers’ — the strata formed can indeed be seen as “each other’s negative constituting principle”.

Alternatively, we may articulate the same logic as a kind of functional equivalence. What being able to extract the signal from the noise in the output enables is the following kind of reasoning: the signal in the output is what the user would extract from the source material if they too were able to access it, and had the editor’s expertise, and plugged it into the same systems... i.e. if they stratified it in the same way, this would be the result. The output is, therefore, the record of the editor’s encounter, articulated.

This constitutes, I think, a second critical distinction in transmission, between *naïve transmission* and what might be called an *articulated transmission* — a transmission that gives an account of itself, and thus can be used to experientially reconstruct the encounter with the source. By articulating the way in which the source was stratified, the editor overcomes the contingencies of an invariably interpretative act, by describing those contingencies so that they can be subtracted.

realise this. Accordingly, what ‘aspects’ end up being one another’s constituting principle is hardly universal — and only a particular view of a ‘work’ can give rise to the two component parts he suggests. There are, for want of an analogy, an infinite number of ways to cut a circle into equal halves.

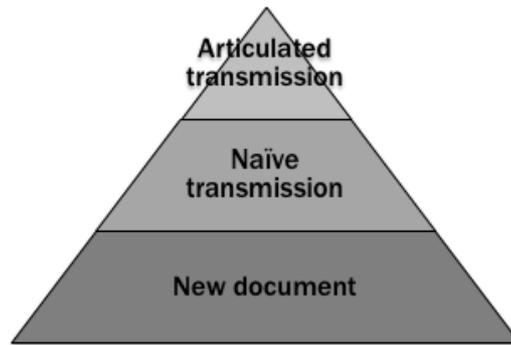


Figure 6: Articulated transmission as a subset of naïve transmission and documents.

2.7. The User–Output Interaction

Though the previous section has touched on this implicitly, this section will deal in more detail with the third interaction in the transmission diagram (Figure 4: see Page 125). As already observed, this interaction is — unlike the first two that may occur largely concurrently (in the ‘editing phase’) — separated in time and space from the interaction that created it (the ‘reception phase’ in the diagram). As such, the output at the point of its reception is not dissimilar to that of the source material at the point of its interaction with the editor: it is an assemblage, whose external connections (especially to its creation) have been severed. Likewise, the user–output interaction obeys precisely the same mechanism as the source–editor interaction: a double-articulation, in which the user stratifies the output, overcoding physical signs with language or another formal representation. As such, the output too must be seen as an indeterminate system, capable of as many determinations (interpretations) as there are systems into which it can be plugged. Given this starting point, there is no reason, objectively speaking, why it should not be subject to any kind of interpretation, and hence no reason it should transmit or represent the source.

Clearly, however, this cannot not be the case. In the above section, we began to look at mechanisms for avoiding this retreat into unconstrained interpretation. The first is that the editor provides an account of how they stratified the source material, and how its forms are recoded in the output — thus the user’s first stratification of the output is between signal (the transmitted forms) and noise (the ‘accidental’ forms of the output). As Sukovic says, “Documentation about representation processes and detailed metadata are often needed to demonstrate that the representation and represented are identical in every important way”. (Sukovic 2008; Pierazzo 2014, 111, makes the same point explicitly.) Put another way, what allows the user to read the output ‘correctly’ — by which we may take to mean, ‘as the editor intended’ — is the identification and plugging-in of the output document to the ‘correct’ (the intended) systems. Unlike the source, which does not mandate any particular interpretations as valid, the output can be explicit about the correct way to interpret it (or, as suggested above, the editor can appeal to convention *in lieu* of being explicit). This can be considered, therefore, as a kind of contract between editor and user: use the output in a specific way and you will get the appropriate result.⁷⁹

As a corollary of this, the forms encoded in the output need to be *extensive*; and for this to be the case, the editor must stratify the source into extensive forms (see above). Intensive assemblages, which lack of extensity implies, are once more subject to any kind of interpretation, while extensive forms are already interpreted. They must also — this is really

⁷⁹ An analogous scenario might be seen in the case of flat-pack furniture. Once it has been sold, Ikea (or whoever) relinquishes all control to the purchaser. Now fundamentally, a flat-pack bookcase is just wood and bolts, and could be assembled into any kind of mess by hitting it hard enough. Assembly instructions act as a kind of contract (not in a legal sense) with the user: build it in the right way and you’ll have your bookcase. (In Deleuzian terms, the flat-pack bookcase is *real*, but *virtual*: it is real as you can buy ‘a bookcase’ — even though it cannot at the point of purchase hold any books — and contemplate its size in relation to your collection; it is building the bookcase that actualises it. Or, as Drucker might say, you need to *perform* the bookcase.)

two sides of the same coin — be formally recoded in the output. This is necessary so that the editor can refer to the formal representation as a whole, i.e. by pointing to the system of encoding in their description of their stratification of the source. Without the output being subject to formal systems, there is only ever an intensive assemblage, and thus no way to proscribe to the user a ‘correct’ interpretation.

In the above sections, a model of transmission has been elaborated, from source to user, via three translations: from source to editor; from editor to output; and from output to user. Transmission, as we have seen, is the forward-propagation of information. It starts with the source and ends up with the user. That the transmission is representative, however, relies on a backwards series of projections. The linguistic (or other) recoding in the user’s head is a specific, already-determined overcoding of the signs present in the output; the output is an encoding of the signs created in the editor’s head; which is, in turn, an overcoding or specific determination of the source. This latter backwards-projection is the most important, as it means that the output represents the source only insofar as the editor determines what the source itself *is*: in other words, what particular assemblage of the source is (re)assembled by the editor. Thus to say the output is the editor’s text, as many have argued (see Pierazzo 2015; Gabler 2012) is quite correct.

This conception of transmission presented here may be called *articulated transmission*. It is critically distinct from the earlier *naïve transmission* insofar as it gives an account of itself. By describing *what the editor did to create it* (how the source was stratified) it is able to give a positive account of the connection between that which it presents and the source material. (*Naïve transmission*, on the other hand, takes it for granted that ‘copying’ automatically leads to representation — somehow.) It occupies, therefore, a level of validity as a representation over and above that of naïve transmission.

There is, however, a further critical distinction that can be drawn (a new pinnacle to the pyramid). Indeed, in one sense there must be: taking articulated transmission as the ‘peak’, as it were, suggests in particular that scholarly editing cannot be distinguished from any transmission that simply articulates its process. (See Figure 7, below.)

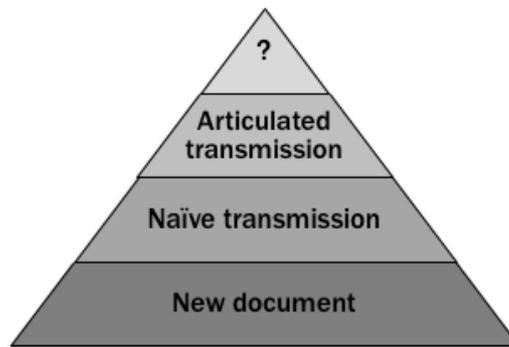


Figure 7: A missing threshold

Indeed, Harvey’s three rules for the editing of documentary texts do little to elevate the result above this conception of articulated transmission: the second, “Say what you are going to do and do it” (Harvey 2001, 13) seems, in fact, to endorse a lack of distinction. Pierazzo’s two models of scholarly editing (see Pierazzo 2015, figures 2.2 and 2.3), in their description of the processes and connections between actors and entities, seem also to lack any mode of distinguishing scholarly editions from articulated transmissions. It is not, therefore, that I am in disagreement with either Harvey or Pierazzo, but that I would nevertheless contend that their conceptions are insufficient to account for scholarly editing, and that a further distinction is required.

The basis of this distinction is an articulated transmission’s giving only a positive account of itself: what the editor did, not all the things that the editor did not do. As we have seen,

editorial interpretation involves plugging in the source document to a variety of systems. But there is no intrinsic reason for choosing which systems. The source, to paraphrase Deleuze and Guattari, can be interpreted in as many ways as there are systems into which it can be plugged (see Deleuze & Guattari 1987, 4–5). Even granting that one of the systems must be the formal system of the output (see page 138) — which alone is not sufficient to stratify the source — or obviously destructive systems such as fire, there are myriad (actually infinitely many) other systems.⁸⁰ Thus, why select specific systems? Why not, after all, plug the manuscripts of *Ulysses* into a paper-aeroplane system? The manuscripts are paper; they would fold. (Obviously, because it is manifestly daft; but there is no essential delimitation of daftness from sense. The question is, rather, why are some things daft?) In general, though, it must be the case that there can be no intrinsic reason — *a propos* of nothing — to connect up the source to particular systems, and thus to interpret it in a particular way. As Sahle has argued (Sahle 2013) there is no objective line to be drawn between what is relevant and what is not.

Articulated transmission, therefore, gives an account of what it is, but without an overarching rationale of *why*; it gives a contingent account of itself, and those contingencies rest, ultimately, with the editor. Essentially, all this does is move subjectivity down a notch: it is an objective account of a subjective operation.

This subjectivity exerts itself because the output is a closed-off system: it must be, in order to be packaged up and sent out into the world to be utilised at a later date. This is not to doubt that any particular component of the source may not be represented, because the source components have been made extensive components, but the entire output, considered as a whole, cannot be said to represent anything beyond the editor's decisions. An articulated

⁸⁰ This is not like the formation of sedimentary rock, where the only force to hand is the downward pressure of water and other layers of sediment.

transmission of *Ulysses* manuscripts transmits components of *Ulysses* (and verifiably), but the transmission as a whole does not necessarily represent *Ulysses* as a whole. While it therefore may be able to account for all its constitutive contingencies, it does not address the biggest contingency: the contingency-of-the-whole. It is, to adapt a phrase, *the truth* and *nothing but the truth*, but there are no grounds for saying that it is the *whole truth*.

There are, it should be noted, two ways out of this impasse. The first is to avoid subjectivity by completely representing the source. The logic of this is outlined by Dahlström:

If it is possible to ‘completely’ identify, formulate, and unambiguously encode every aspect of an artefact, say a document, into that of another document, then by logic the target document ought to be in every aspect equivalent to the departure document. And if it is indeed equivalent, it follows that to a user it is of no importance if she/he gets his/her hands on the one or the other. (Dahlström 2004, 24-5)

But such a notion is easily dismissed: it is what Dahlström describes, borrowing terms from McCarty, as the ‘mimetic fallacy’ and the ‘complete encoding fallacy’. (Dahlström 2004, 24) Pierazzo makes a similar point: “no edition will ever be able to account for all possible facts of a work” (Pierazzo 2015, 56). Likewise Shillingsburg: “[n]o edition is a full representation of that which it attempts to edit. No edition was ever or will ever represent a work adequately.” (Shillingsburg 2006, 154) From the conception of the source as assemblage presented here, it is easy to see why. Not only can the source be plugged into a limitless number of systems to interpret it — thus ‘completeness’ is an infinite number of interpretations —, but some of these are mutually incompatible (one cannot claim that *Ulysses* is a novel *and* a paper aeroplane: the latter implies disregarding the textual inscription, for one thing). Complete encoding is, quite rightly, regarded as a fallacy, and the editor cannot achieve it as a way to avoid subjectivity.

The second approach, therefore, is to appeal to some ‘correct’ interpretation, which would relegate the ‘incorrect’ interpretations (from the legitimate — interpret the text on the basis of authorial intention; to the absurd — *Ulysses* is a paper aeroplane) to irrelevances. Thus Dahlström’s logic would apply: ‘complete’ would, instead, be read as: ‘complete within the bounds of relevance’. But this obviously creates a problem: who decides, and on what basis, what is relevant? What is needed is, therefore, both a regulator of interpretation, and a source of it.

2.8. Constraining Interpretation

The previous section established the necessity (if what we consider scholarly editing is to be distinguished from any arbitrary, albeit rigorously carried-out and articulated, transmission) of some fundamental constraint to possible interpretations. The problem is not so much that that interpretation cannot be regulated: an editor should create, to cite Pierazzo, “informed, circumspect, documented, scholarly interpretation.” (Pierazzo 2011, 466) In particular, ‘informed’ and ‘circumspect’ point to a highly admirable degree of self-regulation on the part of editors. It is not my argument here that editors make outlandish claims (such as *Ulysses* being a paper aeroplane); they do not. However, I would suggest that it is not possible to draw a line, universally, between what counts as outlandish or not; or, more sanely, even within the scholarly realm of informed circumspection, what counts as a ‘correct’ interpretation.

The question, ‘on what grounds do scholarly editors constrain their interpretations’, is therefore a legitimate one. In other words, what do editors appeal to for regulation? This matters because, as outlined above, it is the *regulation* that is really represented by the output of transmission.

One such possible regulator is the author's intention (formulated in any way the editor might choose). While this does indeed function as a regulator — it forces the editor to only plug in systems that correspond to those 'plugged in' by the author — it is unsatisfactory as the *source* of regulation for two reasons. Firstly, it is not necessary to edit according to authorial intention, and many critics (see in particular McKenzie and McGann, above) have questioned the validity of attributing the work entirely to the author. Secondly (which really flows from the first), if it is not necessary, then plugging in the material document to authorship is really only one possibility among many. *Authorship* is then not an external regulator, but as much a part of the source-as-assemblage as any other system: as much an editorial decision, in other words, as any other.

Another possibility is what might be termed 'social agreement' as a position on which to found an external, objective regulator. As Huitfeldt says:

Thus, we may conclude that there is such a thing as objectivity of interpretation. The vast majority of decisions we make in this realm are simply decisions on which all (or most) competent readers agree, or seem likely to agree. (Huitfeldt 2007)

This kind of view is in line with Robinson's suggestion that 'social agreement' is the basis on which, to use his example, an 'i' is an 'i': because we agree it is an 'I', not because it is a line with a dot above it. (Robinson 2009, 44). At the basis of this argument lies an acceptance of the letters of the alphabet being part of an agreed-upon semiotic system. As Pierazzo argues, however, just because lots of people agree on it does not mean it is objective. (Pierazzo 2011, 466) The *letters of the alphabet* is, like authorship, a system into which the document can be plugged — not a universal.

The same can be said of appealing to degree of expertise. Many critics have pointed to the notion of a ‘competent reader’ (including Huitfeldt, above, but also Pierazzo and Bordalejo). According to Bordalejo: “[W]hen I refer to a reading agent, I am talking about a person who has a high-level understanding of the language of the text and who has reached an adult level of reading comprehension. Such a person, I call a competent reader.” (Bordalejo 2013, 67) While, of course, expertise is a necessity in scholarly editing, this does not mean that any degree of competence of expertise is objective. In this regard, it is really ‘social agreement’ at a smaller, more exclusive level.

For competence alone to be sufficient, a ‘general’ competence would have to function in all specific cases. This would imply that all texts are identical, and self-identical. In the former case, expertise in one thing would count as expertise in another (but then, where did expertise in the first thing come from?) Or, in the latter, the expert would already be an expert in that specific material document (and where did that expertise come from?) Even then, it is evident that different texts are not identical to each other, and as Buzzetti and McGann cogently argue, nor are texts self-identical. (Buzzetti and McGann 2007) The consideration here of text as an assemblage likewise implies as much. Thus, the question arises: ‘Which particular competences is the editor employing in this instance?’ and we return once more to a subjective decision. (To shift things up a level, we might posit a higher, universal ‘meta-competence’ that would regulate the choice of specific competences: but this is merely to deny the objectivity of the source material, only to reassert it at the level of a universal competence.)

The above might be seen as a succession of straw-man arguments, but they are all suggestions that have been made for an objective basis for editing. In rehearsing the arguments in this

way, however, I hoped to show that they are not criteria that underpin transmission (specifically scholarly editing) as a universal, but arise concurrently with interpretation — in other words, that each represents a system into which the source document can be plugged. In trying to suggest otherwise, we keep hitting circular arguments, which suggests that the answer has been there all along. We have been trying to find some regulator that is external to the complete system of interpretation (the source-as-assemblage), while simultaneously arguing that there ultimately *is* nothing external to interpretation.

What is being pointed to, ultimately, is the lack of a Platonic ideal. As Bryant quite rightly suggests (in relation to Tanselle's view of literary works), the work "has no tangible presence, but this is not to say that it is not real, nor that it is ... some sort of Platonic ideal." (Bryant 2002, 31) In one sense (and as reading Deleuze and Guattari would certainly suggest), this must be the case as there is no Platonic ideal: there is no ontological separation between the real and the ideal, only a "continuity". (Adkins 2015, 4) While positing a Platonic ideal undoubtedly solves the problem at hand — as a transcendent idea, a "universalized abstraction" (Bryant 2002, 31) it can objectively regulate — its positing as a regulator creates more problems. Prosaically: where is it? How does the editor tap into it? More worryingly, it must precede the author's writing of the work in the first place. (Authorship itself is thus relegated from a creative act to the imperfect writing-down of a pre-existing perfection.)

But if the work — or any other determination of the source material — is ideal in a 'real' sense, as suggested by Bryant, then, again, where is it and how does it act as a regulator? The 'ideal' version would be nothing less than the source-as-assemblage already- and unambiguously-determined: in other words, an already-edited scholarly edition, somewhat boldly claiming to be the, "Real [I]deal". But there cannot be such a scholarly edition, as its own creation would be equally inexplicable. The problem, in its essence, is this: the editor

must create (subjectively) the ideal determination of the source, while simultaneously appealing (objectively) to that ‘ideal’ to regulate its creation. However, I would like to argue here that this is, in fact, the solution.

The basis of this argument can be seen elsewhere. It constitutes what McGann (citing biologists Maturana and Varela) terms an autopoietic system: “All components of the system arise (so to speak) simultaneously and they perform integrated functions.” (McGann 2003, 8) In such a description, we clearly see echoes of the definition of an *assemblage* suggested earlier. Similarly, in analysing the nature of the literary work, Eggert follows a similar set of arguments to that above. Invoking Peircean semiotics, he argues that the work itself does not function as a sign without an interpretant. He eventually concludes that the work itself is not an “Ideal,” but is rather only what he terms a “regulative idea”. (Eggert 2009, 232, 236) I think this logic, and indeed the term itself, functions over and above the case of the literary work: it may encompass any ‘idea’ that gives rise to a particular determination of the source material document (a *work* is one such determination). I should, therefore, like to borrow the term — regulative idea — and extend it beyond works. In the rest of this chapter, I will discuss the construction of regulative ideas, and their functioning.

We may approach the origin and functioning of a regulative idea from two directions (which should not, after all, be surprising given its inherently circular nature). As a top-down regulator it determines into which systems the source-material documents are plugged at the point of their encounter with the editor. As such, it constrains interpretation, putting an end to the interpretative free-for-all that could otherwise take place. In doing this, all that is really happening is the delimitation of the source-as-assemblage: taking a system into which anything and everything could be plugged, *ad infinitum* (and, indeed, *nauseam*), and producing a closed-off system. It is, therefore, an editorial choice like any other, but elevated

a rhetorical notch or two: a “We hold these truths to be self-evident...” that serve as a starting point from which a republic and “government of the people, by the people, for the people” logically flow. The “we hold” is important: it is not that these truths *are* self-evident (others may disagree),⁸¹ but that *in holding them as such*, a stable basis is created from which to proceed.

From the above, it is tempting to view the regulative idea as simply an editorial rationale — which it undoubtedly is, in a sense. But simply treating it as such places no constraints on what the components of that rationale might be (any *articulated transmission* has a rationale — an explanation). Rather it is a particular kind of rationale, a rationale “carved out of the irrational”, “defined by a particular kind of relationship among irrational factors.” (Deleuze 2004, 262: see the epigraph to the chapter) We might, accordingly, say that it is the rationale that underpins the editorial rationale (or the rationale that underpins that... and so on until one hits a single rationale, the *regulative idea*).

The above describes the downward function of the regulative idea; now we must turn to its upwards construction *from* the source material. This is a process of editorial experimentation. The editor repeatedly attempts to plug the source material document into systems and sees whether that produces, firstly, any stratification at all, and, secondly, a rational one. As an example of the first: the editor might find a document with some text on it, and plug it into various linguistic systems (English, French, Mandarin...) to see whether it is capable of stratifying. Assuming (pre-emptively, of course) that the text were in English, the editor would find that the French or Mandarin systems produced no effect. Returning to Massumi’s

⁸¹ The Royalist side in the English Civil War presumably held exactly the opposite (the Divine Right of Kings) to be ‘self-evident’ — leading to different conclusions (to uphold the monarchy).

example of wood-carving, this can be seen as a trying-out of various chisels and finding out which were capable of carving the wood. The second criterion, that it produces a *rational* stratification, is more difficult: rational in this case really means a compatibility with other systems that, applied alone, would also produce a stratification.⁸² Thus *Ulysses-as-a-paper-aeroplane* is ‘incompatible’ with the fact that the English-language system produced a stratification into material document and linguistic text: insisting on the paper aeroplane scenario would mean disregarding the text. Indeed, English *and* Mandarin are mutually incompatible. Similarly, in the case of wood, chiselling a desk is incompatible with bludgeoning it with a sledgehammer:

Artisans and craftsman [in Deleuze’s view] understand this other conception of matter and form, at least implicitly: they tease out a form out of an active material, collaborating with it in the production of a final product rather than commanding it to obey and passively receive a previously defined form. (DeLanda 2006)

Building a regulative idea is an artisanal activity. It is the experimental plugging-in of systems (of making them interact), slowly discovering what is consistent, that eventually delimit the source-assemblage being edited.⁸³

⁸² This, incidentally, implies that a *single interpretative system* applied alone is *always* valid, insofar as there is nothing to contradict it. This notion will be explored more fully in the next chapter.

⁸³ This internal consistency also, I think, can be seen as excluding any arbitrarily conjured systems the editor chooses to impose. The most obvious example of this is to assume that the document text — let us imagine it a novel-length — is actually some form of code. Undoubtedly, given enough text, by stratifying the document on the basis of every *n*th character the editor could make the text say whatever it wants (this does not stop it being the plugging in of a system that produces a stratifying effect). The so-called ‘Bible Code’ (Drosnin 1997), which ‘uncovers’ hidden codes in the bible by doing just this, predicts, among other things, the Kennedy assassination. While, on its own, this must be regarded as a legitimate stratification of the document — why not? — it is its incompatibility with other systems (the history and composition of the bible; its status as a religious text; the improbability of predictions made a couple of thousand years in advance) that makes it stand

Given that the above suggests experimentation as the only mechanism involved, it is necessary to return to the question: *whither expertise?* Earlier in this discussion, I ruled out expertise as an objective hook on which to hang the regulative idea, but this does not make expertise irrelevant. Instead, the notion of experimentation allows a better account of expertise: not as a universal, but as a means to shortcut unlimited experimentation. Put another way, expertise is the result of already having carried out a degree of experimentation. While it is true that, logically, total expertise *in the specific source material at hand* is impossible (as argued above, this essentially means the source material is already edited), expertise allows the incorporation of other, similar material documents — and how they might be connected — into the source-as-semblage. An expert in medieval Latin manuscripts and their creation can ‘plug in’ wholesale the knowledge gathered by having previously edited hundreds of other such manuscripts, into the specific source document they now wish to edit. Expertise can be seen, therefore, as the packaging together of suitable systems, so that they may be selected *en masse*. This does not, therefore, rule out experimentation (the editor must check that it works); it merely streamlines an otherwise laborious process.⁸⁴ While the result of this might be to make a particular stratification appear self-evident to a competent reader, it is really only a surface effect.

From one perspective, therefore, we see the boundary of the assemblage drawn, and from the other constructed as an end-point of logical compatibility derived from experiment. As long as the two match up — the conceptual delimitation of the boundary, and the constructed

out as ridiculous. Mere ‘articulated transmission’, by contrast, does not preclude this kind of cherry-picking, as it demands only an *account* of it. (McKay, cited in Jackson 1997, has produced the same kinds of prediction using the text of *Moby Dick* — which rather proves the point.)

⁸⁴ It might also be argued that expertise allows the plugging in of the source document to systems of which non-experts might have no knowledge.

limit — all is well. Such a system is, of course, tautological; but it is self-validating. A scholarly edition is, therefore, an instance of what I will refer to as a ‘reflexive transmission’: that is, it simultaneously transmits and validates itself as a whole.⁸⁵ (See Figure 8) Or, to put it another way, it represents the coincidence of arbitrary closure (“*I have stopped editing*”) with a theoretical closure (“*It is finished*”). In this, we see the distinction between *reflexive* and *articulated* transmission: the latter only says, “I have finished, and this is what I have done”.

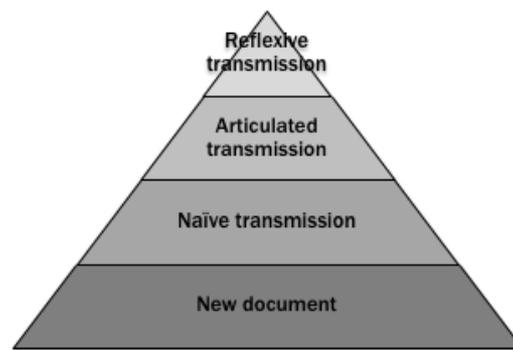


Figure 8: Reflexive transmission as a self-validating form of articulated transmission.

A reflexive transmission’s regulative idea is therefore what the output, in its totality, represents: not the document, *per se*, but a specific, articulated and self-consistent determination of the document. It might be, therefore, a literary work (as an idea), or the text of the document — or any other determination. Of course, to this we must remember a further necessary system into which the source material must be plugged: namely, the output. (The reasons for this are discussed more fully above: see page 138.) Thus the output

⁸⁵ As noted in the introduction to this chapter, this is *not* a definition of scholarly editions, or sufficient account of them; however, ‘being a reflexive transmission’ is a *necessary requirement* of a scholarly edition, if it is to give a complete account of itself. Other forms may be considered reflexive transmissions, without conforming to particular standards for a scholarly edition.

represents the source as it determined and represented⁸⁶ — certainly a tautology, but I think a powerful one, and, indeed, the only possibility.

In particular, its power comes from the wrapping up of contingent arguments such that they always point back to the regulative idea, rather than being passed on as contingencies to the user. It is not to say that it is anything more than subjective, but the subjective buck stops with the editor. The user either buys into the whole regulative idea as a representation, or does not. But if they do buy into the system, it has the effect of underpinning, as a single articulable (i.e. citable) argument, any further claims that the user can make about the source based on the output. It is a region of objectivity carved out from only subjectivity: not universally objective, but functionally so within its own domain. This makes it, in the terms of Deleuze and Guattari, a kind of *signifying regime*.

2.9. Regimes of signs

Following Deleuze and Guattari, I will argue that the effect of the regulative idea applied to the whole system is to create a particular *regime of signs* or semiotic system:⁸⁷ specifically, a *signifying regime*, resulting from the overcoding of the source material as a whole by the edition's interpretation of that source material.

In the fifth 'plateau' of *A Thousand Plateaus*, Deleuze and Guattari argue that there are several kinds of sign regimes — systems that concern and govern the connection between signs. "In contrast to structuralist accounts (whether anthropological, linguistic, or psychoanalytic),

⁸⁶ Bearing in mind that *representation* is really created by the encounter between the output and the user. Alone, the output does nothing.

⁸⁷ Much of this argument is based on the fifth 'plateau' of *A Thousand Plateaus*, entitled *On Several Regimes of Signs* (Deleuze and Guattari 1987, 123–64)

which pursue an ahistorical and universal account of the sign, Deleuze and Guattari argue that signs order and are ordered differently in different regimes.” (Adkins 2015, 83) They identify four such regimes: the signifying regime (that of structuralist linguistics, for instance), but also the pre-signifying, post-signifying, and counter-signifying regimes (Deleuze and Guattari 1987, 123–30) These regimes, they argue, are always in reality mixed, and, moreover, that the four regimes they identify constitute an arbitrary selection from an infinite number of semiotics (these four then may be considered ‘ideal’ types, which do not exist in pure forms). (Deleuze and Guattari 1987, 132)

Though Deleuze and Guattari begin this plateau with a discussion of the signifying regime, they argue that this is “[n]ot only not the first [...] but [they] see no reason to accord it any particular privilege.” (1987, 130) As Adkins observes, the “mistake of structuralism has been to take the signifying regime as the only possible model.” (2015, 84) This, indeed, has been the approach hitherto adopted in this chapter: the avoidance of starting from the perspective that a signifying regime, in which transmission ‘automatically’ (*automagically*) ‘stands for’ the source material, and instead looking at the emergence of particular sign regimes. Assuming a signifying regime from the output is to argue that the output of a transmission will represent the source come-what-may. Accordingly, this analysis can be seen as operating in a largely *pre-signifying* regime, the “regime that continually wards off the hierarchical formations of the signifying regime” (Adkins 2015, 84)

Naïve transmission can be seen as indicative of a kind of post-signifying regime: it has been sent out into the world, isolated from the regime in which it did signify something, and has now taken on a life of its own (as far as anyone can be certain). The naivety, therefore, really comes from the user taking this post-signifying regime and transforming it into a signifying regime: assuming (naively) that it does represent the source. Articulated transmission is rather

a mixture of the signifying and post-signifying: its internal components are subject to a signifying regime, being codified. As Adkins suggests, “the signifying regime is the regime that treats signs as discrete constants amenable to codification” (Adkins 2015, 85) But at a holistic level, there is still a looser, post-signifying connection: the whole is not codified; there is no total, all-encompassing signifying system. The regulative idea, by contrast, creates a ‘reflexive transmission’ by codifying or determining the whole.

As in all signifying regimes, the regulative idea is an overcoding signifier, a “supreme signifier” to which all other signs refer (Deleuze and Guattari 1987, 127) According to Deleuze and Guattari, the “supreme signifier” has two functions: it delimits a closed space on one hand, and on the other interprets it. (Deleuze & Guattari 1987, 114) They associate this most readily with State formation, borrowing terms from Georges Dumézil’s work on Indo-European mythology: sovereignty, they say, has two heads, the “magician-king and the jurist-priest”, one ruling a territory and the other interpreting its laws. (Adkins 2015, 192) The State, for Deleuze and Guattari, is a signifying regime of this kind. But, I think, the parallels with our conception here of a *regulative idea* are clear. As the ‘ruler’, it determines the extensity of the source-assemblage; and as ‘interpreter’ stratifies it. But they are, ultimately, the opposite sides — two heads(!) — of the same coin, mutually implicating in maintaining the determined, closed system of representation. Therefore — and it might otherwise appear a gaping hole — there is no real need to say what the regulative idea *is*: it constitutes the output as a whole, and is constituted by it. As Deleuze and Guattari write:

There is not much to say about the center of signifi-
cance, or the Signifier
in person, because it is a pure abstraction no less than a pure principle;
in other words, it is nothing. Lack or excess, it hardly matters. It comes
to the same thing to say that the sign refers to other signs *ad infinitum*

and that the infinite set of all signs refers to a supreme signifier. (Deleuze & Guattari 1987, 114-5)

It is, therefore, within its own regime, the sign of the thing it represents, which it — itself — determines. In *Theorizing the Digital Scholarly Edition*, Gabler claims quite correctly that the scholarly edition is the editor's text “*in lieu* of the author” (Gabler 2010, 44) — but this only raises the question of how, if it is the editor's text, is it also the text of the source material. In the above argument, we have an answer: because a signifying regime is created in which the text of the source material is also the editor's text. Such a regime establishes the ‘regulative idea’ as the sign of the source, and creates a signifier–signifies relationship between the output and the source.

This signifying regime, it should be reiterated, is not a universal system. It does not preclude other such regimes: someone else can transmit the source, and in doing so create an entirely different regulative idea. Nor does it really exist outside the user's encounter with the output. It may, in more malleable media forms, be reconstituted, or have its totality brought into question by the addition or subtraction of more material — in which case, it is either a new signifying regime, or it falls back to a different mixture of regimes.

2.10. Conclusion

This chapter has set out a model of transmission of textual sources by reproduction, looking at the interactions at each stage of the process as a conceptual model. The point of departure was to make the case that the only viable conception of the output of transmission is an entirely new document. This is, I think, the logical conclusion from both the model set out here, where the source material is conceived as an instance of the Deleuze and Guattari's concept of an *assemblage*, as well as other theoretical positions, most notably McGann's

conception of the 'social text'. That the output of transmission is a new document throws, necessarily I think, into doubt the relationship between the source and the output; it certainly cannot be assumed. It also suggests that to argue that the source and output are versions of the same pre-existing 'ideal' text is to mistake how things actually work: based on Deleuze's notion of *difference* as ontologically preceding *identity*, it can be argued that the ideal work arises from the existence of versions, and not the other way round. The challenge, therefore, is to explain the connection between the source-version and the output-version.

In breaking down the process of transmission into a series of translations between media, based on encounters between each 'actor' considered again as an assemblage, it has been possible to make three critical distinctions between kinds of reproductive transmission. These are (re)illustrated in Figure 9.

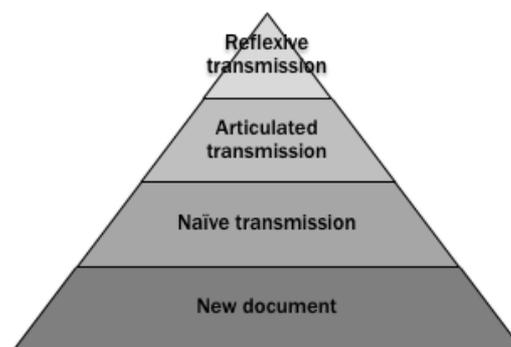


Figure 9: hierarchy of transmission validities⁸⁸

⁸⁸ One reviewer of this thesis put a very fair question: 'Why a hierarchy?' This is a very fair question, especially given the general predominance of non-hierarchical thinking in the work of Deleuze and Guattari. (At first glance, a Venn diagram of transmission validities might seem more appropriate!) In response, as I think this is a decision that needs justifying, I would say that it is not the approach here (or, indeed, of Deleuze and Guattari) to eschew hierarchies — as if they could be prevented — but rather to recognise them, their fundamental non-essentiality, and to explain their genesis. 'Transmission validity' is here presented as a hierarchy as it *is* a hierarchy — from the perspective of a user at the end of the transmission.

At the base (illustrated in red) is the fact that any reproduction constitutes a new document, and with nothing else to go on, it should not be considered as anything else. The distinctions that are drawn are therefore of an epistemological nature: in other words, what reliability as a transmission of the source can a user draw from the output.

The first critical distinction, between an entirely new document and what I have called *naïve transmission*, is predicated on some kind of activity that reproduces the source and a simple claim to be a copy. The naivety really comes about in showing a kind of blind faith in the act of copying to produce something that functions as a reproduction.

The second distinction, between naïve transmission and articulated transmission, is that an articulated transmission gives an account of what *is* represented and *how* it is represented. But this is only a positive account of what *is* represented, but not what is *not*. It gives no guarantees about the whole transmission. As the output is itself necessarily a closed system, the boundary of the source is arbitrarily drawn — or, rather, simply *ad hoc*: it is where the transmission stops. But there is not, or cannot be, an account of where that boundary is, beyond it simply being there.

Reflexive transmission is therefore a subset of articulated transmission: it articulates not only the transmitted components, but gives a global account of — the rationale behind — closing the output off at the point that it does. Determining this is, essentially, a tautological (or autopoietic) construction: what Eggert (and I) term a ‘regulative idea’, created out of the source material and simultaneously delimiting it.

It is the level of *assurance* one may have about the transmission one receives. The pictorial representation may be drawn as a Venn diagram, but it would be the same diagram (interpreted as a conical cross-section) viewed from above. A ‘flat’ diagram of a hierarchy only hides the hierarchy.

In one sense, it is tempting to argue that nothing but reflexive transmission is really transmission at all. Without guaranteeing the whole, there is always the lurking contingency of incompleteness, whose distorting effects cannot ever be pre-emptively countered. But I think that this is overly simplistic, and perhaps reductive. There is a level of transmission — naïve as it may be — that is predicated simply on something having been copied and knowing that it was. To take, again, the example of a photocopier: a person watching reproductions emerging from the copier would assume that something bearing a vague resemblance to the output was inside being copied. Particular kinds of ‘editor’ — especially more mechanical reproduction — seem to invite a greater degree of faith than seemingly more fallible human activity. This is understandable. But it also, I think, points to the user ‘plugging in’ their own assumptions about the source material: “It is a printed page, therefore a photocopier is sufficient to reproduce it.” Such logic can be seen as an appeal to convention; there is no guarantee that it is valid, but the assumption of it goes a long way.

A similar kind of logic may be applied to articulated transmission. In this case, the determining factor might be the apparent ‘natural’ discreteness of component parts. For instance, there is clearly a difference between transmitting a lot of haphazardly collected whole documents and transmitting haphazardly selected words from a single document, though theoretically establishing a dividing line is difficult: as has been argued in this chapter, determining the discreteness of any document or set of documents is an interpretative act. The difference between the two can, perhaps, be seen in the potential for a component part to support its own regulative idea. The question then arising is whether a regulative idea for the whole can apply in concert with myriad, component-based regulative ideas. Such a divide is most strongly felt in digitisation processes that are linear, as opposed to experimentally reflexive, in nature.

The next chapters will develop a mechanistic conception of these ideas more fully, before Chapter 4 looks at real examples of scholarly editions and digitisation projects.

3. Transmission models and parameters

No text is "transferred" wholesale, like a bucket of coal being moved along a conveyor.

— Johanna Drucker

3.1. Introduction

In the previous chapter, I began to outline a conceptualisation of text digitisation, (digital) scholarly editing, and — by extension, and far from explicitly — everything in between, under the generic notion of *textual transmission*. Based on several concepts of Deleuze and Guattari, I articulated the transmission process as a series of interactions or encounters: between source material and editor, editor and output, and output and user. Each stage in this process marks an interpretative act. However, in the interests of making this a general model, a great deal was, if not simplified, then ‘wrapped up’ into the three encounters described. This, I would contend, does not invalidate the conceptualisation. Almost anything with complex underlying mechanics can be modelled as a ‘black box’.

The most prominent black box in the model of textual transmission is what I have labelled ‘the editor’. The editor, as observed, is always an assemblage: a set of component parts — people, pens, paper, fingers typing at a keyboard, computers, their operating systems, and, significantly, knowledge — that, taken together, interact with the source. But assuming the result in any given process were the same, it would not matter quite how the editor worked. They could, in producing an XML document, use Windows or Linux or a Mac, and be fuelled by quinoa or doughnuts: it makes no difference if the result is the same. Or an optical character recognition engine that is as accurate as a human reader need not, in terms of its

result, be distinguished. “A difference,” as the aphorism attributed to William James goes, “that makes no difference is no difference at all.” By considering the transmission process, and especially the ‘editor’ in the abstract — as a black box — it was possible to avoid any pre-emptive distinctions: digital versus print; critical editing versus page facsimiles, and so forth. Everything, as the previous chapter argued, is interpretative activity — whether photographing a document, transcribing it or critically editing a text — and is thus, on this basis, describable by the same abstract mechanisms.

Of course, James’s aphorism evidently has a corollary, which might be put thus: “A difference which does make a difference makes all the difference.” In modelling textual transmission in this way, we saw that the validity of a transmission is not automatic — or, put another way, that transmission does not automatically lead to representation. Thus we identified three *validation thresholds*: a *naïve transmission*, that only claims to have transmitted the source; an *articulated transmission*, that articulates how the source was transmitted; and a *reflexive transmission*, that builds a complete model of what is transmitted, and transmits and articulates how. These are determined by *real* differences in the transmission process, even if, in describing them, their properties were described in terms of abstract mechanisms within the overall ‘editorial black box’.

The aim of this chapter is, therefore, to unpack or decompose the editorial stage of the transmission process to provide an approach to exploring the ways in which how the transmission process takes place creates these differences in representative validity. Validity thresholds are, after all, not intrinsic properties: a scholarly edition is not trustworthy because it is a scholarly edition; it is a scholarly edition because it is trustworthy, because it has a certain validity. If a transmission is to be considered *reflexive*, the general logic of repetition, of ‘editorial experimentation’, must actually happen in the transmission process. This chapter

intends, therefore, to articulate a formal model of transmission processes, in which we can see, as a result of a given configuration, how (or even whether) a given validity threshold is reached.

I do not think it is going too far to argue that textual scholarship — traditionally speaking — has shied away from considering process, instead concentrating on outcomes. This is especially the case when contrasted with digitisation, which talks of workflows and tasks and transformations; and even digital scholarly editing, which is altogether more process-orientated. There are, I would suggest, two reasons for this. The first is the way in which scholarly editing was carried out: as the activity of the ‘lone scholar’. As Pierazzo says:

For centuries, editorial work was more or less the province of a single scholar able to control and supervise all stages of the process up to the point when the manuscript was delivered to a publisher; the publisher was then in charge of typesetting, preparing the layout, copy-editing the manuscript, and so on.

(Pierazzo 2014, 135)

Or even more so, in the case of Shillingsburg, who continues to typeset his own editions using LaTeX. (Shillingsburg 2013). This stands in marked contrast to the typical creation of digital scholarly editions, which, as Shillingsburg pithily describes it, “takes a village” (Shillingsburg 2005, 94). This is, in part, due to the wide range of expertise required quite apart from the ‘editing proper’: programmers, designers, server administrators. This range of discrete tasks, carried out by distinct individuals, makes the whole process of transmission more easily decomposable into stages. The person who knows how to operate the camera captures a digital image of the manuscript; several people transcribe; a computer collates; a programmer transforms some TEI into HTML; a server administrator sets up a web

application... It is easy to see that there is a workflow, and it comprises distinct stages, because the stages are more distinct and concerns separated.

This is not to say that one scholar working alone does not have a 'workflow', *per se*. Greetham's excellent guide to textual scholarship comprises a detailed account of the typical stages that comprise the creation of a scholarly edition (selection, transcription, collation, and so forth). (Greetham 1992) The Lachmann Method can, in one sense, be seen as a set of discrete tasks in a logical order. That said, the precise actions of the editor — as they may be conceptualised — are more difficult to grasp: the editor is still, largely, a black box, and in particular the editor's brain. We cannot, thus far, adequately decompose and model its intricacies. But it is important not to draw distinctions primarily on this basis: a human reader *is* an OCR system, with variable quality in different circumstances.

This aside, I would suggest that there is a second reason for this focus on outcome over process in textual scholarship. As the previous chapter argued, a scholarly edition is a reflexive transmission, which serves to create an identity between the source and the output, establishing a signifying regime. The outcome is that a user, by 'buying into' the established regime, is able to take for granted the relation between source and output. *What is it?* subsumes the question, *How was it made?* We do not, as users, need to be concerned with the second question, because the editor has wrapped up all the contingencies of *how* into the *what*.

But this leads us to two fundamental requirements. Firstly, it is necessary to account for transmissions that do not fulfil this requirement, or may only partially fulfil it. Such cases call for a more full account of the *how*. Secondly, that reflexive transmissions, as much as they mask the notion of process behind that of outcome, must come about somehow: that

is, they are as much a product of a particular process as anything else, even if this may be later taken for granted. What is needed, therefore, is a mode for modelling — that is, providing a generic account of — transmission processes, decomposing stages into sub-stages. In doing so, we can see, by the way stages ‘connect up’ to each other, the effects of any given operation on later operations, the nature of interpretative models and their basis, and, ultimately, determine how these various validity thresholds are, or are not, reached.

This consideration of process is very much a Deleuzian understanding of things. His philosophy, as has already been seen, is based on a rejection of intrinsic identity in favour of ‘becoming’. (See DeLanda 2006; Adkins 2015) As Stagoll notes, to say ‘a tree turns green in the spring’ is, in Deleuze’s terms, mistaken: “[W]e ought not to say ‘the tree became green’ or ‘the tree is now green’ (both of which imply a change in the tree’s essence), but rather ‘the tree greens’ ... [becoming green is] not a disruption of some continuous state, but rather the state is constituted by events ‘underlying’ it.” (Stagoll 2005, 90) In a similar vein, we find Drucker’s view of text as necessarily performative; as a process:

A quality, materiality, whose identity depends on contingencies cannot be mistaken for a self-evident object. For instance, we understand the production of a play as a distributed material event, rather than as a fixed or static thing. We can understand all textual and material production in the same way — as dependent upon interrelated systems and conditions.

(Drucker 2013)

The first section of this chapter aims to articulate a system for modelling transmission as a process. The second section aims to identify various outcomes of a transmission process that arise as a result of specific configurations.

3.2. A transmission model

3.2.1. Decomposing the transmission process

In the previous chapter's analysis, it was suggested that transmission process could be divided into two: 'editing' and 'reception'. In this section, we will deal with the former, which is itself divided into two interactions or encounters. The first comprises an encounter between, on one side, the editor (remembering that the editor is shorthand for a potentially very large assemblage of persons and technology) and on the other, the source material. The second encounter is between the editor and the output, in which the form extracted in the first interaction is recoded and reterritorialised in the output.

However, it is undoubtedly simplistic to consider any of these encounters as a single 'event'. Take the first (source–editor) interaction. If 'the editor' is really — as is especially the case in large digital projects — a team of people, each using computers and different software packages, and cameras and scanners, then it is clear that there is not one single interaction with the source material, but multiple: a succession of encounters. This is even true of one, single editor working alone with a manuscript. Evidently he or she does not 'interact' with the source material all at once, memorising everything, before writing it down in the output. There are intermediate phases: writing notes, making preliminary transcriptions, producing collations, drafts, proofs, all of which precede the second interaction, between the editor and the output. In the conception set out here, then, 'the editor' is always a complex assemblage, and the source–editor encounter always really multiple encounters.

That said, simplified though it might be, the view of transmission as one interaction is not invalid (see the introduction to this chapter). However, by dismantling the source-editor interactions into component stages, we can more precisely see how they operate, and how

specific transmission processes give rise to specific kinds of output and attain certain validity thresholds.

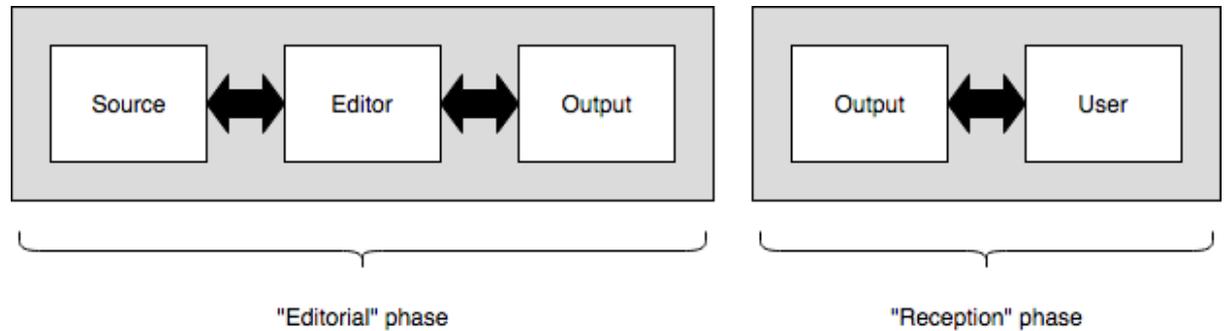


Figure 10: Basic schematic representation of the transmission process, reproduced from Chapter 2.

To begin, therefore, let us first return to the schematic representation of the transmission process from the previous chapter (reproduced here: see Figure 10). At present, we are concerned with the ‘editorial’ stage. This is a macro-level view: the entire editorial process represented as a single stage, comprising two interactions. However, we can break this down into several ‘component stages’. For each component stage we see the same logic as the macro-level ‘stage’ reproduced in microcosm. In other words, each stage is a process of stratification: it takes as its input (or source) an assemblage; it applies an *interpretative model* that stratifies the input; and it reproduces (recodes and reterritorialises) the extracted form, creating as output a new assemblage. This is illustrated in Figure 11.

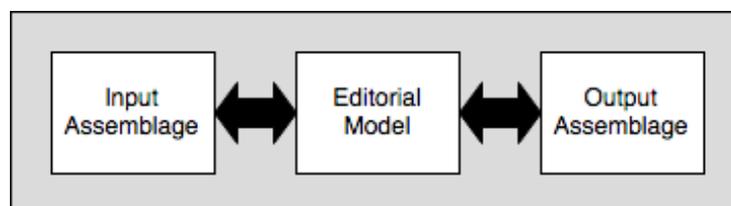


Figure 11: An editorial stage or function.

These intermediate assemblages might be considered examples of what Boot calls a *mesotext* (Boot 2009) — the text ‘in the middle’ of two stages. While Boot conceptualises such mesotexts as, for instance, an editor’s notes (which may be represented), in this instance such ‘meso-assemblages’ can take any form, provided that the form of the output of one stage can be taken up as input to the next stage. They might, therefore, comprise an editor’s rough transcription on paper, or some computer data, or be remembered by the editor from one stage to the next.

Furthermore, it should be noted that, though this generic representation of a stage appears to produce a single output-assemblage — the extracted form — the process of stratification always produces two strata. Strata, as Deleuze and Guattari note, always come in pairs — even if we are only interested in one of them and discard the other. For instance, when we carve wood, each stroke of a plane divides the wood into the object being carved (the output) and the smaller piece and sawdust that are chipped away. When we talk of the output, therefore, it is really the output that we are *interested* in. More than this, even: it is the output to which we have given form (the statue) rather than the ‘formless’ discard (sawdust). Take transcription: when one transcribes the text of the document, the linguistic code is extracted and output in some new form, while the materiality of the document is discarded.⁸⁹ However, in some circumstances, both strata might be relevant. The most obvious example that springs to mind is stratifying a physical document into two surfaces: we might very well be interested in both the *recto* and the *verso*, in which case there are two outputs.

⁸⁹ Indeed, it is sometimes difficult to give specific form to the discarded aspects beyond enumerating them as ‘lost dimensions’ — indeed, to say it has a form might be taking things too far: what is left of a document once we have extracted the text? This is the nature of an assemblage (see Nail, in from previous chapter); it does not need to represent anything (and in this case, we do not particularly care).

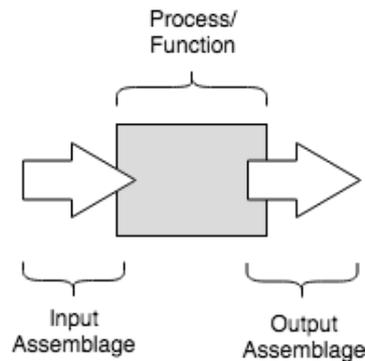


Figure 12: A simplified representation of an editorial stage or function.

Having described the internal logic of a stage, we may simplify the diagram for convenience (see

Figure 12, above). It should be remembered that the grey box represents an action or process, while the white arrows represent the input- and output-assemblages. Considered as such, a stage is simply a function: $f(input) \rightarrow (output, discard)$. What designates a unique function, therefore, is the application of a *single* interpretative model to the input assemblage. Represented in such a way, we have the basic unit with which to model any kind of complex editorial process. A larger process may be broken down into multiple component-functions, where they exist and it is useful for analysis. But, equally usefully, where the combination of functions is extremely complicated (and the details unimportant), its component-functions may be wrapped up and bracketed as a single function.⁹⁰

⁹⁰ This composition/decomposition allows the modelling of complex processes such as transcription — which is really multiple stages of interpretation and reinterpretation at the level of graphemes, lexemes, linguistic structures — as a single operation, unless we are specifically interested in the details: transcription considered as a whole is simply a function that takes a document as an input and outputs a linguistic text. Likewise photographing a page: we can bracket the internal mechanisms of the camera (which *could* be represented if we chose) and concentrate on the input and output.

In Deleuzian terms, we may regard each function as an *event*. An ‘event’ “describe[s] instantaneous productions intrinsic to interactions between various kinds of forces.” (Stagoll 2005, 90) As such, it is a single, one-off interaction, determined by the factors leading up to its enactment, including the input assemblage and the assemblages brought to the interaction as interpretative models. This uniqueness is a vital consideration in later analysis: we must assume an interpretative operation in one circumstance is not the same as if it were carried out at another point in time, or by a different person, or on a different input.

3.2.2. Linear and recursive connections

One of the types of connections between functions is a *linear* connection — that is, functions chained together so that the output of the first stage becomes the input of the second stage.⁹¹ At the broadest level, there must be linear connections between at least some stages: after all, the entire editorial process can be regarded as a single (and therefore linear) stage, as it takes some input (the source) and produces an output. When we say ‘workflow’, this is what we mean.



Figure 13: Functions in a linear chain

As an example, in the above diagram, the first stage could be taken to represent the photographing of a document-surface (the input is a document-surface, the process is the photographing, the output the photograph), while the second stage represents the

⁹¹ We see this very literally in computer programming, such as with the Unix ‘pipe’ operator: `input | do-a | do-b | do-c` (the input is ‘piped’ to `do-a`, the output of `do-a` is the input of `do-b`, etc.) Alternatively, in other programming languages, function calls are nested: `do-c(do-b(do-a(input)))`

transcription from the photograph. This, it should be noted, makes the overall process of ‘producing a transcription’ different from a transcription process whose input was the document itself. The photograph is evidently different from the original document: its materiality, amongst other things, has been ‘discarded’. Stages that are connected in a linear manner are affected by the operation of the previous stage. “In principle, every link in such a chain affects and delimits the nature and quality of the final digital product,” (Dahlström 2009, 173). In the example of a photograph, if the photograph were of a low quality, the transcription might be different. Or, if the photographer is careless and crops the top third of the page, the transcription of that top third cannot be made. Robinson’s idea of a *chain of digitisation* makes just this point: that information is lost, or at least fundamentally altered, by each stage.⁹² (Robinson 1993) Tanselle likewise makes such a point with regard to using photocopies *in lieu* of original documents: the results are “editions of the texts of photocopies, not editions of the texts of the original documents”. (Tanselle 1989, 61)

We therefore see another property of an editorial function: it is a single event, happening once, and in unique circumstances. It must, therefore, necessarily be the application of a pre-determined interpretative model:⁹³ a smash-and-grab raid using the tool at hand. This has several consequences, not least the capacity for a function to fail to produce an effect. In such circumstances, the function may have one of two outcomes: firstly, the output *is* the input-assemblage (it has no effect, produces no stratification), or the input-assemblage is discarded

⁹² Robinson’s argument is more focused on remediations-of-remediations between distinct projects and over time (e.g. taking one remediated form, say microfilm image of a document, and ‘re-remediating’ the microfilm by making a digital scan) — but the logic of his argument stands even at smaller, ‘inter-project’ levels, provided what we consider a *function* is genuinely isolated.

⁹³ If it is *not* pre-determined, it cannot truly be a single interaction: there must have been a first interaction in which the interpretative model was determined.

and there is no output. In the latter case, the function acts as a kind of selection — quite simply because later functions cannot be carried out.

It is for this reason that an overall process that decomposes to purely linear functions tends to produce *naïve* or *articulated* transmissions: there is no point at which interpretation can be based on the source at hand, being necessarily pre-determined. It is possible to *articulate* those functions — which would be representation of a kind — but a lesser, necessarily incomplete, representation: a mapping of ‘output facts’ to ‘source facts’, but without grounds for justifying the choice of facts from the infinite possibilities.

The adaptation of an interpretative model (a necessity for producing a *reflexive transmission*) is therefore, as argued in the previous chapter, based on experimentation and hence repetition: in other words, taking the output of a function and using it as the input for the same function. We might call this a *recursive* stage. To give a simple example, having accidentally cropped the top of the document when photographing it, the photographer might check the photograph, and, noting the problem, take another photograph, having corrected the alignment of the document; only a ‘good’ photograph can be passed to the next function.

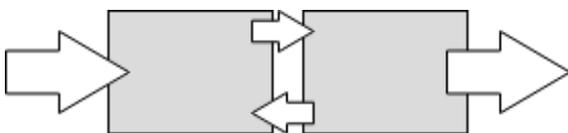


Figure 14: two mutually recursive functions.

Accordingly, what can, at a higher level, be seen as a single recursive function is really multiple. A recursive function must decompose to at least two functions, of which one is (or functions as) an action, and the second a validation (i.e. ‘take photo, check photo’). This

makes the two stages *mutually recursive*.⁹⁴ This general idea of mutually recursive stages which, between them, act as actions and validators, allows us to account for the kind of repetition required for the development of complex adaptive models. A complex model is really a set of simple models connected in a mutually recursive manner. And, of course, it can only be as complex as its constituent functions allow: the ‘take photo, check photo’ is relatively simple, comprising only two functions. Figure 15 shows four functions connected in a particular, mutually-recursive configuration (the function producing the output — whatever operation it might perform — is therefore the final validator).

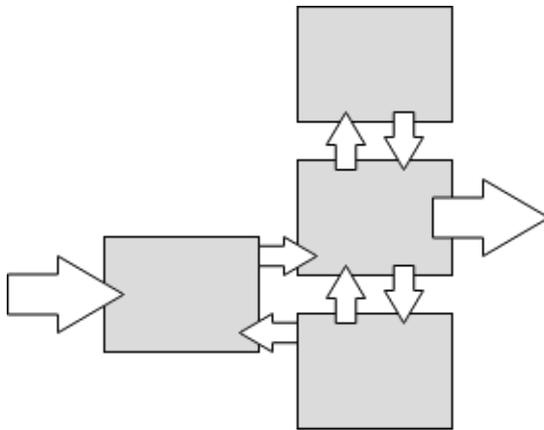


Figure 15: Four functions connected in a mutually recursive fashion.

Transcription is an example of more complex mutually-recursive functions. In the previous chapter, this was described as a kind of repeated experimentation by the editor. Now we may, albeit primitively, model such logic. (‘Primitively’, because we cannot fully decompose the operation of the human brain!)

For example, let us imagine the process of transcribing the text in Figure 16 (below). We might begin by applying a function that extracts the graphemes. This produces two outputs:

⁹⁴ The ‘taking a photo with the page rotated x degrees’ function may even be decomposed into a set of distinct functions (say, one for each degree); some validate and some do not.

some graphemes (“the cut sat on the mat”) and the discarded material aspects (let us assume that this is a discrete document). The ‘interesting’ output — the graphemes — might then be passed to a function that attempts to stratify the graphemes as something meaningful in the English language. This succeeds (“cut” is, after all, a word). This might then be passed to a function that checks the validity of its input against, say, English idioms or poetic rhyme — in which case, there is a failure (“the cut sat on the mat” does not match a common phrase, or its rhyme scheme fails). Whereupon we return to the first, grapheme-extracting function, and ask whether the *u* could be something else (yes: it could be an *a* that has been written quickly). The above chain of processes then repeats, and we produce “the cat sat on the mat”. We might compare this logic to that governing the transcription of Figure 17. In this example, having ‘tested’ the graphemes, we return to the first stage but are unable to resolve the “cut” problem: it really is a *u*, and attempt to find another model that might justify it. For example, could it be a typing error? We might reach for further interpretative functions, taking into account a wider context. Eventually, nothing producing a result, we might simply accept the ambiguity: “the cut [sic] sat on the mat”.

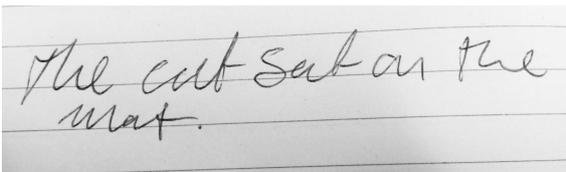


Figure 16: example of transcribing ambiguous handwritten document

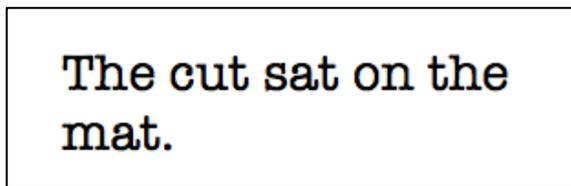


Figure 17: example of transcribing ambiguous printed document (assume it is written on an actual typewriter)

The above example should not, of course, be taken as a definitive description of the process of transcription. Almost invariably, the process is more complex, and involves many more

functions connected to each other in a mutually recursive fashion.⁹⁵ However, as an example it illustrates the potential for mutually recursive interpretative functions to develop a complex model that arises out of interaction with the source itself. This is especially apparent if we compare it to a one-shot linear function, such as (a by modern standards quite primitive) optical character recognition system. Put simply, such a process, having only one ‘attempt’, must simply impose a model: OCR would, for instance, simply produce, “the cut sat on the mat” and ask no further, essentially higher-order, questions, nor, therefore, adapt its model.⁹⁶

What we see here is that, for all the potential to interpret indefinitely (for mutually recursive functions to bounce off each other in perpetuity) there is a terminating condition: what might be regarded as *sense* — that is, each function is essentially satisfied with the validity of the assemblage produced. As such, we may trace a distinct path between the functions (a kind of stack trace, in computing terms) that were actually carried out, and in what order. The myriad possible paths that could be taken are, in Deleuzian terms, the *virtual* (but nonetheless *real*) dimensions of the assemblage, while the *actual* dimension is the path taken. As such, for all its potential complexity (the ‘cat sat on the mat’ example above is highly simplified), there

⁹⁵ It is not the intention here to present a *general* — in the sense of *universal* — account of transcription (or interpretation in general); rather, the manner in which functions are connected and operate are always local and contingent. My contention is that this is an *abstract* account — or, better, a methodology for decomposing actual transmission processes and analysing how they actually work (rather than attempting to foist a pre-determined model on any particular process).

⁹⁶ By the same token, we can envisage methods by which an OCR engine might be more successful: by, for instance, verifying produced words against a dictionary, or by function in a syntax tree, before checking again for other possible matches — in other words, by building in mutually-recursive functions. It is for this reason that I think drawing an ontological distinction between computer-mediated OCR and human transcription is fundamentally misguided: there is only a difference of degrees. In suggesting as much, I am perhaps leaving myself open to accusations that I am downplaying the ‘art’ of interpretation, particularly the role of intuition. But I think even such things as intuition can be considered in these terms: what is intuition, after all, but a cognitive leap that is retroactively (and rationally) validated — taking a jump forward, but then returning to fill in the gaps? An intuitive leap that cannot fill in the gaps is simply asserting something out of thin air.

is a route through the maze, and it is in theory possible to describe it. Whether we do so depends on whether such a description is useful to analysis.

This has one further implication. If we can chart a path (get a stack trace) of all the functions and the order in which they were carried out, one particular function always has the last word. As such, a hierarchical arrangement exists between the strata,⁹⁷ which we saw in the transcription example above: the visual appearance of the u/a grapheme in the handwritten example was subsumed under a linguistic interpretation (the handwriting is a ‘weaker’ stratification that can be overruled by higher-level interpretations), while in the type-written version, the opposite was the case (typed text is not so easily overwritten: the *u* is a *u*, even if other systems disagree). The structure of representation (as seen at the end of the previous chapter) is hierarchical: that of the signifying regime.

For future representations, therefore, a recursive function can be represented as a single function, as illustrated in Figure 18. This marks the *path that was actually taken*. Implicit in this being able to ‘package up’ a complex set of mutually recursive functions in a single stage is the fact that each internal, component-function must have no side effects, or rather, no side-effects of interest (that is, no additional strata that are not simply discarded). By contrast, if a stage contributes some other information to another stage, it must be modelled as such. This will become apparent in the next section.

⁹⁷ Given this, we can (re)produce Van Hulle and Shillingsburg’s ‘Textual Orientations’ (2015; see the conclusion of Chapter 1). What each of these orientations produce is a ‘viable’ stratification. They are, in a sense, archetypes: not blueprints (not a *given* way of approaching text), but an abstract set of plausible stratifications. When they talk of the ‘hypothetical editor’ producing confusion by mixing readings from different versions of a text based on personal taste, they describe the application of *two* forces of stratification, both of which — contradictorily — attempt to have the last word, and thus fail to produce a single, hierarchy of strata.

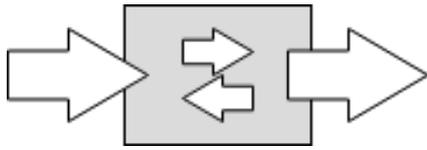


Figure 18: simplified representation of a recursive stage

From this analysis, it is easy to see the relationship between the recursive functions and the kind of logic articulated in the previous chapter with regards to the construction of a *regulative idea* and the notion of *reflexive transmission*. A recursive stage in a transmission process is what establishes — or contributes to the establishment of — a regulative idea; that is, a regulator of interpretation derived from and applied to (in an autopoietic fashion) the input, and not simply imposed upon it. It is not, of course, correct to say that the presence of a recursive phase ensures this kind of representational regime, but that — in all but trivial cases — it is a necessity. At the absolute limit, an entire transmission process (if that is really what happens) can be considered as a single recursive function. From any point in the transmission chain it is possible to return to a previous function, adapting that function's input. Globally reflexive transmissions do just this.⁹⁸

It must also be the case that, once a complex model is developed and applied at one stage, later stages are able to preserve this complexity. This point was, indeed, discussed in the previous chapter: the stratification at one stage, if it is to be maintained, must take into

⁹⁸ There is, as will be discussed later, the possibility of locally reflexive transmission: where individual paths are themselves reflexive, but the whole — the accumulation of paths — is not.

account later processes, including the form of the output and the user's interaction with it. This makes the entire process globally recursive.⁹⁹

This tallies with McCarty's view of modelling as an iterative process. (McCarty 2005, 260) Likewise Pierazzo: "[M]odelling is an ongoing process: as knowledge of the object deepens within the context of the project in general and the editorial work in particular, so it is to be expected that the initial model will need to be refined many times over." (Pierazzo 2015, 26) Considering transmission as a process, therefore, there must be some point at which a model can be developed and refined, rather than simply imposed: that is, some potential within the overall transmission process for recursion.

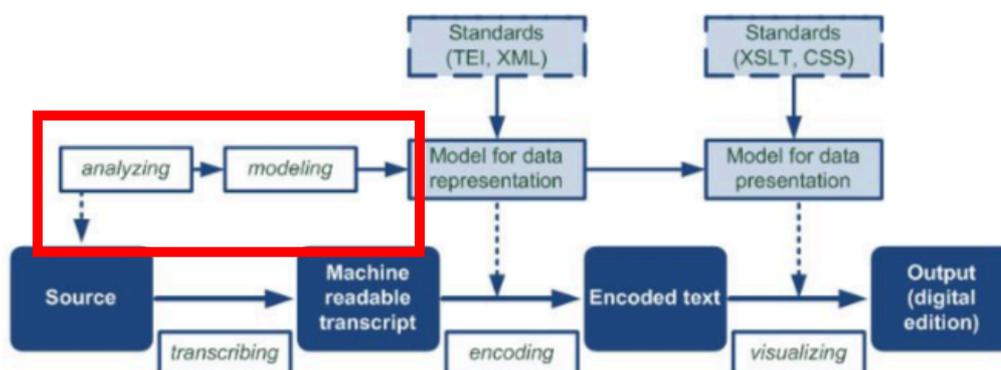


Figure 19: Rehbein and Fritz, workflow of editing (Rehbein and Fritz 2012, 52; reproduced in Pierazzo 2015). The red box (my addition) highlights the recursive loop.

In Rehbein and Fritze's model of an editorial workflow (Figure 19) we can see just such a 'recursive loop' (the 'analyzing' operation points both forwards to the model and backwards to the 'Source'). Where there is no recursion, there is only the application of an already existing model: without such a recursive loop there would only be a linear path —

⁹⁹ As suggested in the previous chapter (with the example of flat-pack furniture), the editor cannot control how a user might interact with the output — but a kind of contract exists: use it in such a way and I (the editor) provide certain guarantees.

transcription and encoding necessarily based on a pre-conceived model, not born out of analysis of the document (i.e. deriving, Figure 19, only from ‘Standards’).

3.2.3. Sequential and parallel paths

Hitherto, we have established a base approach for modelling the ‘editorial’ phase of the transmission process, based on a set of functions or stages that can be chained together (a recursive function, though necessarily comprising component-functions, will, considered as a whole, take an input and produce an output, and thus may be chained like any other stage). A series of chained-together stages can be considered a path, as illustrated in Figure 20. This is a sequential path: a previous stage must be completed before the next stage can occur.

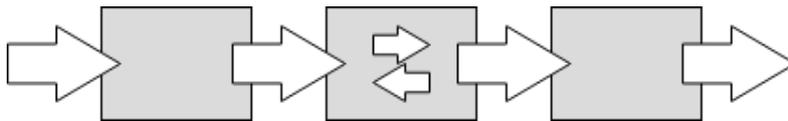


Figure 20: a linear, chained set of stages (the second stage is a recursive stage)

However, paths may also occur in parallel. This occurs when there is some kind of dividing function: a function that produces two outputs of interest. The photographing of a manuscript that divides it into *recto* and *verso* is an obvious example of this. From this point of division, two paths may continue in parallel (see Figure 21).

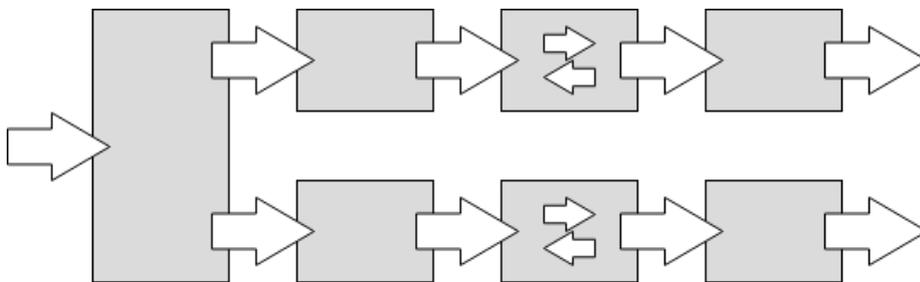


Figure 21: The second stage produces two outputs — for instance, each side of the document — which can follow separate paths of stages

In this example, the first stage represents some preliminary stage (perhaps selection of a document), the second stage marks a division into two assemblages (*recto* and *verso*). The third stage might then be the photographing of each side of the document, which is then transcribed from the photograph, before being output. The important thing to note in this example is the separation of paths, which means that no function in either path has a conception of what is produced in the other.

This is most significant with regards to the recursive transcription function: the complex interpretative model that is created in each of these stages is separated from the other, and as such they are — or at least should be considered as — different models. It should be stressed that this occurs only when the two stages are completely isolated from each other, which may not necessarily be the case. When the same editor — and here I mean a human, expert, *scholarly* editor — transcribes one side and the next, undoubtedly some experience of transcribing one side will spill across into the transcription of the second; a particular feature of handwriting that becomes apparent on one side, for example, might guide interpretation on the other. In this case, however, we are straying from the process illustrated: each function produces a side-effect which spills into its parallel counterpart; and therefore each process is not strictly isolated. Parallel paths are strictly isolated from each other.

For instance, large editorial teams that work independently on transcriptions and never communicate operate on this basis of isolation. It is not that *Person A* transcribing the *recto* side or *Person B* transcribing the *verso* are necessarily ‘incorrect’ in their transcriptions (assuming they do not make avoidable mistakes) — but that an interpretative model for transcribing the *recto* is developed separately from that of the *verso*, because they are developed in isolation. Both are, therefore, correct as far as it goes, but the two cannot be trivially merged together with the assumption that their independently-developed models

will not conflict.¹⁰⁰ (Not, of course, that a scholarly edition is produced in such isolation; it is a more common feature of crowd-sourced transcriptions.)

There are two modes by which strictly parallel paths might affect each other: one is accumulative, the other synthetic. Synthesis involves the merging of two paths into each other, while accumulation is the concatenation or addition of one path to another without altering the first path. The above example of a transcriber, whose experience of transcribing the *verso* ‘spills over’ to their transcription of the *recto* is one of accumulation. *Recto*-experience is accumulated, *added to* the interpretative model applied to the *verso*. (Merging or synthesis is, by contrast, two-directional: the *recto* and *verso* are merged; made subject to a model that is a combination of the two.) This accumulation is illustrated in Figure 22 (below).

¹⁰⁰ As an example, consider a manuscript whose writer has a particular ‘quirk’ (one letter that looks more like another, for instance). If this appears *once* on one side, the transcriber of this side will have only one contextualised example, and is more likely to take the grapheme as it appears than the transcriber of the other side, who might have more examples. Trivially merging the transcriptions of the two sides produces an error with respect to the document as a whole. If they collaborated, however, we might find a conversation such as: “Have you noticed how [author] writes his [letter] to look like [other letter]?” “Aha... that word makes sense now!” This ‘conversation’ represents (trivially) a later, merging operation that normalises or collapses the two transcription models.

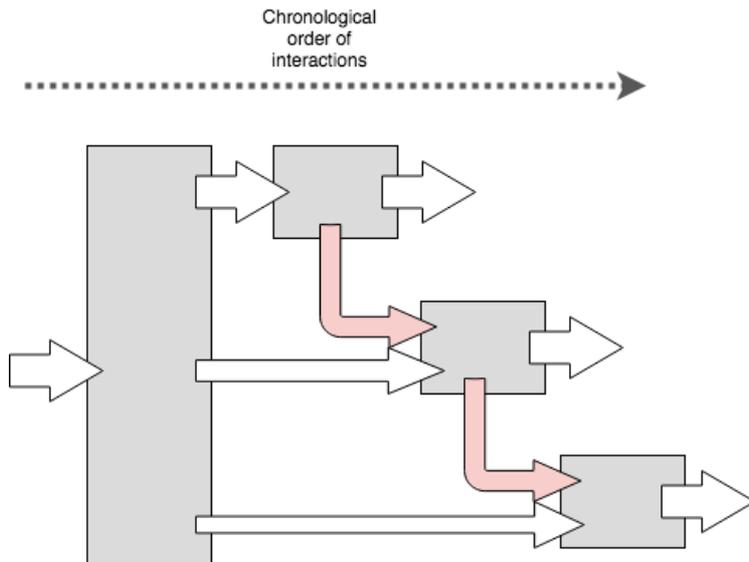


Figure 22: Accumulation of model (red arrows indicate additional assemblage-components passed to next interaction)

Synthesis is, therefore, necessarily a synchronic act: assemblages on both paths need to be considered simultaneously in order to merge them. Accumulation, by contrast, is an asynchronous act: one path is carried out first, and its affects propagate forward to the next path to be carried out (but not the other way round). Indeed, in many cases, it is the synchronicity (or not) of action that determines whether an operation is synthetic or accumulative.

To illustrate this, let us take again the example of transcribing the *verso* and *recto* of a document. In the first path, we transcribe the *recto* side of a document, and — fundamentally — output the resulting assemblage (maybe even going so far as to publish it); as far as the *recto* transcription is concerned, it is complete. However, in this case, the *recto* transcription is carried out with an interpretative model adapted to the *recto*. This, evidently, is different to transcribing each in isolation — but only for the *verso*; the *recto* path has already moved on to the next stage. This is a kind of one-way adaptive model, that establishes earlier (chronologically prior) interpretative actions as a kind of archetype to be applied to later (chronologically subsequent) interpretative actions. This ‘archetype’ model will be discussed in the coming section.

3.2.4. Convergence

Actual convergence of two (or more) paths is necessarily a synthesis. This might seem a strange assertion, as one could readily imagine two paths being joined by the material contents being simply accumulated or concatenated: to get the text of the whole document, for instance, one merely concatenates the text of the *recto* with the text of the *verso*. Really, however, if each side of a document has been interpreted in different ways, a simple concatenation does not produce ‘the text of the whole document’, but rather a juxtaposition of *recto* and *verso*. (Rather like trying to ‘join’ an apple and an orange: you can place them next to each other, but there is still an apple and an orange.) If the two paths can be synthesised together by concatenation, this necessarily assumes the application of the same interpretative model at an earlier stage — in which case, one has already begun synthesis (and thus conceptually merging) at an earlier point. It should be noted that the corollary of this is also true: where two paths are subject to the same interpretative model, the effects of a conceptual synthetic whole (the possibility of merging) are felt.

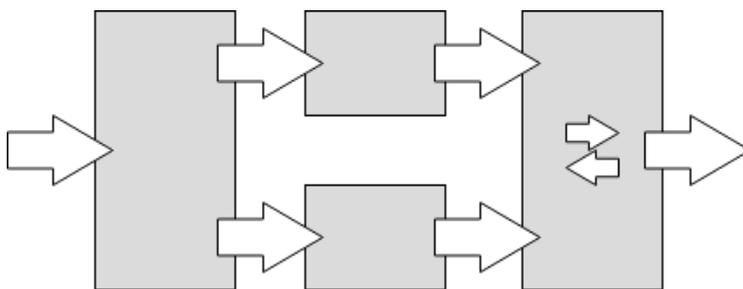


Figure 23: the third stage (for example, a transcription) joins parallel paths, essentially by merging its inputs (the recto and verso). This is qualitatively different to transcribing the two pages separately and concatenating the results.

What, then, is the nature of a merging function? In Deleuzian terms, it takes the extensive assemblages operated on in each distinct path and considers them as components in an intensive assemblage. This new, intensive assemblage is once more rendered extensive by

interpretative action carried out by the merging function. If this merging function is the final stage before publication, it draws a final, delimiting boundary, giving final form to the output.

However, we must consider the case in which there is no convergence, no joining mechanism. This case — as will be seen later — is particularly pertinent in large-scale digitisation projects, where there is no explicit completion and, therefore, no final joining mechanism within the editorial phase. In this case, each parallel branch can be seen as separately sent out to the user. This is the case even when it is ‘published’ within the auspices of a single entity (a single website, for instance). This distinction is illustrated in Figure 24 and Figure 25.

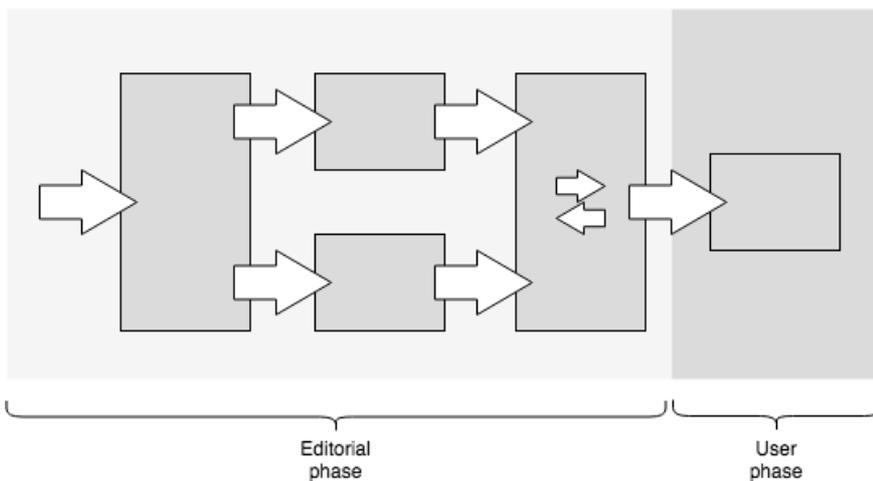


Figure 24: Penultimate function merges two parallel paths into a distinct entity ‘pre-publication’

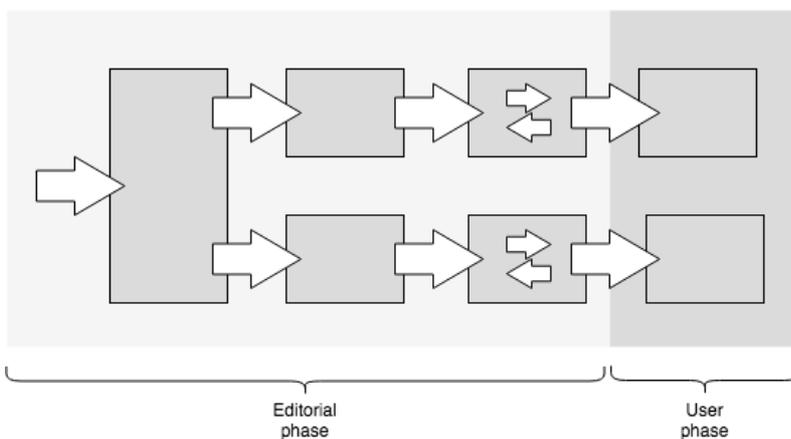


Figure 25: No merging function: two assemblages are ‘published’ as distinct entities to the user.

This conception articulated is useful, as it allows us to model complex processes, particularly at scale; to more easily see when, where, and by whom an interpretative action is taken; and to analyse the effects of these parameters on the resulting output. It is, above all, a realist approach. If, for instance, transcription of each side of a document is carried out by a different person, locked in a different room, we cannot assume that they apply the same interpretation to each — indeed (such is the nature of interpretation), we should assume the opposite. And this situation is different again to having the whole document transcribed by one person as a single interaction. To argue otherwise is, I think, to appeal to some kind of transcending objectivity.

3.3. Model specificity and criticality

From the above, we may begin to outline several properties of models, which will again fall on a spectrum. The first property to consider is that of *model specificity*. This is a characteristic of any interpretative function, and may even be applied to the transmission process as a whole (the overall effect of component functions ‘wrapped’ in a single function). I use the term describe the relation between the interpretative model applied and the assemblage upon which it operates (its specificity). *Specificity* is always determined with respect to the input of the function. An interpretative model can have a high specificity with regards to something else, but a low specificity with respect to the object at hand.

In general, a linear function has a low specificity, being necessarily determined in advance, while a recursive function has a higher specificity, being able to adapt to the given input (i.e. it becomes specific). As the previous chapter argued, there can be no pre-emptive and perfectly adapted model, and hence it must be developed as the result of recursive action.

To give an example in terms of carving wood, a chisel has a high specificity while a sledgehammer has a lower specificity (with respect to wood): a sledgehammer is designed to hit things, while a wood chisel is designed to chisel wood. It is not so much that a sledgehammer can be applied to a wide range of things, but that a sledgehammer (one assemblage) encounters another assemblage purely as a result its ‘hit-ability’. This is again the Deleuzian notion of *affect* (“The capacity to affect and be affected,” as DeLanda says). A sledgehammer hits things that can be hit — which is, essentially, any physical object. The specificity of a wood chisel to wood relates to wood’s ‘carve-ability’, which is not shared by other things (say, a block of iron).

So too with functions in a textual transmission. At one end of the spectrum, we might place the taking of photographs (a sledgehammer of digitisation *par excellence*). As a model, it has a low specificity: it ‘interprets’ anything that light bounces off of, whether a historical manuscript or an elephant. Towards the other end of the spectrum, we might find, for instance, a TEI-XML schema for encoding manuscripts. Evidently, this has a high specificity with respect to manuscripts, but would struggle to interpret an elephant: a high-specificity model applied to the wrong thing produces no output at all (in this, as will be discussed later, we see a conceptualisation for the selection of source material). Of even greater specificity is, of course, a TEI schema specifically tailored to the actual manuscript at hand.

Regarding the specificity of a model, therefore, what matters is, what was the model intended to operate on? And, underpinning this, where did the model come from? How was it created, in interaction with which assemblage, and what is the relation between the assemblage on which the model is based and the assemblage to which it is applied? The specificity of a model can be drawn on a continuum, which is essentially determined by its point of origin (see Figure 26, below).

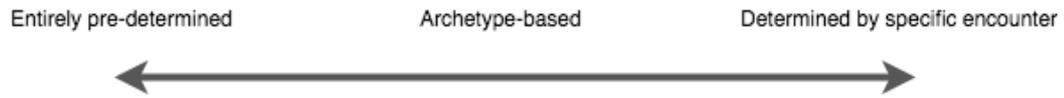


Figure 26: Continuum of model specificity

A model that is entirely pre-determined has a low specificity: it must, being — to put things bluntly — plucked out of thin air as a possible way of interpreting. The *archetype*, by contrast, is determined by an encounter with something that it is assumed will share the same characteristics — whether as a result of a highly generic characteristic (being able to be hit; reflecting light) or because the chosen archetypes are very similar in nature (basing a TEI schema for a medieval manuscript on another medieval manuscript). This distinction is not an absolute, but recognises the difference between an archetype determined within the transmission process itself, or something that might be seen as pre-empting it.

In addition to specificity, we may also consider a model's *criticality*. This is the ability of a model to produce a 'sophisticated' interpretation. A sledgehammer has a low criticality: it can only bludgeon. A wood chisel has a high criticality, but is dependent also on *specificity* in a given encounter (i.e. it produces a sophisticated interpretation when it encounters wood). Specificity and criticality, when both are taken to be high with respect to a given assemblage, are closely intertwined. But a model may be specific without necessarily being highly critical. Or a model may be highly critical, but have a low specificity with respect to the assemblage at hand — in which case, it fails to interpret. Criticality, therefore, is a model's ability to formalise.

We can begin to place various textual interpretative models on these scales of specificity and criticality. Photographing a document, as discussed above, has a relatively low specificity — but not, on reflection, extremely low. In some ways, it is an *archetype*: we already know, having seen photographs of other documents, that a camera is sufficiently specific for the representation of documents (a camera with a large zoom lens is a different matter). However, it ranks much lower in terms of criticality. While it provides what Sutherland and Pierazzo call a “visual description” (Sutherland & Pierazzo 2012), we might question just how good a description it is. This, indeed, depends on the resolution and colour range of the camera (let us concentrate here on digital cameras). A high-resolution, several-million-colour camera is, generally, more specific and critical than a low-resolution, monochrome camera.¹⁰¹ But what it produces is homogeneous — simply a grid of pixels. It does not, therefore, formalise the document (at least in terms of its constituent text: it might formalise the physical extensity of the document). Thus, as Gabler argues, the text of a facsimile can be “merely *seen*”, while in a transcription it can be “*read*”. (Gabler 2007, 198) To me, this initially seemed a strange remark: if one can read the text from the document itself, surely one can also read the text from a facsimile (provided it is of sufficient quality) rather than just ‘seeing’ it. Considered as a degree of criticality, however, this makes sense. The text of a facsimile can be ‘read’ only to the degree that the document itself can be read — which, of course, depends on the expertise of the reader. It does not formalise the text; that is, make it readable *as text*. (Granted, the difference may be slight in the case of a neatly printed or typed document; but could also be massive in the case of medieval handwriting.)

Some form of textual transcription, by contrast, is more critical an enterprise, as it extracts and formalises the text *as text*. However, within this broad rubric there are several degrees of

¹⁰¹ Though, it should be remembered, *specificity* is always *with respect to something*: a monochrome camera deals perfectly well with a monochrome document.

criticality and specificity. As an example, let us consider digital systems for plain-text encoding. Evidently, the ASCII character set is more restrictive (128 characters) compared to Unicode (potentially in the millions), and thus, in general, ASCII will require greater ‘normalisation’ (really, the coercion of document-graphemes into one of the 128 available characters). However, *specificity* also plays a role here. A document produced with a typewriter can easily be encoded in ASCII. The document is already formalised along the same lines; ASCII is thus sufficiently specific. But where this is not the case, Unicode can be seen as more ‘generally’ specific: or, rather, from the wider range of available characters, more precise formalisation can be carried out. As such, both Unicode and ASCII are *archetype models*: they are in a specific case, based on previous (in this case, very generic) experience — in other words, that the editor knows that they can be used to represent text. The interpretative model developed in the application of Unicode to a document-text is the specific characters used (i.e. the actual path taken through a set of mutually recursive functions); and, even with two million characters available, they might still be found wanting. Unicode is not, ultimately, determined by the specific document you are trying to encode, but pre-emptively.

The above argument might have given the impression that a model is the formalisation, but it should be remembered that it is really the application of a formalisation, dependent on, among other things, expertise. An expert encoding a text in Unicode is more likely to recognise, say, the ‘thorn’ character (þ) and encode it as such, rather than as a p. Similarly, a non-German speaker might be able to transcribe the characters of a printed text (they are predominately the same as English), but might interpret the double-s (ß) as a B. It is really a question of expertise: which can be taken here to signify a highly specific and highly critical model. As argued in the previous chapter, there are two ways to conceptualise this kind of expertise. The first, a derivation from first principles, involves a ‘higher-level’ validation (‘It

looks like a B, but “groB” is not a word, and besides it is strange to end a word with a capital letter... Could it be something else?”) — a recursively adaptive function, in other words. Or expertise can be seen as simply the ability to short-cut such thoughts. But, ultimately, given the uniqueness of any input to the function, ‘absolute’ expertise — a highly specific and highly critical model — is only derived from the input itself.

XML encoding — here let us use the example of the Text Encoding Initiative (TEI) — takes things further by allowing the formalisation of other, either non-textual or meta-textual features. Compared to other forms, such as HTML, XML is inherently more critical: HTML has a particular use-case (web presentation: albeit somewhat more semantic thanks to HTML5) and must interpret things according to its own regime.¹⁰² (See, on the relative merits of HTML and XML: Hockey 2004; Renear 2004.) When considering TEI-XML, though, we must be careful to separate the entire TEI (i.e. TEI-all) from either pre-determined subsets (TEI-lite) or specific customisations. In the case of pre-determined subsets, what matters is the relationship between the model assumed by the subset and the object at hand. It might or might not, in other words, be sufficiently specific to the task at hand. Likewise a specific customisation: it matters whether the customisation was created pre-emptively (based on other, similar documents; or, worse, on hunches).¹⁰³ Where the

¹⁰² That HTML does not (typically) ‘fail’ to interpret things that are beyond its scope — basically, everything to do with complex text encoding and presentation — is, I think, the fact that it is the *only* means for presentation on the Web, and thus editors ‘coerce’ their documents into its format. (Added to this is, of course, the fact that HTML is not consumed as such, but rather transformed by the browser — thus styling with CSS classes enable the visual distinguishing of what is, semantically in HTML, identical.)

¹⁰³ TEI-all (versus, say, TEI-lite) can be compared to using Unicode (versus ASCII). It provides a large swathe of options, of which only a particular set are (most likely) ever going to be used. Therefore, the *specific, adapted model* is the elements actually used in the way that they are actually used, as part of the overall structure. As such, conducting a thorough preliminary analysis of the whole document and producing a more restrictive schema based on requirements, or using TEI-all and retrospectively cutting down the schema to what *was* used, are equivalent.

model is highly critical, it is less likely to be sufficiently specific in a given instance; or the contrary, it may be general, and therefore less critical. Model criticality and specificity can, accordingly, be seen as a more generalised account of what Cummings (speaking specifically of the TEI) has described as the “Unmediated Interoperability Fallacy”: the idea that, for a model to be applicable to a wide range of things, it must be suitably generic, encoding only more basic structures; or highly specific, losing interoperability.¹⁰⁴

The above are, it should be very strongly emphasised, intended to illustrate the distinctions between various interpretative models, and particularly the result of developing such a model from interaction with the input itself; it is certainly not intended to be an exhaustive list of possible interpretations, or a particularly deep analysis. That said, I would suggest that all interpretative models applied can be placed on this scale of specificity, and that this position will correlate strongly with the capacity of a model to produce critical, formalised interpretations.

Beyond the extremes of this specificity spectrum (the hyper-specific at one end, and the totally non-interpreting at the other) there lies a substantial range of models that might be considered *archetype-based*. There are, as we have seen, some broadly placed, and non-project-specific, archetypes that are sufficiently generic as to not be questioned: photography is one such example. More interesting examples emerge as we move closer to the *specific* end of the spectrum. These kinds of archetypes (‘true’ archetypes, we might say) are not drawn externally from the transmission process, but emerge from a limited part of it (but not, by definition, the entirety).

¹⁰⁴ From Cummings, “Customising the TEI”:
<http://slides.com/jamescummings/customisingtei/#/3/4>

Such archetype-models might be established, for example, from studying a small subset of the selected documents. When customising an encoding schema (an interpretative model) for a large set of documents, the editor may base their work on a smaller subset — one or two documents — that they take to embody the features of the larger set. That this happens to be the case is necessarily an assumption, and relies on the homogeneity of the full set. (It must be an assumption, as confirming it and making alterations to the schema would involve looking at the entire set, and thus the model would no longer be archetype-based.) When this assumption proves to be incorrect, we find the same two outcomes: failure (the model is unable to interpret; like a wood-chisel on iron, it just bounces off) or an over-determination of the input (i.e. the input is made to fit the model; a sledgehammer wallops whatever it comes into contact with). Of course, if it were possible — if the process enabled a recursive looping-back to the point at which the model were established — this ‘incorrect’ assumption could be rectified.

To give a concrete example, when editing the *Woodman Diary*, the editorial team of which I was a member produced a preliminary TEI schema (using Roma), discarding myriad modules which were felt — on examination of a few pages of the diary — to be unnecessary. Most notably, we threw out elements relating to the mark-up of poetry (there was no poetry in the pages that served as the model archetype). Mid-way through the encoding of the first volume of the diary, however, we came across a section that quoted a poem (a verse from Shakespeare). At this point, three options were open to us: to skip the quoted verse entirely, because it could not be encoded correctly in the prose-oriented TEI schema; to ‘fudge’ matters, simply disregarding the ‘poetry’ interpretation and making it fit the existing model (encoding it as a curiously short paragraph with peculiar line breaks); or, as it was possible within the organisation of the project, to return to Roma, add back in the appropriate models, and encode the verse properly. In this case, being a small and self-contained project,

we chose the latter course. But in doing so, it should be recognised, we ‘re-drew’ the process diagram, transforming two linearly connected functions (generate the schema, encode the data) into a larger recursive function (or, properly considered, two mutually-recursive functions: the ‘encoding the data’ phase serving as a validation of the schema). Equally, though, we could have taken one of the other options — or may have been restricted from doing, given a particular configuration of the transmission process, and the effect of this would have to be recognised as producing a different output, or not producing an output at all.

There are also instances in which archetype-models might be continually accumulating. We saw such examples in the earlier section on inter-path connections: a model takes what has gone before (chronologically) as additional input for new inputs, but without retrospectively altering that which has gone before.

And what is the effect in such cases? This is always difficult to tell. It lies, ultimately, in the difference between the foundation of the model (what was taken as an archetype) and the input at the point of the encounter. But the fundamental point for any model that does not arise recursively, autopoietically from the encounter itself, is that it imposes an interpretation that is based on something else. Thus it is a normalisation, the adoption of a lowest-common denominator, which suggests a lower criticality.

3.4. Selection of documents

In the previous section, an approach was developed to model the transmission process by decomposing it into a series of stages, with the output of one stage forming the input of the successive stage. In doing so, however, we bracketed the question of the first input: the source documents. In Greetham’s excellent and comprehensive overview of the field, *Textual*

Scholarship: an Introduction, over half the book is dedicated to bibliographic activities that pre-empt explicitly editorial activities. (Greetham 1992) Kline, likewise, dedicates preliminary chapters of *A Guide to Documentary Editing* to the collection of source material. (Kline 1995) That said, the selection of documents in textual scholarship is, typically at least, taken to be a bibliographic task that pre-empts 'editing-proper'.

As Greg argues, the meaning of "written or printed signs" on the page is no concern of the bibliographer. (Greg 1966, 247) Applied to the transmission model of textual material elaborated in the previous chapter, this suggests that the physical documentary basis of the source-assemblage pre-empts the application of interpretative systems. In other words, documents are selected first and then an interpretative model applied second. However, this is to suggest that document selection can be carried out independent of interpretation, a view that Eggert rejects: "[B]egin to perform any bibliographical operation involving marks on the page and one finds they are no longer arbitrary. They have to be raised onto a textual plane, to be given meaning, if only a provisional one." (Eggert 1994, 1)

In this, I am strongly inclined to agree with Eggert. To even necessarily speak of a document is to bring certain interpretative actions into play. What we might call a document exists *virtually*, a set of affects that are *actualised* by interaction with an editor. In other words, a document is what happens when an editor gets hold of a piece of paper. (It is something quite different if a child with crayons gets hold of it, whether the piece of paper is a gas bill or a fair copy of *Ulysses*.) Thus, even at a trivial level, an editor who sets out to edit a work must be able to interpret a form of that work from the selected documents.

As such, we must posit that the selection of documents is itself an interpretative act. Even if, as Greg's words might be interpreted to suggest, the bibliographic work is carried out by

someone else at an earlier stage, their actions must be considered as an interpretative foundation that cannot be disregarded: it happened; it is part of the chain of processes. Moreover, in other forms of textual transmission — including many (digital) scholarly editions — the roles of bibliographer and editor are more closely intertwined. In many cases, there may not be any separation at all.

On the surface, this leaves us with something of a paradox: if selection is interpretative, we must interpret a document in order to select it; but to interpret it, it must have already been selected. However, I would suggest that a way out of this paradox is to recognise that, in the first instance, what we interpret is not ‘the documents’ (or whatever is selected) but the world. From this starting point, ‘interpretation’ can be conceptualised in a similar way to the interpretation *of* the material documents discussed in the previous chapter: that is, as an encounter between an interpretative tool (a model) and an intensive assemblage, which stratifies the assemblage and extracts extensities. In this case, the assemblage is the *world*, and it is stratified into *relevant documents* and everything else. Indeed, it can even be suggested that it is this action that ‘creates’ documents in the first place: not that documents do not exist, but that their precise form and boundaries are determined by interpretative action. Like inscribed text, which exists in some way but is elevated to such a status by its actual response to a particular interpretation, so too documents. A document is delimited as a discrete entity by the act of selection, ‘becoming’ a document by its being selected as such.¹⁰⁵ We might call this selective interpretation a *material-delimiting model*. In its complete operation it is much

¹⁰⁵ The previous chapter discussed Foucault’s questioning of apparent ‘unities of discourse’ — what constituted a work or *oeuvre*, for instance. The same is, I think, true of documents. Our notion of what constitutes a ‘whole’ document is always open to debate: is it a single sheet of paper? In which case, are both sides of the page separate documents? (And in this case, is there a fundamental difference between a text written on one side of two sheets of paper rather than both sides of one?) What of pages bound together?

the same as any interpretative stage discussed in the previous section. It is a stage that outputs either documents or nothing.

In some circumstances — where selection is largely self-evident — it may be pragmatic to view this as a single stage or function (it can always be seen as such); in other cases, however, decomposing its operation is necessary to understanding later stages. This is because selection is not necessarily a single ‘event’, even if it can be modelled as such with hindsight. Again, Massumi’s example of chiselling wood is relevant: we might say that a woodworker has selected the ‘relevant’ pieces of wood to fashion a table from a large block, but this selection was hardly carried out in one single action. Likewise the selection of documents: its effect might be to stratify the world, its application — an interaction — can hardly be with the entirety of existence in one fell swoop. There is, therefore, another side to this stratification: the pragmatics of its application. Its ability to stratify a portion of reality is dependent on an encounter between the stratifying model and that portion of reality actually taking place. This is, I think, fairly obvious. An editor must interact with a document (or possible-document) in order to select it. As such, material delimitation has two sides to it: a conceptual model and its pragmatic application; a *what you aim to select* and a *where you look for it*. Thus, instead of viewing the material-delimiting model as a stratification of the world, we may decompose it into a repeated action applied successively to a series of possible-documents, which it either selects or does not.

The kind of material-delimiting model applied evidently has an effect on future transmission stages, which depends largely on the interplay between the conceptual and pragmatic sides of selection. Each side acts as a kind of constraint on the other. In general, the stronger and more specific the conception, the weaker the effects of its pragmatic application. This may therefore be considered a continuum, as illustrated in Figure 27 (below).

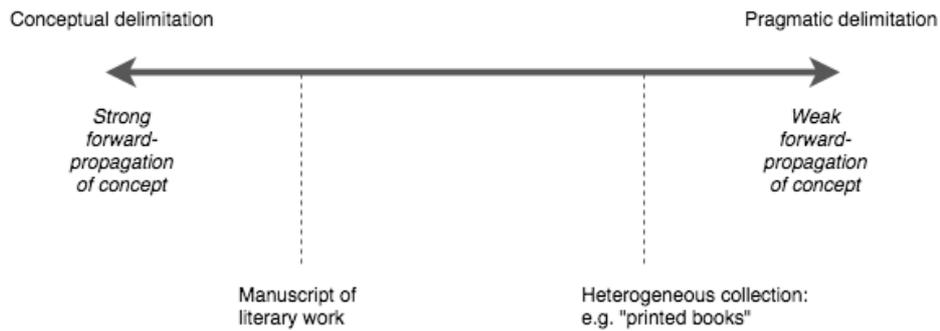


Figure 27: Material delimitation continuum: each material delimitation is in part based on some conceptual notion, and in part on its actual application.

There may be a *strong* conception that results in the selection of a highly restricted set of documents. Aim to select the documents that correspond conceptually to a literary work and the choice is narrow: either a very specific set of documents is found, or it is not. The effect of such specificity is to foster the impression of a kind of self-evidence — but it is only an impression. It is really the effect of previous interpretation: the *Beowulf* manuscript points to *Beowulf* because someone else has previously discovered it, studied it, catalogued it. In such cases, the conceptual determines the pragmatic: if you want to edit *Beowulf*, get the *Beowulf* manuscript (and, once you have the manuscript, edit *Beowulf* from it).

In this situation, the separation between the material-delimiting model and later interpretative models seems somewhat redundant; but it is simply the case of one very strongly implying the other. But, I would argue, this separation ought to be maintained, for two reasons: firstly, that the forward-propagated effects of the former are never total (this would suggest that a document, once selected, is entirely self-evident); and secondly, in other cases there is a clear separation, not just in theory but in practice.

We see this most readily at the other end of the spectrum, where material delimitation is guided by a weaker concept and, as a result, places greater weight on pragmatic application. We see this most readily in mass digitisation projects such as Google Books. A conceptual

delimitation of ‘all printed books’ (to somewhat roughly caricature) is so wide-ranging that, pragmatically, where one actually looks for these books plays an immensely important role. In such circumstances, this material delimitation may be more readily seen as purely preliminary: the conceptual delimitation does not carry through into later interpretative processes to the same degree; it is less definitive. As such, we see a much clearer separation between the *material-delimiting model* and later *interpretative* models.

With regard to this material-delimiting model and its pragmatic application, a number of further points should be considered. The first is that the poles of the continuum (see Figure 27) are not valid points for any real transmission activity. At one end, documents cannot be selected solely in theory (but not in practice). Put simply, an editor may conceptualise what they are looking for, but without wielding the conceptual implement (i.e. actually going out and looking for documents to match) it will produce no documents. And at the other end, the editor can have no way to stratify existence, and will end up selecting everything — setting themselves, in other words, the task of editing the world.¹⁰⁶ There are no ‘purely conceptual’ or ‘purely pragmatic’ modes of document selection.

Secondly, as considered here, the material-delimiting model must, at a minimum, select documents in their entirety. This is because the selection of a document is a real action that operates in a physical dimension. Therefore, it must — short of physically tearing the

¹⁰⁶ This latter point is, perhaps, more contentious. However, I would suggest that even extremely loose conceptualisations such as “all the documents in the world” are still conceptualisations: they separate documents from non-documents (however that might be defined). In this, we must bear in mind that ‘selection’ is really the stratification of an intensity — in this case, existence — and the extraction of extensive components; without a conceptual means of stratification, this cannot take place. More insidiously, the editor may set seemingly non-signifying conceptualisations, such as “all the documents in that box” — but this, too is a conceptualisation, albeit one created by whoever (whether deliberately or accidentally) put the documents in the box.

document in two — either select the whole document or not: the choice is binary.¹⁰⁷ This is not to suggest that later interpretative models cannot extract portions of a document. Indeed, such basic operations as transcription — the stratification of the document into linguistic text and material substrate — ‘extract’ portions of the document. Likewise separating the portion of a document written in one scribe’s hand from another, even if they are physically collocated: one at the top and one at the bottom of the page. But in the model presented here, these are all *interpretative models*, ‘plugged in’ once a document has been selected. This is to take the extensive document and reconsider it as intensive. Moreover, while it is true that more conceptually-based selections tend to elide material delimitation and later interpretation, insofar as these can be separated, what distinguishes them is the purely physical extensities on which the material-delimiting model operates.

3.5. Delimitation regimes

If selection is determined both by a *what* (a conceptual delimitation) and a *how* (a pragmatic delimitation) — both in concert, to varying degrees — what effect does any particular combination have on the outcome? The first thing to note with regards to this analysis is that, while Figure 27 (above) conceptualised the relative importance of the conceptual and pragmatic as a continuum, real selective action is restricted to a set of particular, discrete configurations.

This is a result of selection producing discrete, extensive entities (typically, physically delimited documents), with particular types of connection between them. They can be seen

¹⁰⁷ This raises an interesting question with regards to born-digital documents. To what degree do they really have ‘physical extensity’? Nevertheless, to talk of ‘born-digital *documents*’ is to presume an extensity (even if it requires further mediation to determine it) rather than arbitrary, ‘intense’ bits of data read from a hard-drive. This question is taken up more fully in Appendix A of this thesis.

as instantiating one of three conceptual systems, or *images of thought*, articulated by Deleuze and Guattari in *A Thousand Plateaus*: the *tree*, the *radicle*, and the *rhizome*.

3.5.1. Sequential delimitation

The first is *sequential*. In this kind of assemblage, documents-as-extensities are subject to an overarching sequential (or linear) unity, a macro-level extensity. This linearity may be brought about by a selective mechanism (material-delimiting model) that imposes a logical order to the material documents. Most typically, this means a set of documents connected by the flow of text across documents. A printed novel demonstrates this readily: quite aside from a physical unity (reinforced by binding, covers and so forth), there is a linear flow across page boundaries. Indeed, this would persist should the binding fall apart and the covers disintegrate. The same is therefore evidently true of unbound manuscripts of a novel: interpreting a single linear text across individual documents clearly orders them. We do not even require a linear *text*: any signifying system that orders the documents in some sequential manner is sufficient. A set of manuscripts full of revisions and notes — that is, without some clear overarching textual sequence — can be ordered sequentially by some conception of genesis (an author wrote page *x* before page *y*).¹⁰⁸ A sequential delimitation delimits precisely because there must be a beginning and an end: one arrives at a point where there is nothing left to add in the sequence.

¹⁰⁸ It can, indeed, be argued that a full genetic analysis of a manuscript divides a single surface into multiple ‘documents’. The surface is a synchronic notion, while interpreting an act of composition creates a ‘set’ of documents diachronically. It clearly depends, therefore, on the interpretative system employed whether a surface is one or many documents. The same, incidentally, may be said of cross-hatched writing (see the previous chapter): the linear text posits two ‘documents’ (determined by surface-orientations) which may be sequentially arranged (the author wrote in portrait *first* and landscape *second*).

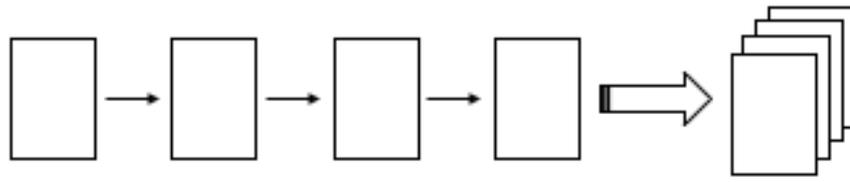


Figure 28: *Sequentially ordered documents delimit a single extensity*

This is, incidentally, the case too for a text contained in a single document. It is a linear delimitation, of length one; the point at which there is nothing to add to the sequence is arrived at by default.

As observed above, the tendency of this sequential connection between documents is to override the degree to which individual documents (however they may be conceptualised) may be considered molar extensities. The *Woodman Diary* is again illustrative. At the level of the *page* (or surface), the discreteness of a particular page is subordinate to the flow of text. At a higher level, the same is true of the two discrete notebooks that comprise the diary: evidently they each have a material unity, but once selected on the basis of textual sequence, a document-assemblage is created that subordinates physical (dis)unity to some semantic interpretation. It is as if — conceptually, in thought — the back cover of the first volume and the front cover of the second volume were removed and the two volumes stitched together. In such circumstances, what we are dealing with is really one document. More precisely, the preliminary delimitation transforms multiple possible documents into what must, with regard to the functioning of the transmission process, be treated as a single entity.

This is true whether the editor assembles the documents in order and delimits the extent of the sequence, or whether the editor selects the set of documents already delimited. In other words, the set of documents is already subject to a *regulative idea*, or the regulative idea is constructed out of the documents themselves. (As noted above, the former is really a

previously carried out example of the latter: someone must have constructed the regulative idea — the work — out of the documents.) But how might this occur in this case? We can model this in terms of the application of a recursive selective function. The editor selects one document in the chain, which feeds back into the selective function, restricting selection of a second document to one whose text continues on from the first (or which precedes it), which feeds back again; the sequence of documents is built up like a game of dominoes.

From this, we may suggest several properties of a linear delimitation. Like all delimitations (as we will see) it is in its operation an *accumulation* of material. In this case, it is ‘open’ for accumulation in only one dimension — linearly, along the *x*-axis. There is no possibility within this regime to extend in any other dimension. At the same time, it marks the *synthesis* of documents into a single entity, under a regulative idea that is constructed out of it. Moreover, being restricted to expansion in a single dimension, it is easy to identify a point at which the chain — and thus a regulative idea — is complete: i.e. when there is no further document to be accumulated.

This reduction of multiple molecular assemblages into a molar extensity undoubtedly makes the resulting assemblage an instance of a *tree* or *root*-like system. The tree is one of three *images of thought* described by Deleuze and Guattari in *A Thousand Plateaus*. (Stagoll, 2005, rather than ‘image of thought, describes an “arborescent schema” — perhaps a more usefully prosaic term.) A tree is a hierarchical structure, proceeding from a root, by a process of division: “Binary logic is the spiritual reality of the root-tree.” (Deleuze and Guattari 1987)

According to Robinson and McGuire, the tree or *arborescent model*:

is a fundamental intellectual model for much of western thought, stemming originally from Aristotle’s “classic theory of categories”, which in essence propounds that entities are placed into the same category, by

rational division, according to an objective assessment of shared characteristics (Robinson and Maguire, 2009)

As a structure, therefore, it can only lead to the replication and repetition of thought: it leads, to borrow Adkins' terms, to *good sense* (a single direction of thought) and *common sense* (stable entities). (Adkins 2015, 22). As such, the tree is a "marvel of stable, hierarchical organisation. Lines of descent are always clear, as is the process of differentiation." (Adkins 2015, 23) Deleuze and Guattari give a number of examples: Chomskian linguistics, with its grammar-trees, divides the sentence into clauses, and then into words, each with a given function; biological classification; and Saussurean structuralism, dividing the sign into signifier and signified (Deleuze and Guattari 1987, 6) Trees also pervade such fields as information science and the digital humanities. The *Functional Requirement for Bibliographic Records* is a hierarchical system;¹⁰⁹ as is the OHCO (Ordered Hierarchy of Content Objects) hypothesis of text (De Rose *et al.* 1990).¹¹⁰

The tree is also, as the previous chapter suggests, the structure of the *signifying regime*. The *regulative idea* created is the root node of this system, from which all other interpretation flows, hierarchically subordinate to it. This is the regime in which representation, as opposed to merely heterogeneous connection, is possible. But, most importantly, it should be asserted that this is not a universal, pre-established truth, but rather the result of specific interactions between assemblages and the potential of one to overpower another. Trees are created by interpretation that delimits a given assemblage, gives it form, and elevates it to a signifying

¹⁰⁹ The *work* (the intellectual creation) is realised in an *expression* (a specific, though 'abstract', form), which is embodied in a *manifestation* (for instance, a particular edition, of which there may be many copies printed) and the *item* (a single instance: a particular copy of that edition). (International Federation of Library Associations and Institutions 1998)

¹¹⁰ The OHCO model has been highly critiqued; see Renear *et al.* (1993); Huitfeldt (1994); Pierazzo and Stokes (2010).

status. It is the take-over of a manuscript that under interpretation yields (say) the text of *Beowulf* by the *regulative idea*, the work (the signifier), *Beowulf*. And *Beowulf*— the regulative idea — both interprets and delimits: it is the critical, cutting edge of a model whose first effect is a binary division of the world into *Beowulf* and not-*Beowulf*¹¹ (which includes, it should be noted, editions of *Beowulf*: anything interpretable by the model). It is thus, in the terms established earlier, a highly critical and specific model.

So far, we have really only talked about material delimitation — the selection of documents — its consequences are already numerous for later transmission stages. As a strongly conceptual delimitation model, it shuts down the pragmatic possibilities of later stages, at least without undermining the regulative idea created (in which case, it no longer functions as a regulative idea).

Accordingly, any interpretative systems into which the material assemblage can be plugged — the subsequent ‘modelling’ delimitation of the whole interpretative assemblage — must be done wholesale. That is, the editor is constrained to the entire set of documents (the ‘work’, if the term is appropriate) as a work, and not as individual documents. Modelled as a process, then, this means a single path. Evidently, photographing the document demands the invocation of multiple paths; and the editor must work independently to a degree on separate documents or sections. But there should also be a later merging function that joins these paths and re-establishes unity under the regulative idea.

¹¹ This, incidentally, is the reason hierarchical structures such as FRBR are restrictive: they demand to model things that are amenable to it. For instance, Pierazzo outlines many problems with FRBR with relation to manuscripts: the elision of *manifestation* and *item*, given the uniqueness of manuscripts; and the problems of transmission and circulation. (Pierazzo 2014, 63) These problems can be seen as arising from the application of a hierarchical model outside the domain for which it was intended. And potential solutions (such as the FRBR-object-oriented taxonomy) can, by the same token, be seen as the *adaptation* of a model (recursively, as it were) to the object in question.

Nor can there be the possibility of later accumulating more documents under the regulative idea. While it was stated earlier that a linear delimitation has a single dimension, it is not that the dimension existed previously as a natural property. Rather, the linear dimension exists as a kind of possibility: one of many ways in which the documents might be configured. It is the effect of a particular kind of interpretation; but neither can it be overridden without also overriding that original interpretation. Adding another document that is not part of the sequence would be to add another dimension.¹¹² By the same token, neither is it possible to later remove documents from the sequence.

Most significantly, the text (*the* text, the text by which the documents were selected and sequenced) is elevated hierarchically above other ‘texts’ — such as marginalia — on the page (see Figure 29, below). Of course, the hierarchy could be reversed, by selecting documents based on marginalia, though this would give rise to a quite different regime (assuming the accumulated marginalia could be given a discrete form). Likewise, the documents, materiality and bibliographic codes; they become subject to the text. Or, more precisely, they become subject to the regulative idea that was constructed on the basis of the text (and to which the text itself is also subordinate).

¹¹² To give a trivial example, it would be to delimit an assemblage of ‘*Beowulf* and some other papers I found’: the whole could no longer be considered *Beowulf*, or interpreted as such. Of course, it might be possible to construct an alternative regime, though it would not be the regulative idea of *Beowulf*, but something radically different.

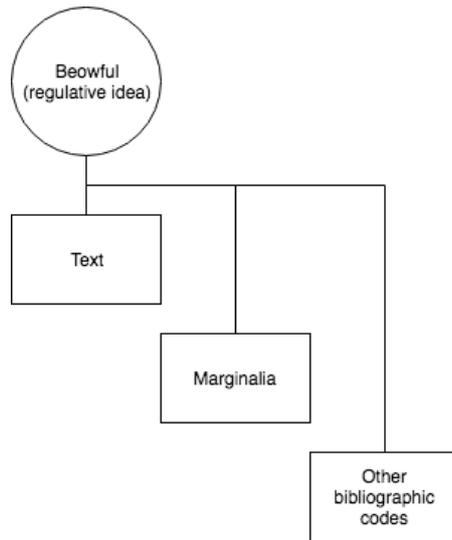


Figure 29: Hierarchical importance of the selective regime

Further interpretative functions in the transmission process contribute to the development of this hierarchy: each stratification marks a binary division between that which is formalised and extracted and that which is not. The regulative idea is thus unfolded throughout the transmission process. But — and this is the most important thing — each further stratification must be carried out in line with the regulative idea as it is established up to that point (i.e. recursively, with reference to it). To do otherwise is to undermine the regulative idea. This is the distinction between a reflexive transmission and an articulated transmission. At any stage in the transmission process, the editor, acting unilaterally (i.e. without reference to the regulative idea as hitherto developed), can interpret in any way they wish, and articulate those actions. But if those actions undo the regulative idea, the status of the regulative idea as ‘that which is represented’ is no longer valid. Representation is reduced to being only partial.

Though in the above discussion I have used the example of the *text* being the primary mechanism for producing a sequence of documents, this could in reality be any system that produces a single, logical order. If the first draft of a novel were written in no particularly

coherent order (so that the sequence of the text jumped from the beginning, to the end, to the middle, say) it would be equally possible to delimit this as an ordering of the text (if it could be constructed) or as an ordering of the pages (a bound notebook keeps its pages in order). While the former, if we had to apply a label to the regulative idea, would be ‘the text’ (of the novel), the latter would be ‘the notebook’ (which contained a draft of the novel).

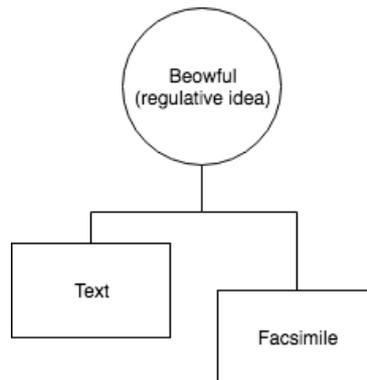


Figure 30: 'Pluralistic' model (one conception is always necessarily hierarchically elevated)

Of course, digital systems of reproduction (most notably the Web) allow either perspective to be dynamically generated; but these, however easy it is to instantiate on-the-fly, must be seen as functioning in parallel — as separate versions of the document/text. They cannot be collapsed down to a single logical order. The same is true of digital editions or archives that provide a text alongside a page facsimile, allowing a user a view of the transcribed text and ‘access’ to the visual aspects of the page. There are two ways of approaching this scenario: one is to suggest that, actually, two regimes — two regulative ideas — are created, one for each view. This is, I think, true; each stands legitimately on its own, and the hierarchical relation is levelled, at least to a degree. (See Figure 30, above.) In this case, therefore, the regulative idea really points to a plurality: the work really *is* the linguistic text abstracted from the page, *and* the page itself. But it is a plurality that is contained by and subordinated to a singularity.

In the next section, we will deal with delimitation regimes that build in plurality by default.

3.5.2. Parallel delimitation

The second kind of delineating-model creates a *parallel* relationship between documents. This is the relationship arising from interpretation of the documents on a basis of identity. Typically, this is the result of interpreting the same text encoded in the documents concerned — in other words, that the documents contain versions of the same text.¹¹³ As before, we may consider this in terms of process. It is all very well to say that the editor ‘selects versions of a text’, but this is dependent on the text itself having been already formalised, and the status of each version as a version of the *text* having been already established. But how this happens needs explaining. The first stage, then, is the delimitation of a single ‘version’ (as a linear sequence, discussed above). This version is then selected, and is fed back into a recursive selective function as a model for further selection. In other words, the first version serves as an archetype for the selection — on the basis of interpreted textual similarity — for selection of other linearly sequenced ‘versions’. This makes for a very different regime. Each version closes down the possibility of extension along the x-axis (once a version is complete, it cannot be added to), but there is no closure along the y-axis: there is always the possibility of adding a new version, assuming it exists. There is, therefore, a *conceptual* synthesis (marked by a closure along the x-axis), at the same time as *actual* accumulation along the y-axis.

The material-delimiting model therefore does not produce a single unity, but two tiers of unity: a macro-level conceptual unity (which would typically be considered the *work*) and the component unities of individual versions. The macro-level unity, a regulative idea,

¹¹³ Or, more conventionally put, that the documents instantiate the same work — though, as the approach taken in this thesis suggests, this implies a work that exists as a transcendent phenomenon somehow ‘awaiting’ instantiation, rather than being the result of interpreted similarity between documents.

unfolds over time as more versions are added. But the component-unities are themselves fully delimited, and thus subject to their own regulative ideas.

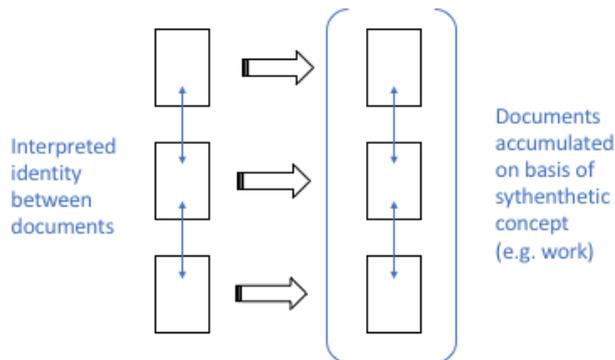


Figure 31: Parallel delimitation determined by identity between documents, transmitted as independent versions conceptually synthesised (the bracket denotes a conceptual 'work')

This kind of system is called a *radicle* by Deleuze and Guattari. Such a system is not, intrinsically, different from a tree-system: it obeys the same logic of systematic division based on characteristics at each level of a hierarchy. According to Deleuze and Guattari, the radicle, or “fascicular root”,¹¹⁴ is a tree-like system in which the “principle root has aborted, or its tip has been destroyed; an immediate, indefinite multiplicity of secondary roots grafts onto it and undergoes a flourishing development.” (Deleuze and Guattari 1987, 6). While it may appear to be ‘multiple’, Deleuze and Guattari argue that this is really a kind of pseudo-multiplicity: “We must ask if reflexive, spiritual reality does not compensate for this state of things by demanding an even more comprehensive secret unity, or more extensive totality.” (Deleuze and Guattari 1987, 6) It is, therefore, not so much a multiplicity as a plurality.

¹¹⁴ A *fascicle* is defined by the Oxford English Dictionary as, “A bunch [or] bundle” — in other words, the multiple ‘tied together’ to produce unity (*fascēs* being the Latin for ‘bundle’). Pertinently for this debate, a *fascicle* is also a single segment of a book or work published in instalments (that is, a unified literary work distributed in its physical dimension over multiple ‘objects’) (OED).

They provide several examples of radicle system, such as works by William Burroughs using his “cut-up method” — the cutting up of existing linear texts on paper, and rearranging them to make a new work. “Folding of one text onto another [...] constitutes multiple and even adventitious roots [...] implying] a supplementary dimension to that of the texts under consideration.” Despite the apparently ‘multiple’ nature of the new work, “[i]n this supplementary dimension [...] unity continues [in] its spiritual labour.” (Deleuze and Guattari 1987, 6) Likewise Joyce, whose words “shatter the unity of the word, even language, only to posit a cyclical unity of the sentence, text or knowledge.” (Deleuze and Guattari, 1987) The resulting unity must be all the more totalising, to hold together a quasi-multiplicity of elements, via the creation of ever higher numbers of dimensions. There is no true break, therefore, with tree-like division.

We see this situation readily in the example above: there is, at one level, an accumulation of versions, but by appealing to — actually constructing — the notion of a version, a higher unity is constructed, ‘tying together’ in a bundle the component parts; and from this higher unity, there is a hierarchical division into multiple versions.

But what does this mean for the transmission process? The principle question is, to what extent the macro- (work-) level regulative idea necessarily overrides the component- (version-) level regulative ideas. What latitude is granted to later interpretation? Evidently, there must be some constraint applied by the macro-level regulative idea: the grounds on which the component parts were tied together cannot be undone without destroying the basis for their selection. At the same time, each component version can invoke its own transmission path, which can remain distinct and subject to its own regulative idea, which may be developed by interaction with that component alone. The component-level regulative idea may, therefore, be freely developed, but one of its initial inputs is the macro-level regulative idea. As such,

each version is a reflexive transmission — or can be, provided the regulative idea is maintained, unrolled throughout process — but it makes a slightly different claim to the regulative idea of that of a single document in isolation: it is the document/text of that version insofar as it is a version of the text/work. And, significantly, the transmission path of each version may be retained into the output. In this case, the macro-level regulative idea asserts a plurality: in other words, that the work *is* its multiple versions. A *variorum edition* is one such example.

Conversely, there is the possibility that the versions can be later ‘collapsed’ into one single version (i.e. by creating a critical edition). This, obviously, demands a merging function. It operates by stratifying the versions themselves on some basis (e.g. final authorial intent, ‘best text’), and discards the non-selected variants (or hierarchically subordinates them, by placing them in a critical apparatus). Such an operation is an actual, as opposed to merely conceptual, synthesis, and marks a return to a root-structure. Of course, the intention to perform such an operation invariably places greater restriction on previous interpretative processes: each version must be treated in a homogeneous way, if they are to be merged (e.g. for a collation to be effected).

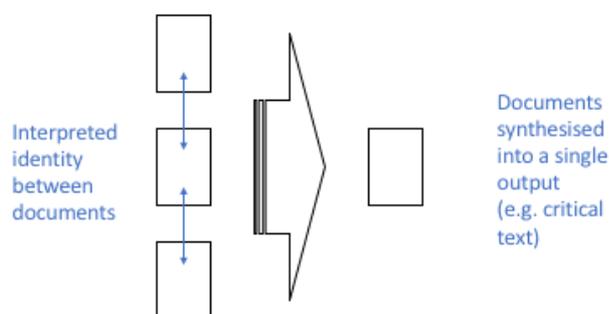


Figure 32: actual synthesis

Finally, there is the question of completeness. A linear delimitation, as we have already seen, is necessarily complete in its one dimension. A parallel delimitation, however, is extensible

along the *y*-axis, and thus has the potential to be extended (and at a later point in time) if another version were discovered. Of course, a truly reflexive transmission, carried out exhaustively, would have uncovered all versions: thus the totality of the work would be present (the regulative idea would be complete; *y*-axis additions exhausted).¹¹⁵ Even so, I would argue that the addition of a newly uncovered document, while it undoubtedly adds to the plurality of the texts, does not radically change the regime: as a hierarchical system, the regulative idea still subordinates all versions, including any newly added versions. In the case of a critical edition, where the versions have already been collapsed, this depends on the status of the new document. If (assuming, for example, that we were editing on the basis of final authorial intent), the new version were shown to be non-authoritative, it could simply be added to the critical apparatus without disruption to the text. In other words, it does not challenge the hierarchy already imposed. (On the contrary, a *more* authoritative version would challenge the imposed order, and demand a new critical version.)¹¹⁶

3.5.3. Disjunctive delimitation

A disjunctive delimitation selects documents (or sequentially delimited sets) on the basis of a disjunctive relationship between them — in other words, a *non*-identity. It is the type of preliminary delimitation seen most typically in archival and other digitisation projects.

This, it should be noted, is the default position in the selection of any set of documents: outside of any interpretative activity, the relationship between any documents in a set is

¹¹⁵ It might be argued that, with regards to a linear sequence, there is the possibility of uncovering a document that could be attached to either end; but at the same time, there are other clues available, not least the completeness of the text. But a set of versions of the same text can offer no clues as to how many other versions exist.

¹¹⁶ This is, of course, intended to illustrate the notion of completeness; whether the editor could, in practice, add more versions or modify existing versions depends on the medium.

always one of difference rather than identity.¹¹⁷ The point is, rather, that such a preliminary delimitation is unable, based on an interpretation, to connect the documents either sequentially (as an ordered set) or in parallel (as multiple versions). Accordingly, it is distinguished from these previous types of delimitation the fact that the selected documents cannot be collapsed down or synthesised into a single entity, either a single sequential ‘document’ or a single critical version. It is, therefore, accumulative by default, and hence much more determined by the pragmatics of selection.

As argued earlier, however, selection can never be determined entirely by pragmatics. There needs to be some criterion by which to stratify the world; to determine, if nothing else, what counts as a document. Even with such a definition in place, to select documents entirely at random is, if not completely absurd, an edge case which I think can be safely ignored. But this does raise the question of how documents are selected conceptually, and what kind of regime is created in theory and practice by its application. The easiest way to begin to conceptualise this is to suggest that the editor determines a selection domain: not specific documents, but a conceptual space from which to select.

Any criterion can, in principle, describe a conceptual domain. It may be documents written by the same person, or documents relating to a specific subject. Wider domains might include the contents of a particular archive or collection, or, as in the case of Google Books, all printed books. A domain might be described as the intersection of multiple criteria, as in the *Letters of 1916*: “epistolary documents written between November 1915 and October 1916 with a tangible connection to Ireland” (Schreibman *et al.* 2017) The narrowness or broadness of the domain determines the weight placed on the pragmatism of selection from that domain.

¹¹⁷ See the previous chapter for a fuller discussion of difference preceding identity in the works of Deleuze (especially *Difference and Repetition*, 1968/1994) and Deleuze and Guattari.

Documents written by a single author is more conceptually focused, leaving less room for pragmatism, than ‘all printed books’, which is much more dependent on how actual selection is carried out.

But considering a domain in such a way is to assume that it already exists as a recognisable entity, that is, taking an already-extensive view of things. If a domain already exists, to be simply pointed to by the editor, it is necessary to ask how such a thing was determined in the first place. Whether it is the editor, or a bibliographer, or an archivist, someone must have determined the extensity of the conceptual space from which to select. We can again model this as a process.

To begin with, we may suggest that documents form heterogeneous connections with numerous other entities. As discussed in the previous chapter, it connects to an author, to language, to myriad other concepts (to Ireland, to a printing press, to an archive). These potential connections are the *virtual* components of the document: real connections, whose force is felt once they are *actualised*, that is, once we begin to consider them. The process of delimiting a domain therefore is rather like tracing the paths of a network. We start with one document, and isolate a potential connection — say, to the author. This author authored numerous other documents, which become part of the network. It is the author that connects the authored documents. A domain is therefore delimited by shared connections to a third entity; the third entity, in other words, becomes the selective function to select other documents. This perspective is nothing more than taking an intensive view of a domain.¹¹⁸

¹¹⁸ This idea was explored in the previous chapter. There are two ways of looking at a document: the extensive, which *internalises* authorship as a *property* of the document, and the intensive, which views authorship as the *external* connection to an author. The extensive view sees authorship as a pre-established category — as natural and self-evident — rather than something created by the process of authorship, or *re-created* as the result of the editor re-establishing the connection between the author and the document.

As such, it is open to expansion in any number of dimensions, or in the number of dimensions created by the editor through the heterogeneous connections to various entities.¹¹⁹

This kind of system is what Deleuze and Guattari term a *rhizome*. (The first chapter of *A Thousand Plateaus*, which carries the title ‘Rhizomes’, is a discussion of this ‘image of thought’ and its contrast with hierarchical tree systems.) *Rhizome* is a botanical term, used to designate underground roots of certain plants that spread and produce new shoots in multiple directions from any point, and, when broken, to (re)assemble and produce multiple new plants. Bulbs and tubers, for instance ginger or turmeric, are rhizomes.

In fields of information science, the concept of the rhizome is the one which is most particularly associated with Deleuze and Guattari. According to Burnett:

[T]he rhizome is an acentered, nonhierarchical, nonsignifying system without a General and without an organizing memory of central automaton, defined solely by a circulation of states (Burnett 1993)

Similar concepts have been elaborated elsewhere, especially in the realm of knowledge organisation. Radford discusses Eco’s idea that a “library can be conceptualized as a labyrinth of texts that contains the possibilities for new arrangements.” (Radford 1992, 419) As we shall see, this notion of a “[net where] every point can be connected with every other point, and, where the connections are not yet designed” (Eco 1984, 81) very closely parallels the DeleuzoGuattarian rhizome. Indeed, Eco goes as far as to reference Deleuze and Guattari in his classification of labyrinths (Eco 1984); Eco’s ‘net’-type labyrinth clearly mirrors the

¹¹⁹ That is to say, *pre-emptively*, there are no dimensions, or an infinite number of possible dimensions. Deleuze’s approach is to treat such a system as topographical space, rather than Euclidean space with a fixed number of dimensions. (See DeLanda 2005.)

rhizome, while ‘linear’-like and ‘maze’-like labyrinths may be seen as equivalent to Deleuze and Guattari’s *trees* and *radicles*.¹²⁰ According to Bogue:

[A] rhizome, as Deleuze and Guattari explain . . . is the antithesis of a root-tree structure, or ‘arborescence’, the structural model which has dominated western thought [...] Arborescences are hierarchical, stratified totalities which impose limited and regulated connections between their components. Rhizomes, by contrast, are non-hierarchical, horizontal multiplicities which cannot be subsumed within a unified structure, whose components form random, unregulated networks in which any element may be connected with any other element.

(Bogue 1989, 107)

As such, Adkins states, “the rhizome continually creates the new [...] it does not follow a linear pattern of growth or reproduction.” (2015, 23) Instead, “non-hierarchical linkages [are] made pragmatically as they are needed, horizontally or across any number of levels, and linking elements of disparate nature when appropriate, crossing categories.” (Robinson, L. 2010, 605)

¹²⁰ In literature, this conception of a library (and, thereby, knowledge organisation) is described by Borges in the short story *The Library of Babel*, in which a hypothetical library of arbitrarily interconnected rooms contains books with every conceivable ordering of alphabetical characters. Latour’s *Actor-Network Theory* (ANT) also presents a similar proposition. As a basis for social science research, it presents a ‘material-semiotic network’ in which actions by human and non-human actors may be studied as part of a non-hierarchical network of interactions. (See Latour, *Reassembling the Social*, 2005.)

In the introduction to this work, Latour questions the validity of attributing aspects of a given scenario or formation to a ‘social aspect,’ as if ‘the social’ were an external and singular phenomenon; instead, he argues that ‘social’ elements are inseparable, a fundamental part of a complex network that gives rise to formations. Latour also recognised the similarity of this conception with that of the rhizome, remarking that ANT could just as well have been called an “actant-rhizome [sic] ontology”. The rhizome as a concept can thus be seen as finding an application in Latour’s theory.

Deleuze and Guattari go on to outline several characteristics of rhizomatic systems. They are, as we have seen, characterised by *connection* (external connections of phenomena, rather than internalised properties) and *heterogeneity* (the connection of many types of thing). The third principle is that of *multiplicity*. Hierarchical organisation subsumes the multiple into the one, by creating an additional, higher dimension (“Unity always operates in an empty dimension supplementary to that of the system considered”: Deleuze and Guattari 1987, 9). The rhizome, by contrast, has no extra dimension: it exists on what Deleuze and Guattari call a *plane of consistency*. A rhizome is also capable of *asignifying rupture*: it can be broken at any point, without dissolving back into a kind of dualism. Most pertinently, a rhizome must be *mapped* and not *traced*. For Deleuze and Guattari, a tracing is the imposition of pre-determined forms, which is antithetical to the rhizome; a map, by contrast, describes the rhizome as it is: the map “is entirely oriented towards an experimentation in contact with the real”. (Deleuze and Guattari 1987, 13)

To give a concrete example (in terms of documents) of how a rhizomatic system might be formed, let us take academic papers. Starting with one paper, we might look up all the papers that it cites and connect them (imagine drawing a network diagram). From each of these papers, we may draw further connections to the papers that they, in turn, cited; and continue indefinitely. (In all likelihood, we will end up with a complex network of papers citing each other, rather than a tree diagram in which all citations point back eventually to an ‘*Ur-paper*’.) At the same time, we might map other entities and connections: we can add authors to the network and connect them to papers; or keywords (see Figure 33).



Figure 33: Section of the conference logo for Digital Humanities 2014. Authors are connected by co-occurrence of keywords in papers accepted to the conference. (Dario Rodighiero 2014, <https://dh2014.org/2014/06/26/dh2014-conference-identity/>)

It is clear is that this process can be continued indefinitely, by adding any number of heterogeneous connections. But how, then, do we call a halt to indefinite expansion in every dimension? How, in other words, can we delimit a domain? The answer is, by isolating one of the connecting entities and elevating it to the level of a substantive. By adding a supplementary, higher dimension in which this entity may operate to hierarchically subsume other components; by transforming an entity into an (albeit preliminary) regulative idea. For example, intensively considered, there are lots of academic papers that are connected on the basis of authorship to (say) Peter Robinson; elevating Peter Robinson-the-author to a regulative idea creates an extensity, internalising the academic papers as ‘properties’ of Peter Robinson. Accordingly, the rhizome is made subject to a higher unity (under the ‘Robinson-

regime’); it becomes a radical system, a domain from which we can select. Thus the accumulation of documents again becomes a kind of conceptual synthesis.

From this initial position, there are two possibilities. The first is that the entire domain is selected. In this case, the domain criterion functions as the regulative idea of a hierarchical system. If we select all the academic papers of Peter Robinson, what is represented by that selection is, if we had to give it a title, *Peter Robinson’s Academic Papers*.¹²¹ But, the domain being only a preliminary space from which to select, with much left to pragmatic action, there is the possibility that we do *not* select the entire domain. What kind of system is then produced?

From one perspective, it cannot be denied that any papers selected *are papers of Peter Robinson* — by which we mean, they are connected by authorship *to* Peter Robinson. We can also, of course, enumerate the papers that are selected. In terms of validity threshold, therefore, we reach an *articulated transmission*: we can say *what* is selected. But we cannot validate any relationship between what is selected and the encompassing whole. We might have conceptually elevated Peter Robinson to a regulative idea, but that regulator has been defied; it does not regulate in practice. What is selected is not, or cannot be, reflexively validated against the regulative idea we have created, because that would reveal its failure to regulate. Peter Robinson as a selective function should stratify the world into ‘written by Peter Robinson’ and everything else; but it does not. This, it should be noted, is the same as selecting an arbitrary set of features from a document; the features are part of the document,

¹²¹ In this example, it might be suggested that there are really *two* criteria: authorship and some notion of ‘type’ (i.e. academic papers), and that any selected paper represents an intersection of the two criteria (really, a *connection* to both as entities). By carrying this selection to its completeness, it actually is the *intersection* that is elevated to single, regulative idea.

but together do not *represent*; or, rather, only partially represent, and to a degree that cannot be determined.

Moreover, it is impossible, preceding only from what *is* selected, to invoke a new regulative idea — at least, not one that is reflexively validated, as this would involve a return to the point at which we stratify the world, in order to ensure that the regulative idea has indeed selected what it should.

What we have, therefore, is a kind of rhizomatic system: limited in some way, but not to the extent that that limitation signifies. That said, once we stop selecting, an arbitrary limit is reached. Furthermore, we may posit that limit as an entity in itself (it is the ‘collection’, for want of a better word) and connect it up to the selected documents. Evidently, it is connected to all the documents that are selected. The collection is therefore a selective function or model for selecting itself. And thus, if we insist on making the collection signifying, what we are doing is elevating the collection itself to a substantive — which is tautological.

From this, however, it is possible to suggest another way of conceptualising a complete, reflexively-validated selection (again, let us say, of Peter Robinson’s papers). It is the coincidence in a mapping of Peter Robinson-as-author and the-collection-as-substantive: that is, they are one and the same point; the collection *is* Peter Robinson’s papers. Accordingly, it is possible to envisage a scale of degrees to which the collection-as-entity and the selective model coincide. At one end, they completely coincide, and at the other, not at all.¹²² This, I think, corresponds well with intuition. It is plain that a collection of the works of Shakespeare that happens to omit a sonnet or two is closer to being the complete works of

¹²² The latter is the plainly daft case of intending to go out and select the papers of Peter Robinson and coming back with only papers by Hans Walter Gabler.

Shakespeare than a collection comprising only *Titus Andronicus*, *Romeo and Juliet* and *King Lear*. There are, therefore, what we might think of as degrees of tautology: the extent to which a collection represents some other entity, or only itself. This suggests that representation itself is a matter of degrees, and in theory quantifiable; though this, in turn, requires a reflexive analysis of the domain to determine its size. That said, if it must be put in absolute terms, only a completely represented domain represents that domain; a little tautology goes a long way.

To conclude this section, it must be noted that this completeness (or not) is applicable at the level of the whole: whether a radicular or rhizomatic system, it does not (by default at least) override the extensity of the component parts. That is, whether we selected all the documents or not, what *are* selected are still discrete documents, and may thus be made subject to their own regulative ideas. As seen in the previous section, this is dependent on the degree to which the global regulative idea overrides the regulative of each document: that is, the extent to which the collection as a whole is made to signify; the extent to which it is treated as a collection, rather than a set of ‘collecteds’ that happen to collocated.¹²³

Such a distinction begins to develop in the process of transmission. Evidently, each document, being discrete, invokes its own transmission path. What matters is the relative independence of each path: whether there is some kind of global, pre-emptively defined model applied to each document, or each path recursively adapts its model to the document at hand. The latter is collecteds-oriented, the former collection-oriented. Dahlström draws the same conclusion in his comparison of mass digitisation and what he called ‘critical

¹²³ It is the difference between copies of *Titus Andronicus*, *Romeo and Juliet*, and *King Lear* being put in a box (de-emphasising the collection) and being bound in a volume with “The Works of Shakespeare” printed on the front.

digitisation': the former tends towards homogeneity, the latter recognises more fully the heterogeneous nature of each object. (Dahlström 2012, 463) (I will return to these approaches to digitisation in the next chapter.)

We find the same distinction arising once the transmission reaches the user. To what extent does the user treat the collection as a singular whole, or a set of discrete entities? To what extent does the form of the output enable one or the other? Or even, by nature of the form itself, imply one over the other? These questions form the subject of the next section.

3.5.4. Hybrids and subdivisions

In the above, I have given an account of three types of structure that are created by various grounds for selecting and accumulating documents. These should not be regarded as absolute or fixed. A rhizomatic system, as Deleuze and Guattari observe, can, at various points, contain tree structures; or a tree begin to partially sprawl into a rhizome. (Deleuze & Guattari 1987, 14–5) We can easily see how such things might happen. A parallel delimitation is a plurality of sequentially-delimited documents. A disjunctive delimitation might create a rhizomatic structure that, at various points, structures components hierarchically: containing sequential and parallel delimitations as components. See, for instance, the *Walt Whitman Archive* (Folsom & Price, eds.) which contains multiple versions of *Leaves of Grass* as well as a plethora of other documents. Documents may, within an overarching structure, have numerous — and far from universal — relations with each other.

These multiple structures may, moreover, arise from any point in the transmission process, though the point-of-selection is perhaps the most pertinent. It has already been argued that disjunctive delimitations produce a domain from which to select, but we must question whether that domain itself comprises internal structures. (Indeed, the pragmatics of selection

— ‘it depends where you look’ — have already been considered to some extent.) The point is that, having delimited a domain, we should not fall into the trap of considering it simply as a ‘pool for selection’ where all documents are equal. This is particularly the case where a domain transcends previously-established boundaries: for instance, the manuscripts of a given author, or even better (for being less discrete) papers on a given topic, that are held in multiple repositories. In such circumstances, while the ‘selection pool’ might be established conceptually, there are two pragmatic steps to actual selection: the selection of a repository, and then the selection of documents from that repository. What the former suggests is an opting-into the selection policy of the repository itself. In such cases, therefore, the structure finally produced can be seen as having two principle tiers, hierarchically arranged: documents, linked rhizomatically to a concept, but sub-divided into ‘sub-collections’ determined by their source: trees-within-rhizomes. Of course, the effects of this depend on individual circumstances.

3.6. Points of synthesis and merging

The previous sections looked at the process of selection and the various organisational forms (Deleuzian ‘images of thought’) instantiated. Depending on the regime created by selection, a greater or lesser degree of freedom is opened up in later stages of the transmission process: for instance, whether processes are strictly sequential, or may open parallel paths of process. It also has an effect on whether paths are merged, and the nature of the outcome of this operation.

Let us take the simplest example, that of a linear delimitation (several documents that, in a linear order, comprise a single entity). A linear delimitation is, as argued above, synthesised from the outset. As such, the transmission process is predominately sequential: there is only one thing, one assemblage, to operate on. The output is also, by the same token, singular: a

signifying entity is delimited and input, and a signifying entity emerges. By deciding to edit *Beowulf*, selecting the documents that comprise *Beowulf*, and actually editing *Beowulf*, what is produced is a representation of *Beowulf*. This does not, of course, mean that there can be no temporary division in practice — an editor can photograph each page as a distinct entity (forking the transmission path) and transcribe each page separately — but for the notion of *Beowulf* to hold, these must once more be merged back into a single entity.

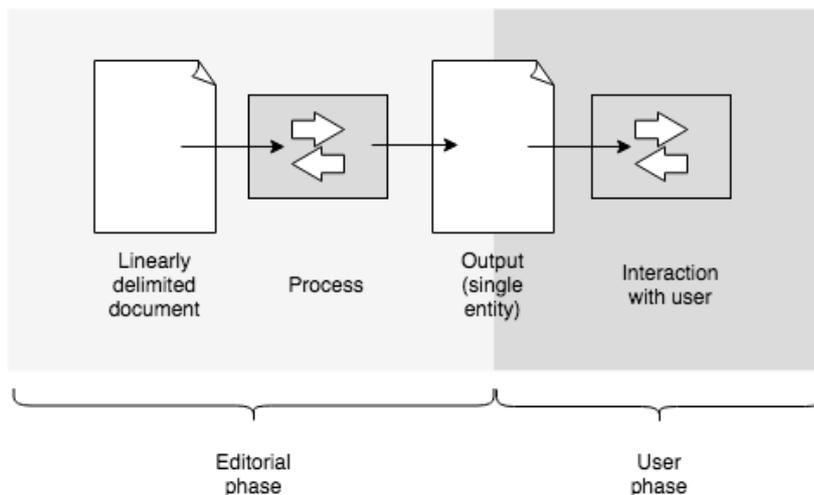


Figure 34: Sequential delimitation and linear transmission produces a single entity; transmitted to the user as a single entity.

A parallel delimitation, from the outset, selects a plurality — multiple versions — and hence invokes multiple paths. At some point, these may or may not be merged (a critical edition merges; a variorum edition does not). This is even more true of a disjunctive delimitation, as the higher-level unity is only ever conceptual. The question then arises, at what point is a merging possible? In what circumstances is it even necessary? To begin with, we may identify several potential points of merger.

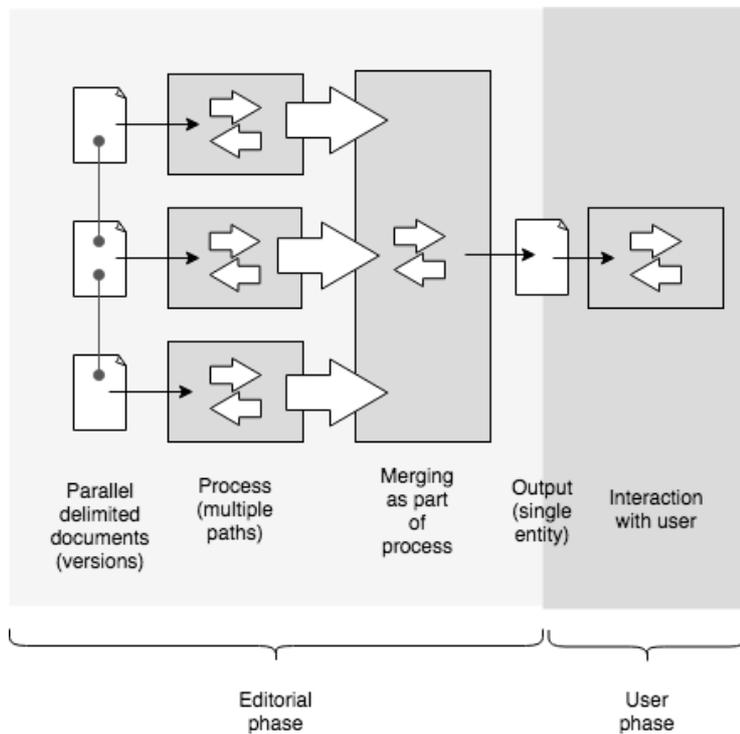


Figure 35: merging as part of the transmission process

The first, as described above, is merging at the point of selection (see Figure 34). Secondly, a merging function may operate at any intermediate point in the transmission process: for instance, an editor collates multiple versions, and then produces a critical text. (See Figure 35, above.)

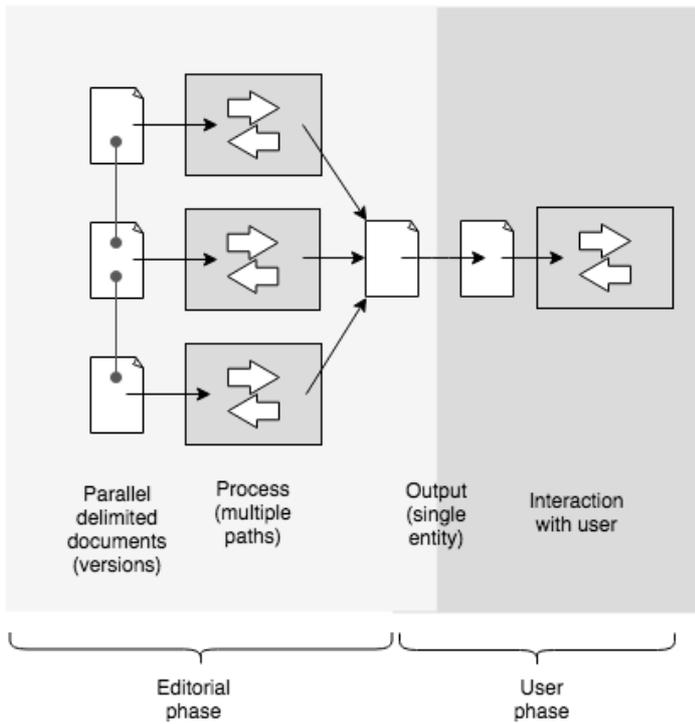


Figure 36: Merging at point of publication

The third is a merging at the point of output (as in Figure 36). This is the case, for example, when multiple distinct entities are processed separately, but merged by being joined in a single form: a book, or a website. A book, by its nature, is a single entity, and must be published as such. It thus requires a merging of paths. What is less clear is whether that merging necessarily overrides the extensity of its component parts: whether it is really a synthesis or simply a kind of concatenation. I will return to this question shortly.

There are, however, further possibilities, of which we find the most ready examples in the digital medium, notably the web.¹²⁴ Due to the non-fixed nature of the web, there is no requirement for the output to be a single, merged entity. In terms of process, transmission

¹²⁴ CD-ROMs, by contrast, obey the same logic as books in this regard. They are discrete, fixed entities.

paths that are carried out diachronically may continue on into the ‘user’ phase. (See Figure 37, below.) In simple terms, various component parts may be published additively, as and when the editorial phase of that particular path is complete. Google Books is a prime example of this: digitised books are added piecemeal once they are digitised. Thus, for the user, there is not a ‘complete’ holistic entity (*The Google Books*), only a kind of rubric — a domain name, a common system and interface — under which the component books are collocated.

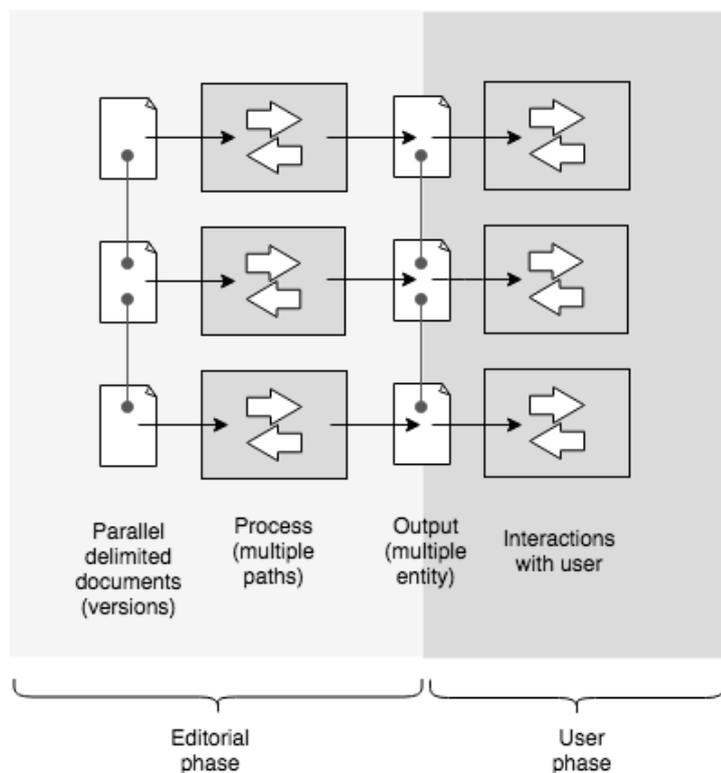


Figure 37: Conceptual unity of parallel delimited documents, transmitted independently (no merging)

In the absence of a merging function, a particular kind of regime is created in the hands of the user. Each component is a discrete entity, and may be treated as such. Provided that it *is*, in the hands of the user, treated as such, there is no question arises of what the ‘whole’ might signify. In terms of validation thresholds, therefore, it is perfectly possible in this case to have a mixed regime. Each component entity may be reflexively validated (that is, completely represented on its own terms), while the whole is merely an articulated transmission (it articulates automatically, the presence of an entity confirming the transmission of that

entity).¹²⁵ Against this, we must weigh other factors: chiefly, the degree of homogenisation carried out in the transmission process. Homogenisation, properly considered, is a kind of conceptual merging — it makes the component parts subject to a whole, even if that ‘whole’ is not available to the user.

More significantly, we must consider a fifth potential point of merging: a merging invoked by the user. (See Figure 38, below.) This is the result of the user considering the whole as a signifying entity, and not just an accumulation of components. That is, taking what is there at a given stage of incompleteness — a rhizomatic system dependent on the contingencies of the whole operation — and elevating it to a substantive. This is again most visible in the digital medium, where myriad tools facilitate the use of the collection as a holistic entity: topic modelling, word frequency analysis — any kind of what has, after Moretti, been termed *distant reading* (Moretti 2013). Even full-text search operates on the whole collection as a collection.

¹²⁵ It is as if one has an accumulation of mini-scholarly-editions in a single entity, while the single entity itself is not a scholarly edition.

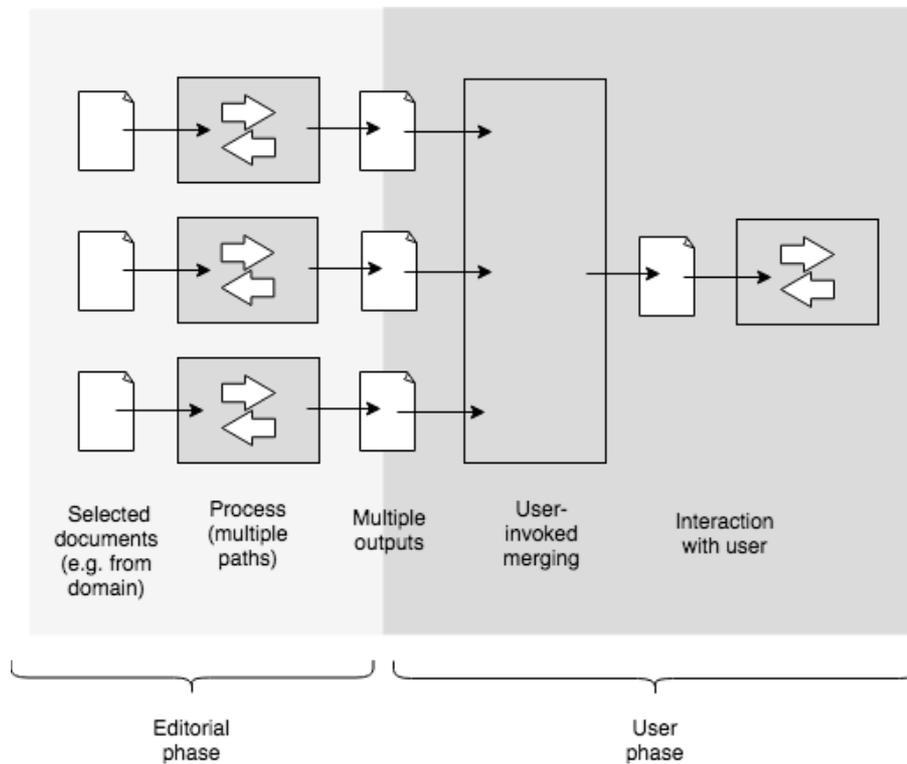


Figure 38: user-invoked merger

Of course, quite what a user can do (and, indeed, whether that was foreseen by the editor) depends on the form of the output. This will be discussed in the final section.

3.7. The output, spatiality and movement

In the above section, I suggested that the form of the output in part determined the extent to which the component parts of a transmission were conceived of as an extensive unity, or, alternatively, as simply a ‘container’ for independent components. Evidently, if by selection or earlier processes, a unity is already created, this is a moot point: there is already a unity. However, if in the case of a parallel or disjunctive delimitation, in which there is no merging in the editorial phase, the form of the output — governing *how* the user can interact with it — is significant. In this regard, we can ask a number of questions. To what extent are the component parts isolated from each other? Does a user interact with them as discrete entities, or as part of a larger whole? And, in the latter case, how does this larger whole present itself?

Is a user able to move from entity to entity in some kind of logical order (like turning the pages of a book)? Or is access to each entity through some hierarchical structure such as an index page? And, as noted above, in the digital realm, are there tools that provide holistic overviews? As Deegan and Sutherland argue, referring to McLuhan's famous dictum:

The representation structures of any and all technologies . . . have implications for the formulation of knowledge and the reformulation of our critical engagements: that means of storage and reproduction are related; that the medium is, after all, the message. (Deegan and Sutherland, 2009, 5)

The answers to these (and undoubtedly more) questions gives an idea of the kinds of configurations that might occur. In all cases, however, it is not simply a case of the form of the output, but the interaction between the form (and its potentials) and the entities encapsulated therein. Thus to ascribe absolute properties to books or the Web or any other form is, I think, overly reductive. Moreover, it is not my intention here to provide a wide conceptual analysis of various media types: most of the distinctions between print and digital, with particular reference to the latter, have been well-rehearsed elsewhere (including in the Literature Review of this thesis). However, and only insofar as the medium can be isolated from the content, we may attempt a tentative analysis.

In this, Deleuze and Guattari's concept of *smooth* and *striated space* are useful. These concepts might be seen as continuations of ideas previously established, most notably the differences between rhizomatic and tree structures, molecular intensities and molar extensities, the abstract and the discrete, the continuous and the stratified. Treating *space* as a plane, they attempt to analyse how various structures enclose, divide and fragment this space, or conversely open it to new movements and flows. In general, striated space is organ-ised (separated into 'regions' representing things), while smooth space is characterised as the *Body*

Without Organs.¹²⁶ In smooth space, the *point* or fixed location is subordinate to the *line* or movement; in striated space, the *point* dominates, subordinating movement (it is movement from one point to another, which constrains movement).

In their illustration of types of space, Deleuze and Guattari begin with *fabric*. Woven fabric is an example of striated space. Horizontal and vertical elements are interleaved, but always according to a predetermined pattern, with one of the dimensions being fixed: the resulting fabric may be of any length, but its width is fixed by the size of the loom. (Deleuze and Guattari 1987, 525) Its homogeneity of form is, as Adkins observes, what makes it striated: homogeneity does not lead to a smooth space, because the very notion of homogeneity presupposes a rigid arrangement of molar entities. (Adkins 2015) Smooth space is, by contrast, heterogeneous and comprising molecular entities. Deleuze and Guattari give the example of felt as a smooth space: a web of intersecting fibres with no preordained model or structure, and no fixed width or height. Indeed, woven fabric and felt may be the two closest examples to purely striated or smooth space. Further examples of textiles illustrate various degrees of mixed space. Embroidery or crochet produces a smoother space than woven fabric: it is extensible in all directions, and heterogeneous at a local level (the colour thread or type of stitch used); at the same time, it is striated by having a centre, and being subject to an overarching pattern (Deleuze and Guattari 1987, 525). Patchwork, by contrast, exhibits localised striations (the boundaries between individual pieces of fabric), but at the global scale is considered a smoother space: there is no overarching model, rather the possibility of infinitely extension in all directions, using any kind of material. (Deleuze and Guattari 1987, 526)

¹²⁶ A notably weird phrase borrowed by Deleuze and Guattari from Artaud. The 'Body without Organs' is intensive space, in which flows are free to circulate without constraint.

This kind of understanding is useful, as it allows us to consider not only the output in isolation — and not only one *form* of the output in the case that there are multiple — but the effects of some of the processes insofar as they produce a discernible effect. It is relatively easy, for instance, to map different kinds of transmissions onto these fabric models. A set of versions of a work, for instance, may be recognised as starting from an overarching pattern (the ‘work’ as a selective model). At the same time, each version may be subject to some global model (i.e. normalised) or treated as (relatively) unique. A normalised model is more striated, closer to woven fabric. In contrast, a very open-ended disjunctive delimitation is a smoother space — akin to patchwork — allowing expansion in many directions, and not subjecting the components to an overarching pattern (patches of a patchwork quilt are accumulated into a whole, not ‘synthesised’).

To begin, it must be recognised that no textual space is smooth. It is, firstly, subject to some overarching model — even if that is de-emphasised to a degree producing a relatively smooth space. Secondly, texts predominately have subdivisions, whether physically (pages in a book, for instance) or logically instantiated (discrete web pages). What is important to this discussion are the possibilities of movement enabled by a given medium. Relations between entities — the component parts, and the whole — might well be conceptualised by an editor, but what matters from the user’s perspective is how they are instantiated in the output. *Using* the output of a transmission is a process, and thus how the user moves from one interaction — with a given portion of the output — to another evidently shapes the nature of the whole. An increasingly smooth space enables more heterogeneous paths of movement, versus a relatively striated space that closes down such avenues.

It is important to recognise that the possibilities of various kinds of movement are always created, either intentionally by the editor or as a *de facto* result of employing a given medium.

A space that is, in the first instance, striated along all kinds of subdivisions, is not returned to some pre-striated state by possibilities of movement that render the space more smooth; rather, ‘re-smoothing’ is the further imposition of a different model, a different organisation. Deleuze and Guattari’s ‘maritime model’ of space suggests how smooth space may become striated, and the inverse (though they note these are not symmetrical operations).

They describe the sea as “a smooth space par excellence” (Deleuze and Guattari 1987, 529): a boat, for instance, is free to move in any direction on the surface. Deleuze and Guattari go on to describe the gradual striation of the sea: by the advent of shipping lanes and territorial waters, that, while operating in a virtual dimension, nevertheless overcode the physical stratum of the sea, and constrict movement. The same is true of ports, which striate the sea into trade routes (the line or trajectory subordinated to the point: *from* one port *to* another port). At the same time, they also suggest how the sea may be once more smoothed. A submarine, for example, may overcome the constraints of certain striations (such as shipping lanes and territorial waters) — thus ‘re-smoothing’ the space. Though, as Adkins notes, this kind of re-smoothing operation may be seen not as a return to the pre-striated smooth space (Adkins 2015, 235); instead, it is “for the purpose of controlling striated space more completely” (Deleuze and Guattari 1987, 530): the effacement of striations asserts control over the sea as a whole. The same understanding can, I think, be applied to transmission outputs. The effect of allowing heterogeneous movements between entities is to smoothen the space; but with the additional effect of rendering the whole more totalised, as will be seen in the coming discussion.

Returning to discuss transmission outputs, we may begin by distinguishing two sides to a ‘movement possibility’. The first is the availability of an enabling mechanism; and the second,

a degree of facility — how easy it is, in other words, to actually carry this out. These two combine to produce a continuum. At one end, we might place physical juxtaposition of two component entities (either on the same page or facing pages of a book, or on a single web page): this provides a mechanism to move from one to the other, and makes it very easy to do so. At the other end, a complete separation with no possibility of movement. The latter is an absurdly absolute position, and almost impossible to realise (without actively hiding information — which is rather contrary to the purpose of textual transmission). In between, we may envisage any number of possible modes of movement, and the relative ease with which they might be carried out.

Let us begin with printed books. In both the space they delimit and the movement possible within it, they emphasise the unity of the whole over the discreteness of component parts; they are relatively extensive assemblages. As a space, it is both restricted to a single dimension, and has clear boundaries at either end: “[T]here is a cover, clearly indicating its beginning and its end.” (Sutherland & Pierazzo 2012) A sense of the scale of the whole is afforded. Moreover, though other kinds of movement are clearly possible, the ‘default’ mode of movement is a linear one. The content is structured in a linear fashion, in a “predetermined sequence of [usually numbered] pages”. (Sutherland & Pierazzo 2012) This makes the book, as Gabler notes, the ideal form for reading (Gabler 2010, 43) — that is, the linear consumption of linear text.

Other kinds of movement, though the mechanisms exist, are subordinated by the comparative difficulty of carrying them out. Compared to turning the page and continuing to read, using one finger to keep the current page open while consulting an index and then

flicking to another page is evidently more difficult.¹²⁷ Where there is a linear structure, there is a greater sense of a discrete whole (as argued in the section on selection, above, being able to place documents in order suggests an overarching structure — that which orders in the first place). This is an impression, I would argue, that is intrinsic to a book; even when the contents are in no particular order, linearity appears as a kind of unwritten law.¹²⁸ Kline, for one, is eloquent in arguing that editors of print editions (particularly of correspondence) should think very carefully about how they order the constituent letters, foreseeing the potential of an assumed-linearity to mislead the user. (Kline 1995, Chapter 2).

There is, then, in a book an assumption — however subtle — that it encapsulates some kind of discrete whole. What, then, does this mean for a book containing multiple, accumulated entities — as in, for example, a print scholarly edition of a historical figure’s correspondence? Are they independent entities accumulated, or components subsumed into a totalising whole? This depends: a comprehensive index, for instance, allows more direct access to a given entity, and tips the balance slightly towards the *collecteds* — because it is easier to access discrete entities as such. Into this equation we must also add the tendency of a printed book to require a great deal of normalisation — a greater deal than digital, at least; without producing costly page facsimiles, there is a minimum requirement to render a text in a form that can be typeset. Whether there is a strict normalisation (of, say, spelling) depends much more on the transmission process. The answer to the above question, therefore, is somewhere in between on the *collection*–versus–*collecteds* scale. We might think of it therefore as a mixture of smooth and striated space.

¹²⁷ We might compare this to a scroll, which makes ‘random access’ orders of magnitude more difficult. If this thesis were printed as a scroll, the instruction to view this footnote would entail scrolling right to the end of the document!

¹²⁸ A book, I think, imparts a feeling of intentionality; its covers suggest completeness, and the ordering of its pages suggests linear content. (Reference books and cookery books are, perhaps, two counterexamples to this latter point.)

The digital medium (by which, almost *de facto*, we mean the Web) is more difficult to assess, not least because it is capable of a far greater number of configurations. In general — leaving aside more recent, highly dynamic pages — the Web has two default mechanisms for movement. One, as discussed above, is actual juxtaposition: putting multiple things on the same page. The potential in this regard is much higher than a book. As McGann notes, facing pages editions (the juxtaposition of two versions of a text) are limited to two versions. (McGann 1983, 32) Within the obvious limitations of screen size and usability aside, a digital environment has no such constraints. Witness, for instance, the Versioning Machine (Schreibman 2002, 2010, 2016) or the TEI Critical Apparatus Toolkit (Burghart *et al.* 2016). The other movement mechanism on the Web is the hyperlink, which enables arbitrary connections between pages. This creates, at least potentially, a much smoother space, as movement is possible in virtually any direction. But, most significantly, the default (in the case of no links) is no movement, unlike a book. Linear, book-like connections (with a forward and back button, for instance) can be used where it makes sense (i.e. there *is* a linearity), but not otherwise.¹²⁹ As such, the discreteness of component parts — not as subsumed into a totalising unity — may be more strictly maintained.

The most rigid of structures is a hierarchy, with access to component parts strictly via an index page. It clearly suggests some kind of unity (an index page provides a strong notion of the extensity of the collection), but only to a degree. It might be considered more like a box

¹²⁹ It is worthwhile pointing out that a previously determined linear system (i.e. determined by earlier processes) can be *undone* by deliberately not invoking some mode of linear movement. If a linear text, divided into pages, were presented without a ‘next page’ button, the linearity of the text would be undone to a degree. An output might even be wilfully misleading, for instance by having a ‘next’ button that skipped three pages at a time. Not, of course, that anyone in their right mind would produce such a thing; but it serves to illustrate the point that the output does not *de facto* replicate the entity’s regime — it must always be re-enacted.

in an archive: as a container, but one which is largely non-signifying. Thus the space is a kind of patchwork: smooth with regards to its heterogeneity, but with firmly fixed local striations.

To this must be added the fact that the boundaries of a website are not only considerably more plastic (they can be expanded by the addition of new material) but also less defined. For one thing, though most websites do have a starting page of some kind, direct access to a component entity is possible (via its URL, or a search engine).¹³⁰ The starting page acts as a kind of gateway, giving the sense that you as user are ‘entering’ something discrete; but if that gateway can be evaded, this sense is lost. As such, the user interacts starting from the point of entry, again downplaying the sense of the collection as a whole.

We might compare this to a more interconnected kind of structure: a less hierarchical, more rhizomatic, network-like model. Let us take the example of a digital edition of letters. These may, preliminarily, be structured hierarchically (an index page, with a link to each letter on its own page). The editor might add links to the next letter in a correspondence; or add in-line links from named entities, as they occur in a text, to another page listing all letters mentioning that named entity. This opens more avenues for movement, creating a smoother space, lessening the boundaries between entities and, to a degree, merging them.¹³¹ Most significantly, these connections are not haphazardly created: they signify; they have significance in the overarching conception of the collection. This is an instance of re-smoothing space that Deleuze and Guattari present in the maritime model (see above). It

¹³⁰ Granted, a reference (such as a citation) might point directly to a *page* of a book, and the user can jump straight to the page — but doing so requires at least picking up the book, and thus appreciating its extent.

¹³¹ As observed earlier, this is a spectrum. Entities are not *as merged* as being placed on a single page, but are more so than if they had to be accessed by moving up and then back down a hierarchy.

is not smoothening as a return to a relatively non-signifying space, but smoothening that, to a degree, reinforces the space as a signifying extensity.

This tendency is seen even more with the potential of digital tools to aggregate and perform analysis across a collection. We can see a tool such as Google Ngrams (which will be discussed in the next chapter) operating essentially to break down the boundaries between individual ‘Google books’; the effect is to transform the *collected* books into a *collection* of books. In fact, I would suggest, any tool that takes the entire collection as an object produces such an effect, to a lesser or greater degree — again, the specifics matter. I will, therefore, return to specific examples in later chapters.

3.8. Conclusion

This chapter has had two objectives. The first was to give a fuller account of textual transmission as a process, and in such a way that a variety of different approaches can be directly compared. The purpose of such an approach is not to be simply reductive, but to avoid the drawing of arbitrary distinctions. As argued above, it is not as if human transcription and OCR do different things in terms of the result: they both ‘transcribe’ text. The difference lies in the complexity and adaptability to the material at hand that a process suggests. A significant assertion of this chapter is that complexity (and particularly adaptability) are the result of *repetition*, of recursively applied models. Human interpretation involves, almost by default, repetition: interpretation and validation in turn, repeatedly. This contrasts with more mechanical, linear operations, which compound the effects of previous operations (whether they are ‘mistakes’ or not; a mistake being, after all, only a mistake *against* some validation). What matters is, where the interpretative model applied comes from, and how it is applied. Using this conception, we can study specific transmission

processes, and the affordances they provide for repetition and validation, and therefore give an account of the validity of the output.

The second section of this chapter presented selection as an interpretative act, as a process: it is neither magic, self-evident nor pre-emptive; even when selection has previously occurred, it must have happened as the result of a process that has real effects. The following section outlined various regimes that are produced as a result of applying a particular selective function. The aim of this analysis is not the creation of ontological categories, or binary divisions. Rather, it is to suggest that where such categories might appear to exist, they too are rooted in specific action. The ‘categories’ described are significant as a result of their effects on later stages of the transmission and output. Indeed, linear and parallel delimitations might well be seen as special cases of a disjunctive delimitation (‘the continuity of text’ or ‘identity of text’ could be considered as external, ‘third’ entities to which we can connect documents). They are distinguished here because they create different presumptions of unity that affect later stages of the transmission process.

In the next chapter, the concepts and approach developed here will be put to work on various kinds of (scholarly) textual transmission: different types of scholarly edition, digital collections and archives, mass digitisation, and critical digitisation. The aim is to suggest that these are distinguished from each other not (necessarily) by scholarship, accuracy — and certainly not any claims to objectivity — but by the specifics of their process and the types of regime thereby created.

4. (Digital) Scholarly Editions and Digitisation

4.1. Introduction

Thus far in this thesis, all manner of scholarly editions and textual digitisation have been encapsulated under the generic rubric of textual transmission. In this chapter, I will look more specifically at a range of transmission types.

In doing so, we shall make use of the concepts articulated in the previous two chapters. The foremost conception is that of transmission as a process, with the constitution of a particular process giving rise to particular results. Among the possible parameters for any transmission are a number of *validity thresholds* (naïve, articulated, and reflexive) which describe degrees of representation; an approach to modelling transmission stages (linear, recursive; operating sequentially or in parallel); three types of preliminary delimitation, linked to three Deleuzian schema (or images of thought); and a mode for considering user interaction with the output. These concepts will now be put to use.

It should be noted that none of these analyses are intended to provide in-depth accounts, but rather to illustrate the role of process in creating different systems and regimes of representation.

4.2. Scholarly Editions

Scholarly editions, for all their undoubted complexity, are perhaps the easiest things to map onto this abstract picture of the transmission process. To see this, we must understand the purpose of a scholarly edition. If it can be very briefly summarised, it is as such: to transmit source material such that the output is functionally equivalent to the source. This ‘functional

equivalence' evidently mandates that scholarly editions are governed by a *regulative idea*. As discussed in Chapter 2, the regulative idea is *that which is represented*, be it a document or work. As also argued (and as Eggert makes very clear), a regulative idea is not a universal, transcendent or objective notion. It is established by the editor through interaction with the source material. It is not simply asserted out of nothing. Rather, it is a validated assertion: an experimentally validated assertion.

The regulative idea marks the limit of a textual assemblage. It is established simultaneously in two directions: from the outside, where it is posited as a limit; and from the inside, where the assemblage is constructed outwards towards that limit and sustains it as such.¹³² The result is what Deleuze and Guattari refer to as a *signifying regime*. All the component parts of the assemblage refer back to the regulative idea, sustaining it and, simultaneously, being subject to it. This is necessary to the functioning of a scholarly edition as it constrains interpretation, preventing it being arbitrary, and allows completion within the terms laid out by the regulative idea. Scholarly editions thus create *tree* or *radicular* systems; the root node is the regulative idea.

Production of such a system requires a *reflexive transmission* — that is, a transmission that simultaneously determines its components and validates them. As such, the process of producing a scholarly edition can be modelled, from the outside, as a single recursive function. This does not mean, of course, that this function cannot be decomposed into constituent functions; but these constituent functions are arranged in a mutually-recursive, rather than linear, manner. In other words, at any stage in the transmission process, there is

¹³² We may think of it, perhaps, as the keystone of an arch. From one side, it holds the arch together and determines its structure. But the keystone does not float in mid-air awaiting the construction of the rest of the arch: the arch must be built, from the ground up, to support the keystone.

the possibility — in the event of a contradiction arising — of returning to the beginning and experimenting once more.

This stands in contrast to other kinds of textual digitisation, which precede in a more linear fashion and thus have no mechanism (or, at least, not a global mechanism) for reflexive experimentation. It also means that analysing the precise processes behind the creation of a scholarly edition is more difficult: tracing a path through myriad functions that bounce off each other repeatedly is evidently more complex than following a neatly laid-out linear path; at the same time, as it stands to reason that we must be able to encapsulate the entire process in a single recursive function, such precise analysis is not required to account for the output.

I shall begin this discussion by looking at editions of works, probably the most (until recently) prominent form of scholarly edition. These will then be contrasted with documentary editions of various shades and complexities.

4.2.1. Editions of works

Understanding scholarly editions of works (by which, typically, we mean literary works, though there is no reason to consider this a particularly defining restriction) involves, firstly, conceptualising what is meant by a work. Indeed, as the argument presented here will suggest, a great many of the debates seen in the history of textual scholarship arise from scholars taking as a starting point various — if not radically different — notions of what a work is. For Shillingsburg, the work is “the message or experience implied by the authoritative versions of literary writing.” (Shillingsburg 1996, 173) Such a view tallies with a general idea of *authorial intent* — of which, it should be stated, there are myriad shades and positions. The Greg-Bowers-Tanselle school, formerly the mainstay of Anglo-American editing, takes a very strident position, arguing that the work should be recognised as the *final authorial intent* (i.e.

the author's last word on the matter). Implicit in this perspective is the fact that a work is singular. Other versions of the work that might exist are non-authorial, or the author's previous intent (subsequently modified).

Such a view was challenged by the likes of McGann and McKenzie, who stressed that a work is always produced by any number of social forces, of which the author was only one. (McKenzie 1985; McGann 1983) Likewise, other schools of editing (the German school being most prominent) held a particular suspicion of the kind of eclectic editing that the theory of final-authorial intent mandated, particularly on the basis that it is ahistorical. (See Eggert 2009, Chapter 8 for a fuller discussion of the German editorial position.) These perspectives have, undoubtedly, led to a more deliberately bipartisan view, as exemplified by Shillingsburg (writing some three decades after the more Anglo-American-centric view quoted above):

[The work] first, [is] a category into which we place all texts that appear to be versions of the same artistic unit, including all editions and printings regardless of accuracy or authority. And, second, "work" is conceptually that which is implied by the authoritative texts. The second definition leaves open the question what is meant by authoritative, but each archivist or editor has to articulate that meaning in order to limit the range of documents to be collected or represented. (Shillingsburg 2013, 13)

In such a view — particularly the requirement for the editor to 'articulate' what they mean by a work — are the seeds of an alternative conception. The work is neither a particular extant version of the text, nor the author's intent *per se*, but both simultaneously — and, it should be noted, any other self-consistently valid interpretation. Or, put better, the work is determined by the limits 'articulated' by the editor. This gives us a way out of the impasse

that Foucault presents: “A theory of the work does not exist, and the empirical task of those who naively undertake the editing of works often suffers in the absence of such a theory.” (Foucault 1969, 104) The key word in this is *empirical*: that an absolute, universal notion of the work that would allow its definitive establishment is impossible. There is nothing definitive about an edited work; it is only valid within the parameters set out by the editor. This tallies with Gabler’s view that an edited work is always the editor’s. (Gabler 2007; 2012)

As argued in Chapter 2, from a Deleuzian perspective, this view is preferable to positing the work as an Ideal, as separating the real and ideal in this way impedes analysis (Where, we must ask, *is* the ideal? How do we get at it?¹³³) A Deleuzian ontology has no such distinction between real and ideal; everything is real, being is “univocal” (Deleuze & Guattari 1987). Rather, the divide is between the *virtual* and the *actual* (which are both real). The virtual encompasses the myriad potentialities of an assemblage — its *affects*, what it could produce in interaction with something — and the actual, that which is actualised by a specific set of circumstances, i.e. the history of its interactions.

From this perspective, Gabler’s argument that, “it is documents that we have, and documents only” (Gabler 2007, 199) is true in one sense, particularly if read as, *documents are all that we (as editors) can interact with*. The work, then, is a *virtual* component of documents: a potential. And there are as many ways of instantiating that work — of rendering it *actual* — as there are modes of interpreting the documents. The work actualised is a particular interpretation of documents. In this respect, then, Gabler’s argument is open to question. Without something external to the system (i.e. an editor, with an “editorial gaze”: Eggert

¹³³ Moreover, what does this say about authorship? In my view, it reduces a complex, creative act to simply instantiating a pre-existing ‘perfect’ entity (and, almost by definition, not doing a very good job of it).

2009, 214), there is not a work; or, rather, the work exists, but is virtual (a possibility; or, better, a probability of being interpreted in a certain way by a given interpreter). This is, indeed, the same conclusion arrived at by Eggert (from the perspective of Peircean semiotics): the sign of the work being not just (as in the case of a bipartite semiotics, such as that of Saussure) a signifier and signified, but necessarily involving a third term, the interpretant. (Eggert 2009, 234)

In suggesting that there are as many ways of actualising the work as there are modes of interpretation, it is important to recognise that not all interpretations are self-consistent; nor are many interpretations mutually compatible. Constraining interpretation is therefore required. What this means, as argued in Chapter 2, is experimentally deriving a coherent and cohesive rationale. At the very base of this is the assumption that what is being interpreted *is* a work. In doing so, it should be noted, a little authorial intent goes a long way; suggesting it was written by someone (as Foucault's 'author-function' suggests) implies a modicum of intention, which fundamentally restricts later interpretation.¹³⁴

But asserting a work does more than this. It suggests that the documentary evidence should be interpreted as a set of — typically — linguistic signs. I am not arguing here that bibliographic codes (as McKenzie says) should be ignored as a source of meaning. Rather, that the notion of a literary work implies reproducibility; or, in Deleuzian terms, a high

¹³⁴ In this I am inclined to agree with Robinson over Gabler. Seeing a work in a document involves asserting from the outset a degree of intentionality: specifically, an intent to communicate with written language. This closes down other potential interpretations that would otherwise be valid and self-consistent, e.g. "*Ulysses* is a fancy doorstep". Of course, *Ulysses* could be *actualised* as a doorstep — it is a heavy book! — but it becomes inconsistent as soon as a little authorial intent is allowed. (Why would someone spend ten years producing such a fancy doorstep? Why the all the writing if it's just to hold open a door?) Surely — Occam's razor to the fore — if Joyce intended 'doorstop' as an interpretation, he'd have saved saved lots of effort by selling bricks instead of writing a long novel.

propensity of signs towards deterritorialisation and reterritorialisation (Adkins 2015, 50–3). Language exhibits just this tendency (which is why it can overcode the physical stratum): words may be ‘lifted off’ (decoded and deterritorialised) of a page and recoded and reterritorialised on another page (or in thought), in a way that, say, the texture of the page may not.¹³⁵

This may seem like attaching an *a priori* definition to a work (one that transcends interpretation). However, it is also possible to approach this from another angle: that of process. As seen in the previous chapter, a work can be considered, at the point of the selection of documents, a *parallel delimitation*: that is, it selects documents (or fixed linear sequences of documents) on the basis of identity. But this is really to place the outcome before the process. Instead, we might take one document and use it as the basis for a selective function — a function to select other documents, whose application to other documents either interprets an identity or does not. And interpreting identity is really to assert a functional equivalence: the same signs may be extracted from both; which points, inevitably, to extractable, and hence prominently linguistic, signs.¹³⁶

¹³⁵ This raises some problems in editing for works that have some intrinsic, non-linguistic (and especially visual) form. Guillaume Apollinaire’s ‘Horse’ calligramme (a poem written to visually represent the shape of a horse) is cited by Pierazzo (2014) as one such example. The text itself may be deterritorialised, but the shape of the horse not so easily. Likewise, Neuber argues that interpreting the poetry of Stefan Georg requires a consideration of the colours and letter-forms used. (Neuber 2016) The problem here is that consideration of colour and letter-form mean that editing can no longer reduce graphemes to purely alphabetical abstract signs. Either we posit a wider range of discrete signs (red-‘A’, being different from blue-‘A’) — producing a critical distinction where there is only a *partial* distinction — or we reduce them to identical ‘A’s. The problem, essentially, is that these letterform-colour combinations are really rhizomatic: a heterogeneous assemblage of colour and abstract sign, that cannot be separated without overly strong distinction or excessive normalisation. In this case, a *pluralistic* approach is fitting (see below).

¹³⁶ What, it may be asked, are we to make of a work that exists only in a single version? Clearly an identity between documents cannot be extracted. In such cases, I would suggest, the *presumption* that other versions could exist (even as a possibility) is enough to suggest that a work is linguistic in nature. We might consider, for instance, a single extant version of a work, where all the other versions have been lost.

From the above, then, it is possible to see the relative lack of difficulty in conceptualising the transmission process behind the production of an edition of a work. Firstly, in terms of selection, there is a natural restriction in one dimension (the length of the work), and a constrained ability to accumulate versions in another dimension. The starting basis is, therefore, a highly restrictive proto-regulative idea. How an editor chooses to proceed from then on is a matter of applying an editorial rationale and validating it — in other words, further developing the regulative idea. Let us take as an example the production of a critical edition of a work according to final authorial intent.

The first step is, evidently, transcription of each document. This is a recursive process (see the previous chapter) that invokes a new transmission path for each version. Collating versions can be regarded as a process that accumulates and partially merges these paths. Determining final authorial intention is, likewise, a process of stratification: it selects (or elevates) those portions that are judged to be final-authorial and discards (or relegates to a critical apparatus) other versions. Where such stratification cannot produce a result one way or the other, the editor falls back to a base-text or ‘copy-text’ (the default version, in Greg’s terminology; see Greg 1950). The resulting text is then typeset or digitally encoded (if it were not already) and published, along with textual notes that, in essence, explain the procedures of its creation. The user then takes up the edition and, interpreting it according to the editor’s intent,¹³⁷ experiences the work.

¹³⁷ As noted in Chapter 2, this cannot be guaranteed. The user might employ the edition as a doorstep. While the edition has moved beyond the editor’s control, both convention (it is a book, not a doorstep!) and explicit textual notes act as a kind of contract with the user: *read this*, bearing in mind *these* interpretations, and the result is (functionally) the work.

What is important about each of these stages is that they are necessarily encompassed in a recursive whole: from any point, it is possible to reach back, adjust the components and restrictions of the regulative idea. This includes selection, and, significantly, the form of the output: whether the possible limitations of the output are taken into account in earlier stages, or the editor tries something and finds they have no way to represent it and returns to the earlier stage, the result is the same (it is experience that short-circuits the try/fail loop). Experimentation, repetition and validation are what produce reflexively-valid transmissions: a regulative idea in a reflexively-validated transmission can never be pre-emptive, never simply asserted. The regulative idea is really, therefore, the process of transmission — the interpretative models applied at each stage — condensed and elevated to the status of a signifier (against which, reciprocally, the process is itself validated as consistent and coherent).¹³⁸

This means, most significantly, that no particular interpretative process has any universal, empirical validity; only a claim to be self-consistent. We see this most readily in the controversy surrounding Gabler's edition of *Ulysses* (1984). In Eggert's view, much of the criticism of this edition stemmed from misunderstandings of Gabler's editorial methodology, in particular its divergence from the Greg-Bowers 'copy-text' conception. Gabler's edition, based on "the final stage of authorial development" (Eggert 2009, 167), involved an entirely different conception. It is a self-contained, self-coherent regime, producing a regulative idea, which in turn stratifies the documentary sources in a different way (giving preference to different versions) than a 'copy-text' theory would. But this does not invalidate such an approach. Neither Gabler's nor the Anglo-American approaches have a claim to empirically

¹³⁸ We might see the regulative idea as the point at which the produced output wraps up all the threads of its production and 'cuts loose', becoming a free-standing entity.

determine the work *Ulysses*; only a claim on each side to produce a valid regulative idea, to be internally consistent.

Accordingly, it should be observed that no editorial school or approach to editing can be simply applied as a set of transformative processes (i.e. linearly, non-recursively). A school is, as argued in the Literature Review of this thesis, a set of codified practices that act as a shortcut to experimentation in a given instance. This does not, as we have seen, make them universally valid; an approach must always be validated with regard to the work it is employed to edit. To do otherwise is to not produce a reflexively-valid transmission.¹³⁹

In the above examples, we have dealt primarily with editorial processes that produce a single, critical text, under a regulative idea that produces a *tree* structure. As suggested in the previous chapter, it is not essential that a parallel delimitation be collapsed down to a single entity. Each path may be kept relatively isolated, producing what is typically considered a variorum edition: an edition “which displays all textual variants of a work on the reading page.” (Bryant 2002, 20-21) This produces something closer to a *radicular* structure: several, independent versions — each a discrete unity — but brought together under the auspices of a ‘higher’ root or regulative idea. In this case, the regulative idea is constructed on a pluralistic notion of the work. It is still, however, an edition of the work. There must, after all, be some rational grounds for accumulating the versions. Such an edition may, indeed, also include a critical text of the work. The digital ‘Hyperstack Edition’ of *St Patrick’s Confessio* (2011), for

¹³⁹ Indeed, it is only by attempting to use a given approach on something new, and finding it does not fit, that any development of editorial theory might arise. Greg’s ‘copy-text’ theory was, for instance, an adaptation of classical and medieval editing to the very different kind of textual transmission of the works of Shakespeare, in particular its distinguishing between *substantives* and *accidentals*. (Gabler 2012, 25)

instance, presents multiple manuscript versions, as well as print-based critical texts of the work.

The ability to produce such editions is, of course, also dependent on the medium of the output. The digital medium is invariably more suited to the task. (In the Literature Review, it was argued that this kind of edition was rare in print — and single-text critical editions more common — because of the restrictions of the medium.) The digital medium also allows (at least more easily) a pluralistic approach to the same version, most typically by including a transcribed text and a facsimile. This provides numerous advantages, not least, to paraphrase Gabler (2011), to read (as text) what can only be seen in the facsimile, but also to see what cannot be easily de/reterritorialised and de/recoded in some formal way.

4.2.2. Editions of documents

Documentary editions have a long history, even if (previously at least) seen as a lesser form of editing. (See Kline 1985; Pierazzo 2011; 2013) The rise in prominence in documentary editions has been attributed to the advent of the digital medium (Pierazzo 2013), in particular its representative capacity (full, high-resolution page facsimiles) and scope for plurality (facsimiles alongside transcribed texts). A documentary approach has been most obviously seen in the editing of historical documents, where the uniqueness and historical contingency of the material document is deemed significant, and not simply a ‘vessel’ for a linguistic text. That said, with editions of literary works, such as the *Jane Austen’s Fiction Manuscripts* (Sutherland ed. 2010), a document- (as opposed to work-) orientation has been adopted.

Analysing the process behind documentary editions is more difficult than editions of works, for many reasons. Firstly, as we have seen, a literary work presumes (in general) a particular form and level of closure. It is assumed to be primarily (if not totally) linguistic in form, and

is thus more easily ‘lifted’ from the material support — whereas a document *is* the material and linguistic text in a heterogeneous assemblage. Secondly, a work points to a natural closure, which is far from the case with documents. Granted, a single document might have fixed dimensions, but even then it is potentially problematic: does each side of the page count as a separate document? To what degree must two pages be physically connected before they count as a single entity? And so forth. (This discussion has already been carried out in the previous chapter, so I shall refrain from rehashing it here.) The point is that documents are able to form a range of connections to other documents, based on any number of interpreted aspects. To analyse this in terms of structure, and the kind of regulative idea necessary to determine and validate that structure, requires the consideration of a number of possibilities. We must, therefore, consider more closely the establishment of a regulative idea, and what it really represents.

In general, documentary editions can be split into two broad types based on the regime created by their material delimitation (see previous chapter): these may be *sequential* (multiple documents that reduce to a single entity: a tree) or *disjunctive* (a collection of individually discrete documents: a radicle or rhizome).¹⁴⁰ These regimes, it is evident, will produce very different regulative ideas. A single entity may take the limits of that entity as the preliminary basis for a regulative idea (though it is not without problems). A set of multiple accumulated entities must do more to construct a coherent set of bounds. In either case, there is much less conceptual coherence to begin with, and hence a greater focus on the pragmatic application of process.

¹⁴⁰ Typically, documentary editions do not deal with *parallel* delimitations, not least because editing on the basis of uniqueness and historical contingency somewhat butts heads with asserting identity between documents. Kline does give some examples, particularly letters produced in duplicate or triplicate. (Kline 1998, 97) In this case, she advocates many specific criteria for choosing one version, or, ultimately, some form of critical emendation.

Let us begin with editions of *sequentially delimited* documents. In fact, it is better to begin with the most trivial case: a single document — that is, one page with clearly-defined boundaries determined by its physical extent. Selection in such circumstances is therefore verging on the self-evident. ‘The document’ (as proto-regulative idea) is thus coterminous with its physicality: without tearing it in half, it must be selected on that basis (even if the ostensible rationale for its selection — as, indeed, would typically be the case — were its textual content). Because the document is selected as a physical entity, the act of selection does not prioritise any particular dimension or aspect: text, marginalia, other bibliographic codes; all are equal, part of a heterogeneous assemblage — though one subordinated, as it were, to the physical document.

This stands in contrast to other, even sequentially delimited, sets — unless they are selected on the basis of physical unity (or a reconstruction of a physical unity: say, signs of broken bindings, being together in an envelope, or page numbers that would connect the pages in a given order independent of other content). As argued in the previous chapter, when a set of documents are selected on the basis of some other interpretation, most notably a continuing text, the materiality of the document and other bibliographic codes, though selected, are necessarily subordinated to the continuing text (or whatever links and orders the documents). Put another way, what is selected is, primarily, the connection between the documents, which results, secondarily, in the selection of actual documents: everything within the material bounds is determined by an immaterial connection.

As such, the initial basis for the regulative idea is not the documents themselves. Though documentary editions deliberately take an interest, and attempt to represent, other non-textual codes of the document, their inclusion is as a result of being ‘dragged in’ by the textual

connection, not as a result of their own intrinsic interest (if intrinsic interest were the point, selection could include other documents outside the connecting text). Accordingly, what is represented by such an edition is not the documents, *per se*, but the text in its documentary context. I do not think this point is particularly contentious in most cases, though in some cases tensions between physical and textual extensities arise.

One such example can be found in the digital edition of *William Godwin's Diary* (Oxford University 2010). (This example is drawn from an article I wrote for the RIDE journal; see Hadden 2015) According to the edition's editorial statement, the diary comprises thirty-two octavo notebooks, held in the Abinger Collection of the Bodleian Library. Evidently, each of the thirty-two volumes has a physical extensity, and must be selected initially as a complete entities. What connects the individual documents (aside from the rather looser connection of residing in the same place in the Bodleian) is the conceptual and textual: conceptually, they comprise Godwin's diary; and the containing text (clearly dated) orders the volumes. This raises an obvious question: what do we mean, first and foremost, by 'Godwin's diary'? It cannot be 'the thirty-two notebooks' as physical entities; rather, the thirty-two notebooks in their physical embodiment are dragged to the party by the interpretation of their textual content as Godwin's diary. If the two happen to be precisely conterminous, there is no problem. However, the edition excludes some documents, most notably the manuscript MS. Abinger e. 33 (described as a "Supplement to the journal / 1793 / Mar. 23") — presumably on the basis that it is not 'conceptually diary'. On the other hand, the so-called '1796 list' (in which Godwin lists his most common acquaintances) — not 'conceptually diary' either — *is* included. The inclusion or exclusion appears to be, therefore, on a physical basis; though this physical basis is itself derived from a conceptual unity and textual continuity.¹⁴¹

¹⁴¹ At a lower level, there are also unresolvable ambiguities. In certain instances, running out of space between the lines designating a single day's entry, Godwin continues writing into

These, in the grand scheme of things, are undoubtedly minor issues. However, it does illustrate a fundamental point: documents are inherently rhizomatic assemblages, with heterogeneous connections between text, material, and other features. Interpretation stratifies such an assemblage, necessarily ordering, under the regulative idea, the strata produced. Such stratification can be seen, in some way, as unpicking (or deprioritising) some kinds of connection, such that a hierarchical structure is created. Rather like combing tangled hair, in some places tiny knots are formed and tighten. There is no unambiguous way to produce a hierarchy that does not break apart something. Again, a solution here is to assert plurality: it is the material *and* the text, represented separately. Very often (indeed, as *Godwin* does) providing page-facsimiles is the solution: if the user is concerned about a small ambiguity resulting from formalisation, there is, in parallel, an ‘unformalised’ version available.

Thus far, we have said very little about the transmission process beyond selection. Documentary editions, in particular in their digital incarnations, involve myriad functions (selection of features, transcription, descriptive mark-up, transformations and presentation), into which it is not my intention here to delve too deeply. What is important, however, as with editions of works, is producing a signifying system (a reflexive transmission), which again means a fully recursive process. That said, there is in a documentary edition perhaps a greater requirement to focus on the needs of the user: a document is not simply reducible to

the next day’s space. The requirements of formalisation means that these entries are encoded as part of the day to which they relate, not the date whose space they happen to occupy; though with the loss of documentary features, such as line breaks, that are otherwise recorded. It is tempting to think of this as a problem of overlapping hierarchies — and thus simply the result of XML encoding; however, though overlapping hierarchies is precisely the problem, it is really a conceptual one.

its text in the same way as a work, and thus determining which features to select is a more open question.¹⁴²

I shall now consider editions of documents determined by a *disjunctive* delimitation: that is, a set of documents that are not reducible to a single entity, but that, rather, maintain themselves as distinct. When taken together (i.e. made into a single entity) a radicular or rhizomatic structure is created. In this case, I will deal with radicular editions; or, more specifically, what it takes to fashion a (signifying) radicle out of an otherwise (asignifying) rhizome.

As discussed in the previous chapter, this kind of system begins by articulating not a discrete document or sequential set of documents reducible to a single discrete entity, but by determining a *domain*: a pool from which to select. This domain is the proto-regulative idea. Equally, we might view this arrangement inside-out (that is, intensively rather than extensively). In this case, the domain is really the set of documents that are connected in some way to a third entity (say, connected by authorship to a given individual); the elevation of this ‘third entity’ to the status of a signifier makes it a regulative idea, establishing the limit of the domain. As also seen, a radicular regulative idea — indeed, all regulative ideas — accounts for all its components and creates a signifying, representative regime. This means, if we draw a map of connections (e.g. a network of relations), the *collection created* must map perfectly onto the *connecting entity*.

¹⁴² In some ways, this might be seen as the (relative) absence of convention. A critical edition of a work can be described tersely: “final-authorial intent, with Manuscript X as the copy-text” — read in conjunction with the critical apparatus — provides a good articulation of what the edition presents. Documentary editing has some terms — *diplomatic*, *ultra-diplomatic* — but these describe something of a continuum, rather than clearly distinct categories. (Pierazzo 2013, makes this argument.) There is, accordingly, a greater need to spell out to the user exactly what is presented, and how.

However, a radicular system is characterised by two levels of regulative idea. There is the *global* regulative idea, that selects documents, and the *local* regulative ideas of each documents.¹⁴³ It asserts a plurality by default. Accordingly, there may be some tension between the two, especially when the global regulative idea necessarily selects whole documents that only partially correspond to its selective regime.¹⁴⁴ This tension is also felt in the transmission process. Should, for example, the original spellings in component documents be maintained (respecting the regulative idea — the historical contingency — of the individual document), or normalised in the name of global regulative idea (thus allowing, for instance, searching for terms in the collection)? Inevitably, this tension must be resolved one way or the other (the result can be seen on a spectrum between the collection as a whole and the collected components); though granting too much latitude to components may serve to undermine the rigour of the whole.

This kind of radicular documentary edition may, in theory, take as its basis any connecting ‘third entity’ that can be elevated to the role of a regulative idea: official papers, government records, correspondence of a given individual (or between two individuals), private writings (as long as they are discrete entities, rather than, say, a draft of a single work or a diary which reduce to a single conceptual entity), and business and financial records. (Examples taken from Kline 1998.) The examples here are largely of correspondence, though the theoretical implications should be applicable to any kind of document organised in such a regime.

¹⁴³ We might consider this akin to feudal system. Whereas *tree* systems assert direct control (i.e. at the top is the king, who rules the peasants), a *radicular* system introduces a new layer: the barons. They rule their own land and control the peasants therein, while at the same time being nominally subjects of the king.

¹⁴⁴ In this regard, it is no different (except by degrees) to the sequential delimitation of documents seen above: the documents are selected on the basis of a sequential text, but this drags in all kinds other aspects as a result of the requirement to select whole documents.

Within documentary editions of this kind, Kline identifies a number of classifications that are typically employed: of these, *definitive* and *authoritative* are, she says, qualitative, while *comprehensive* and *selective* are quantitative. (Kline 1998, 70) From the perspective of this thesis, however, this qualitative *versus* quantitative distinction needs further investigation. It is true that *comprehensive* and *selection* are quantitative in a certain sense, but really only in relation to each other (having delimited some selective domain, a *comprehensive* approach would take everything, while a *selective* approach takes some subset). Instead, I would suggest, these terms are also qualitative: they point to a particular regime of representation. *Comprehensive* is a quality — rather than ‘how many’, it points rather to being complete within the bounds of the regulative idea. *Selective*, on the other hand, is more complex: it denotes (at its most arbitrary) a lack of completion, but more probably a modification of the original regulative idea.¹⁴⁵

The distinction — or, rather, degrees of distinction — between comprehensive and selective may be seen in the three categories proposed by Simon (which take correspondence as an example):

- [1] The comprehensive publication of a series of letters chosen from some larger group (such as all letters exchanged by Thomas Jefferson and John Adams); [2] comprehensive coverage of some narrowly defined topic; and, last, [3] truly selective publication, in which editorial

¹⁴⁵ *Definitive*, as Kline notes, is seldom now used by editors, who fully recognise the hubris of such claims: no edition is “all-inclusive in scope but also so rigorous in its textual methods that the reader would never have reason to consult the original material” (Kline 1998, 71) That said, this is an external perspective; internally, by definition, a radical edition must be *definitive* within its own bounds — definitive according to its own definition. *Authoritative* is, likewise, seldom used, and particularly vague: at best, it is a synonym for *scholarly*, as Kline suggests.

judgement, not some predetermined factual criterion, is responsible for every choice. (Kline 1998, 72; my numbering)

To this, I think must be added another selection type [0]: the complete letters of an individual (say, Jefferson).

We can see, through closer analysis, how these modes of selection give rise to very different regulative ideas. Case [0] clearly delimits a domain, which can be identified as (to give it a title) *The Complete Letters of Thomas Jefferson*. Considered intensively, this is a connection on the basis of authorship between documents and Thomas Jefferson. Or, more correctly, the dual connection between documents and Thomas Jefferson, and between documents and the concept 'letter' (presumably, Jefferson wrote other things). The regulative idea thus represents the union or synthesis as a single concept of 'letter' and Jefferson.¹⁴⁶ Case [1] takes the above, adds a third term (John Adams), and synthesises a new regulative idea. Such cases are relatively unproblematic.

Case [2], "comprehensive coverage of a narrowly-defined topic" is more complex. (It is, incidentally, unclear from Kline's discussion whether this is intended to be a topic in isolation — Topic X —, or a topic synthesised with authorship: the letters of Jefferson about Topic X. This should not matter; we can synthesise, or not.) This is because a topic demands a greater deal of interpretative force: its boundaries are more fluid. Thus the onus on the editor, when elevating the topic to a regulative idea, is to be very exacting. A further difficulty arises, however, when we consider that a letter may comprise more than one topic. Indeed, the

¹⁴⁶ Evidently, *accumulating* the two entities would give rise to a rather different arrangement: all letters ever, and everything written by Jefferson. That this synthesis seems obvious is, I would suggest, because we know beforehand that Jefferson wrote letters. But in delimiting a domain to function as a regulative idea, we must be sure that it does, indeed, produce a single, closed set.

selection-topic may form the smallest part of the letter. It is in this instance that we see again the tension between the global regulative idea (the topic) and local regulative ideas (individual, historically unique letters that are partially about the topic, but largely not). Complete submission to the global regulative idea would see the topic-relevant portions of each letter extracted (an act of historical-editing vandalism), while allowing absolute freedom the local regulative ideas brings into question the very nature of selection in the first place (if nearly anything goes, why are not other documents with some, though lesser, claim to topicality not selected?). A balance, of course, must be struck, typically in favour of the document. This does demonstrate, however, the extra force required by a weaker global regulative idea in order to regulate: the ‘extra force’ is applied in asserting that a letter is about the topic, when it is only partially (i.e. an overdetermination of its content).

The final case [3] given by Simon is the most problematic. It is hard, at least initially, to see how any arbitrary selection can really be made subject to a regulative idea — it is really a non-rationale. However, it may be possible, though the regulative idea is radically reconfigured. For instance, selection from an individual’s correspondence may be made on the basis of what is ‘interesting’ (i.e. relevant in a historical context) about that individual: e.g. Jefferson as President of the United States, or Jefferson as statesman; but not Jefferson who occasionally wrote a letter to his tailor. This is really positing, and synthesising, connections to a further entity — though, like a topic, a relatively unclear one. At some point, the editor must determine the extensity of such a domain (for instance, by giving a date range: Jefferson-as-statesman means 1775–1800).

In these examples, we see increasing strain being placed on a global regulative idea, and increasingly ‘overbearing’ force applied to compensate. In such cases, the user may struggle to grasp the relation between what is present and what the edition ostensibly claims to

represent. Kline (1998, 80–2) is particularly eloquent in highlighting potential pitfalls and misinterpretations that may arise as a result of selection ambiguities. This may be further complicated by external pressures, such as time, cost and the limits of the chosen medium. (Pierazzo 2011) In the second half of this chapter, we will see potential regulative ideas strained beyond breaking point, at which the whole notion of representation is brought into question.

Before continuing, it must again be observed that little has been said of the transmission process beyond selection. For the most part, this is because the ‘crunch point’ for the creation of such radicle-system editions is the point of selection. That said, in terms of process there is a considerable degree of latitude. Each document, being a self-contained extensity, can have its own transmission path — that is, being edited as a document and not as part of the collection. Questions then arise as to how isolated these paths are. Are they maintained in isolation, recursively developing their own interpretative models, only to be accumulated at the end?¹⁴⁷ Do the paths merge, or accumulate inputs to interpretative models? While an interpretative model cannot (if it is to be a reflexive transmission) be predetermined, is it determined by an analysis of all the documents (i.e. global, normalising model — see the earlier example of normalised spelling)? All are possibilities, and the precise details in any situation clearly have an effect, particularly in determining the relationships between the component parts and the whole.

¹⁴⁷ A radicular edition must be accumulated or merged before publication; if some documents are published before others (as is possible via the Web), the regulative idea from the user’s perspective is not enforced — at least until completion.

4.2.3. Genetic editions

Genetic editions form a further category of scholarly edition, which edit a work as a creative act, and via its documentary evidence. There is, indeed, little to add beyond that already discussed in terms of editing works and documents; though there are some important considerations, particularly with regard to the medium of the output and requisite stratifications in the transmission process.

A genetic edition attempts to capture the creation of a work, not just across documents but within individual documents as layers of revision. (Grésillon 1994, 188; Bryant 2002, 21–2)

The first step is, of course, to select and order the documentary evidence. While typically a sequential delimitation, this raises a degree of difficulty in that sequences may fork. The important stage of a genetic transmission is, however, the initial temporal stratification of a document surface (a division of materiality into discrete time-frames). Of course, temporal stratification also takes place within a more typical critical edition, though the time-frames created are synonymous with individual witnesses; the temporal aspect is merely an ordering of the witnesses. The temporal stratification for a genetic study, by contrast, occurs within and across documents. To a certain degree, therefore, genetic criticism could be regarded as a temporal analysis of documents; but its actual ‘units’ of selection (if that is an appropriate term) are really time-periods rather than material surfaces.

A genetic edition must also pay a great deal more attention to the form of the output. It does not produce, after all, a neat linear text, or even multiple versions thereof. In print, the difficulty has been the constraints of the medium with regard to representing chronology. While it may be possible to replicate documentary features — in a super-diplomatic edition, for instance — the immaterial temporal dimension, which is actually what the editor intends to convey, can only be represented with recourse to either language (a prose description) or

some formal metalanguage (such as the sophisticated array of punctuation marks in Gabler's 'Critical and Synoptic Edition' of *Ulysses* — not, as typically considered, a genetic edition, but a good example of the effort required to represent temporal dimensions on a page in linear text). Even then, the number of temporal strata is limited in print, and is typically restricted to assigning a temporal ordering to zones on the page.

As such, we see a possible effect of the medium on the kind of regulative idea that can be created: in the editor's mind's eye: she could perhaps picture the author scratching out a text word by word, but to represent this as a genetic text requires the aggregation of fluid (or, relatively, more fluid) time into discrete chunks. The greater enabling factor of the digital medium is in the possible delineation of these temporal chunks (in units smaller than page zones, for instance), and tools for linking them in a non-static form. After all, what a genetic edition aims to represent is the 'non-stasis' of the document; a non-static form of presentation enables the editor to represent this more fully.

4.3. Defining the scholarly edition

Thus far in this analysis, I have come perilously close to articulating what might be seen as a definition of a scholarly edition. In one sense, I think this is true. However, it is not a normative definition, saying what a scholarly edition should be, or how editors should edit, or a guide to best practice. Rather it can be seen as a kind of minimum threshold for a scholarly edition; or, even better, a dividing line that can be drawn between scholarly editions and other kinds of textual transmissions. This division is independent of any particular feature or medium or approach; of whether the editor should edit works or documents; and certainly has nothing to do with degrees of objectivity (as argued at length in Chapter 2, appeals to objectivity are invariably flawed).

Put in the briefest way possible, a scholarly edition should be a reflexive transmission. The implications of this have, I think, been well established. Firstly, only a reflexive transmission can create a signifying regime (a tree or radicle): that is, a regime in which there is a relationship of representation between the output and the source. This is a process that involves the creation of a regulative idea, establishing it by interpretation, but also validating the interpretation that it creates for consistency. This consistency demands a completeness within the bounds created. (This, indeed, is one of the criteria of a scholarly edition for reviews in the *RIDE* journal: Sahle *et al.* 2014). This also requires a fully recursive transmission process, as well as a clear articulation of those processes as an editorial statement. This makes a scholarly edition a freestanding, self-validating entity; the process of its creation is bound up with and validated by what it *is*; it can be cast loose into the world and function. It creates, and is regulated by, a regulative idea.

In the second half of this chapter, I will look at other forms of textual transmission. These are distinguished from scholarly editions precisely on the grounds suggested above. In some way, in other words, they are marked by incompleteness — or an inability to validate completeness. As such, they are unable to completely ‘condense’ the process of creation and cut free; or, put another way, an explanation of what the output *is* cannot be separated from the process of creation. Within this much broader category, however, we find very different regimes: for instance, small pockets of signifying representation within an overall non-signifying whole; or not.

4.4. Textual digitisation

In this second section, I will deal with types of textual transmission that are, typically, not considered scholarly editions. In doing so, we will see precisely the dividing line suggested above: that these forms of transmission (for various reasons) produce a transmission that

cannot be made subject to a regulative idea, and thus do not produce a signifying regime. As argued in the previous chapter, this does not mean that there is not some degree of representation; rather (and this very much depends on the approach) there can be some representation, in localised pockets, but the whole is not representative of anything beyond itself. This is most typical of digital collections, if only because the scales often involved are beyond the capacity of the book; the medium is beside the point, except insofar as certain things are possible with certain media.

4.4.1. Digital archives

That said, the first examples here may be considered as being subject to a regulative idea; at least, a reasonable argument can be made that this is the case. I have particularly in mind so-called ‘digital archives’, which should be considered as fundamentally different to ‘archival digitisation’ — the digitisation of the contents of an extant physical archive. This is a distinction illustrated by Price:

In the past, an archive has referred to a collection of material objects rather than digital surrogates. This type of archive may be described in finding aids but its materials are rarely edited and annotated as a whole. In a digital environment, archive has gradually come to mean a purposeful collection of surrogates. (Price 2009)

Projects such as the *Walt Whitman Archive* (Folsom & Price, eds.), the *William Blake Archive* and the *Rossetti Archive* (McGann) may be considered exemplars of this kind of digital archive. As Price observes, they tend to have “the care of treatment and annotation of an edition and the inclusiveness of an archive.” (Price 2009) They comprise, in essence, a heterogeneous collection of documents produced by a given author, along with extensive scholarly notes, page facsimiles and commentaries. As such, they embody a very similar

structure to radicle-structured documentary editions seen above, that is, two tiers of regulative idea: the global, in this case the author (Whitman or Blake or Rossetti), and the local, the regulative idea of an individual document. Where they differ is, of course, in the heterogeneity of the documents, which firmly shifts the relation between the local and the global in favour of the local.

This being the case, it may be asked to what extent the global regulative idea really regulates — that is, constrains selection in a self-consistent way. In an absolute sense, I would have to suggest that it does not. *Heterogeneity of type* marks a radically open dimension (compared to the more restrictive ‘the letters of’). A global regulative idea that was self-consistent would, therefore, have to select everything ever written by the subject in question. There is another, less absolute way of looking at this, however, which is to suggest that (say) Whitman-as-regulative idea ‘orders’ documents by relevance. If we consider Whitman as a poet, his poetry occupies a more ‘central’ position, while his miscellaneous letters tend towards the periphery. This ‘gradient of relevance’ means that, provided the centre is complete, completeness at the periphery is less of a concern (as it is less important in terms of Whitman as a poet). Indeed, Deleuze and Guattari describe just such a system, composed of concentric circles. (Deleuze & Guattari 1987, *On several regimes of signs*) If this argument is to hold water, however, it is essential that the heterogeneity and independence of the component parts is maintained; there can be no attempt to treat the collection-as-a-whole as a signifying entity. This includes, as described in the previous chapter, presenting the component parts in the output in such a way that maintains their discreteness; that is, maintaining a patchwork-like blend of smooth and striated space.

Gabler describes a similarly-structured kind of collection, which he (and Eggert 2005) term a 'work site'. Instead of placing an author at its centre, it instead places a literary work. Its being structured in such a way, evidently makes it for Gabler a kind of 'edition':

It is distinct [among digital research sites] in that it does not primarily reach out centrifugally from a kernel or theme or research question [...] Instead, holding its core within itself, it is centripetal. It stores the comprehensive record of the material evidence of its research object at its own digital center. In incorporating in digital transcription all extant traces and bodies of writing, text, and textual transmission of that research object, it is commensurate with the scholarly edition in the pre-digital paperbound medium. (Gabler 2015)

Again, the same caveats articulated above apply: the completeness at the centre is vital; as is the maintenance of heterogeneity in the transmission process of the component parts. The whole is not a homogeneous signifying entity.

The same argument may be applied to so-called *Thematic Research Collections*, a term used by Palmer (2004) to describe a wide range of heterogeneous materials collected on the basis of a theme:

The thematic framework allows for a coherent aggregation of content. All the materials included assist in research and study on the theme. This coherence is generally anchored by a core set of primary sources. The capabilities of networked, digital technology make it possible to bring together extensive corpuses of primary materials and to combine those with any number of related works. (Palmer 2004)

In this case, however, I think the argument holds less strength. As discussed earlier, a *theme* is an altogether more nebulous affair than, say, authorship. The ‘core set’, as Palmer terms it, is thus determined in a more selective (that is to say, interpretative) manner, and thus it does not ‘anchor’ in the same way as an author, or provide such a gradient of relevance. In terms of the structure created, therefore, it is more rhizomatic than radicular: it articulates a domain, and myriad connections within that domain, without *signifying*. In other words, the totality of the collection cannot be seen as representing — in a consistent, validated way — that theme, only an arbitrary portion of it.

It is, in relation to the theme as a signifier, the same as selecting arbitrary pages (or, even more absurdly, all the vowels) of a literary work. Evidently, they have their origin in that literary work, and thus partially represent it, but not in a way that can be validated as representative of the work as a whole. Of course, there is a fundamental difference to ‘selecting the vowels’: the selected component parts have their own internal self-consistency (under local regulative ideas), provided they are edited as such. Thus the transmission of each component represents that component, even as the collection of components is only partially representative of the theme. Or, put in terms of the validity thresholds outlined in Chapter 2, the component parts are (or may be¹⁴⁸) reflexively-validated transmissions, while the whole is only an articulated transmission.

4.4.2. Critical digitisation

At this point, we switch from types of digitisation that, in general, are the purview of textual scholars, to digitisation more typically associated with archives. One of the principal aims, if not the principal aim, of such activities is digitisation for preservation. As argued in Chapter

¹⁴⁸ Each component part invokes its own transmission path, in which processes may be connected recursively, allowing the development of a self-consistent, *local* regulative idea.

2, this in no way invalidates the notion of digitisation as *transmission*: even if preservation is the objective, and even if it locked away somewhere, it must at some point be accessed — at which point, it must transmit the source material. Another objective of such digitisation is more evidently concerned with transmission: that is, digitisation with a view to widen accessibility, typically by making the digitised forms available on the Web. As a corollary — particularly when transcriptions of sources and detailed metadata are digitised in addition to high quality page-facsimiles — the obvious divisions between such approaches and ‘scholarly editing proper’ are undoubtedly effaced to a degree.

Critical digitisation is a term used by Dahlström (Dahlström & Hansson 2008; Dahlström 2009, 2012) to denote a particular kind of digitisation within an archival setting. In particular, Dahlström uses the term to contrast digitisation procedures within an archival or museum setting to those of libraries. In Dahlström (2012) it is explicitly portrayed as a counterpoint to so-called *mass digitisation* (which will be considered later):

[I]t concentrates on what is unique and contingent in the document, whereas mass digitisation is quantitative in its design to capture what are common, regular, foreseeable traits in large amounts of documents and data. In consequence, then, critical digitisation normally has to develop project specific practices and tools, tailoring them to the qualities of the documents in the particular collection. (Dahlström 2012, 463)

In this, then, we see very clearly a process that is primarily recursive in nature. The development of ‘specific practices and tools’ points to an analysis of the source material and an adaptation of process to the specifics. It “concentrate[es] on what makes a document unique” (Dahlström 2012, 463) In this sense, it is evidently not a linear process, which applies pre-emptively established interpretative models. Being “primarily manual” (Dahlström 2012, 463), activities such as transcription can be subject to multi-level interpretation and

validation (see the previous chapter's discussion of transcription as the application of a set of mutually-recursive interpretative functions). As such, we may well regard the output as reflexively valid. Or, to be more precise, each output (each digitised document) as a reflexively validated representation of the original.

What, then, can be said of the accumulated documents, assuming they are grouped together in a digital repository, or made available through a common Web-based interface? I do not think it is going too far to ask whether critical digitisation creates a 'collection' at all. By this, I mean that its orientation is so far towards the uniqueness and contingency of each individual item that there is neither an attempt to consider what the collection-as-a-whole might signify, nor an overdetermination of the constituent 'collecteds' in the name of a higher, global logic. We can consider this in terms of the selection process. It is evidently disjunctive, taking a connection to a 'third entity' of some kind (a particular archive, for instance), without elevating that third entity to any kind of signifying status.

In other words, there is no global regulative idea; only a regulative idea developed for each component part. Considered in terms of the kind of space it creates, it is a patchwork *par excellence*: a set of heterogeneous objects, represented in all their heterogeneity, with no overall governing pattern.¹⁴⁹ Accordingly, incompleteness does not matter; or, rather, the whole notion of completeness does not apply, as no question of a 'whole' ever came into play.

I think this approach can best be regarded as a perpetuation of a kind of 'archival logic' into the digital space. This is not, of course, to suggest that archives are objective, free of any inherent bias in their selection — a point Derrida most famously makes. (Derrida 1995)

¹⁴⁹ If so minded, we might consider the outcome as a lot of little scholarly editions that happen to be in the same place.

However, the effect of incorporating any given document into an archive does not (relatively speaking) affect that document, insofar as its uniqueness as an artefact is concerned.¹⁵⁰ Evidently, a digital repository is not as a-signifying as an archive, just as it does not contain the original, unmediated document. But, given the approach suggested by critical digitisation, it is of the same order. Implicit in this being the case is, of course, an onus on the user not to treat the digital ‘collection’ as if it were a meaningful whole. The one potential hazard might be the operation of full-text search, which necessarily treats the collection, such as it is at any given point, as a totality in which to search (as seen in the previous chapter, it marks a re-smoothing of the space in the name of a supposed ‘whole’).

4.4.3. Large-scale digitisation

The final two sections of this chapter are concerned with “large-scale digitisation” (Coyle 2006) and mass digitisation. The approaches denoted by these terms are relative amorphous, and the boundaries between the two rather indistinct. (Gooding *et al.* 2013, 630). While mass digitisation can be seen as inherently based on automated processes, it is likewise the case, albeit to a lesser degree, that large-scale digitisation projects are reliant on automation. The *Australian Newspaper Digitisation Program* is a case-in-point. But other projects, such as *Transcribe Bentham*, are less automated, instead crowd-sourcing transcriptions. (Causer *et al.* 2012) However, I would contend that this is not necessarily grounds for distinction. While it is undeniably true that crowd-sourcing transcriptions — even from non-experts — produces better quality results than using OCR (in the case of handwritten documents, by several orders of magnitude), crowd-sourcing can be seen in some sense as the

¹⁵⁰ Of course, *relatively*. But, I think, there is a quite significant difference between placing documents together in an archival box (a relatively neutral practicality for storage) and binding them together in a book with a title on the cover (which more strongly suggests a signifying unity). Likewise archival cataloguing, which can be seen more as a finding aid than a particular attempt to classify or impose some ontological system — though, again, by degrees.

‘mechanisation’ of human activity. The question is not so much *who does it?* — a computer or a person — but the relationship of that activity to the overall transmission process. Transcription in a scholarly editing project is a recursive operation, geared towards the production of a cohesive whole. This means that, though pages or sections must be transcribed at different times, no page is transcribed in isolation. However, crowd-sourcing tasks are, as Rockwell suggests, small and, most significantly, autonomous. (Rockwell 2012, 147) It therefore represents the insertion of a recursive operation — a local recursive operation in an isolated transmission path dealing with a single document — into an otherwise linear work-flow. What characterises automation is, I would suggest, not freedom from human intervention,¹⁵¹ but a linear, rather than recursive, transmission process: the difference between the production line and an artisan.

There is, likewise, in both mass digitisation and large-scale digitisation an orientation towards digitising text for analysis — that is, treating the text as data. True, other uses, such as preservation and dissemination are possibilities, and indeed more likely to be a consideration in the case of large-scale digitisation, though the degree to which these two intended uses conflicts is potentially problematic. (Gooding *et al.* 2013, 636) In general, both approaches are distinguished from archival digitisation, in that preservation is not the primary impulse (even if, to a degree, it is a side-effect).

The most useful distinction between large-scale and mass digitisation is that made by Coyle, for whom the difference is not so much scale as scope. Large-scale digitisation is “more discriminating than mass-digitization projects. Although [large-scale digitisation projects]

¹⁵¹ What counts as ‘human intervention’ is, of course, difficult to determine. Does building a machine to carry out a process not count as ‘human intervention’? Or operating that machine? In which case, is there a minimum amount of manual involvement that tips the balance from ‘human’ to ‘automated’?

create a lot of scanned pages, they are concerned about the creation of collections and about reproducing complete sets of documents.” (Coyle 2006, 642) Very often, this is based “on public and academic institutions undertaking projects to digitize their own collections” (Gooding *et al.* 2013, 630) Large-scale digitisation, then, operates on a discrete set of documents (however it might be defined), and accordingly comes — in-built, as it were — with a notion of completion. This is, therefore, where I think the distinction may be most usefully drawn: large-scale digitisation produces “representative and comprehensive datasets,” which is essential to its utility (Clausner *et al.* 2013, 123). As we have seen throughout this thesis, representation is predicated on completeness.

This being the case, for any large-scale digitisation project we must ask: what is the collection? How is it defined? By whom? Where do its boundaries lie? These questions are important, as a collection is a creation; myriad historical processes underlie its assembly. If a collection seems natural or self-evident, this is only because someone has determined it as such. “Unities of discourse,” as Foucault reminds us, are discursively created. (Foucault 1972) We may therefore consider a collection in terms of how it itself was delimited, in the same way as any other unity hitherto described.

In general, we may take a collection to be the result of sequential and disjunctive delimitations, and, in most cases, a complex hybrid of the two. These hybrids can have very different characteristics, which affect how we might conceive of completeness. *Transcribe Bentham* can be seen as, at the highest level, a disjunctive delimitation: it takes Bentham-as-author and elevates him to the level of a signifier. Within this overall conceptualisation, however, we might find other delimitation regimes operating. Some documents are connected together in a sequential delimitation (for instance, the manuscript pages of a work); where there are multiple versions of a work, these in turn may have a parallel

relationship; and others, such as letters, are standalone and relatively discrete. Accordingly, individual documents within the collection have very different statuses; how they are treated in the transmission process is, therefore, important.

It also raises a question of what completeness would mean. As seen in the earlier discussion of digital archives, it points to the possibility of two (or more) tiers of completion: ‘important’ documents with respect to Bentham — in this case we might consider his major works — occupy a more central position while other standalone notes or letters might be considered more peripheral. Completeness at the centre is thus more important, and more likely to have an effect on the overall collection, than completeness at the periphery.

Other collections that have formed the basis of large-scale digitisation projects, such as the *Australian Newspaper Digitisation Program* (National Library of Australia), almost exactly invert the structure of *Bentham*. At the highest level, newspaper titles are chosen — a disjunctive delimitation, pointing to the twin entities of ‘Australia’ and ‘newspapers’.¹⁵² Within the scope of an individual title, there is a sequential delimitation: newspapers evidently have an order, determined by chronology, which sets clear bounds at each end of the sequence (as well as making very evident any gaps). Within each individual issue of a newspaper, there are then disjunctively-related articles. This allows the possibility of what we might call functional sub-completion: assuming one title is digitised, that title alone may be the subject of study, and it would not matter that others titles had not yet been digitised. This, of course, depends on how the transmission process is carried out.

¹⁵² See *Selection Criteria* for more details.

In general, therefore, we may say that, though large-scale digitisation processes are dependent in a large degree to reducing everything to the level of the individual document (linearly processed), how the internal structures of the collection — how documents are conceptually related to each other — is important. But, significantly, these structures need to be re-instantiated as a part of the process: that is, there must be a merging operation. Such a merging operation is a recursive function, which conceptualises and creates a meaningful whole (as a discrete component of the overall collection). At the same time, we may ask of any such operation the degree to which it synthesises its components (considering, and possibly correcting, interpretations made in previously isolated transmission paths) or merely accumulates them. In some cases, particularly where there is little ambiguity — printed newspapers most clearly fit this mould — the synthesis and accumulation are essentially the same; in other cases they may not be.

Another area in which there is potential for sub-divisions is the creation of a *conceptual* collection from multiple physical collections. *Bentham* again is a case-in-point. Its source material is gathered from the separate Bentham collections of University College London and the British Library.¹⁵³ Given that, prior to their conceptual merger as the object of a digitisation project, they were distinct collections, it is relevant to ask to what extent the selection policies of each of these collections differed. This is not intended as a criticism, but simply something that must be borne in mind: when ‘piggy-backing’ on previous selection policies, the effects of those policies — and the accumulation of two different, even potentially conflicting, policies — cannot be disregarded. (In the case of *Bentham*, the transcription project was preceded by years of editorial and bibliographic activity, so we can assume due consideration has been given.)

¹⁵³ Thus it achieves, conceptually, what any alumnus of UCL recognises as a profound challenge: crossing the Euston road.

Thus far in this section, we have explored the process of transmission chiefly from the perspective of selection and the kinds of delimitation created. We may now turn to the transmission process ‘proper’. It has already been suggested that much of the transmission work-flow for large-scale digitisations is linear in character. However, this begins from initial stages that have, firstly, delimited some of the structures inherent in the collection, and secondly, provided the opportunity to consider appropriate methods for later stages. Thus, though interpretative functions are applied in a linear sequence, and are, consequently, and with respect to the actual document they process, pre-determined, they may be based on some archetype. A project may look at a small subset of documents, and base interpretative models on that subset. Looking at the TEI elements provided for users on the *Bentham* transcription site, we can see that they are adapted to ‘Bentham-y’ documents — presumably based on modelling a subset of documents. This makes the interpretative models both more specific and more critical than if they had been entirely pre-determined. At the same time, the TEI elements available are not based on the specific document being marked-up, making the result (potentially at least) excessively normalising if a given document were to depart from the archetype-model.¹⁵⁴

Such an archetype-based interpretative model can also be applied to OCR. As Holley (2009) suggests, some OCR engines can be trained on a subset of the collection, substantially improving their accuracy. This training set is, again, an archetype-model. It does not fully adapt the model to a new, unseen document in the same way that human transcription can

¹⁵⁴ If — and it seems unlikely — a manuscript was found in which Bentham lapsed into verse form, the *Bentham* TEI model would either fail or be forced to ignore the poetry (sticking with paragraphs instead, for instance).

adapt,¹⁵⁵ but the model is better fitted to the unseen documents than entirely pre-established models. Moreover, an OCR engine can accumulate archetypes, using completed (and checked) documents as further training data. (Figure 13 in the previous chapter illustrates precisely this.)

Of particular importance in this kind of linear work-flow is the scope of what Causer, with reference to *Bentham*, terms a “unit of transcription”. (Causer *et al.* 2012, 122) In more general terms, we might think instead of “units of operation” — that is, the ‘unit’ for which there is a distinct transmission path. If OCR is used, this matters somewhat less: the OCR of two pages accumulated is the same as if it were a single, large page. However, as we have seen, with human transcription this is different: pages transcribed in isolation cannot draw on other pages to aid in interpretation, nor correct an interpretation of a word in the light of a more contextually evident repetition of the word on a later page. Of equal importance, and for the same reason, are the situation and effects of any recursive, validating function. In some cases, these validators are constrained locally to a transmission path (that is, they validate, say, the transcription of a single document with reference only to that document). Others have a wider remit, merging paths and synthesising a whole (which means validating against other documents within that whole). In general, however, such validation in a linear process is a distinct function (a checking stage) rather than a trigger to return and re-apply earlier functions, which points to a rather more limited scope than other such validations.

It is clear from the above discussion that, in large-scale digitisation projects, a large number of parameters in the transmission process are left open — and that these have a great effect

¹⁵⁵ Even then, human transcription is not *absolutely* adaptive. Experts are obviously more ‘adapted’ to their areas of expertise than non-experts. Expertise is, as argued in the previous chapter, a very complex archetype-model.

on the nature of the final product. This process- (over product-) orientation is, I think, inherent to large-scale digitisation. This is, perhaps, something of a chicken and egg problem: whether heterogeneous collections cannot be reduced to a tightly-bound regulative idea that would itself determine a process; or whether the size of the collection practically prevents the necessary analysis to do so. Either way, the resulting process is predominately linear, and divided into a large number of parallel transmission paths. This, interestingly, creates a structure that is almost exactly the inverse of critical digitisation. At the global level, provided the project is complete within its bounds (it digitises the whole initial collection), the result is reflexively valid — it is a representative collection. At a local level, however, the best level of validation available is to articulate the process; how each document is representative is unclear, as the processes of its digitisation are predominately determined in advance.¹⁵⁶

Accordingly, it is evident that large-scale digitisation projects favour a data-analysis-oriented approach, rather than serving as the basis for reading.¹⁵⁷ Achieving both, and at scale, is extremely difficult, as Gooding notes (Gooding *et al.* 2013, 637) — if, indeed, the two are even mutually compatible without taking a pluralistic approach. One cannot simultaneously normalise in favour of the collection, and not normalise to preserve the uniqueness of the component. The kind of linear process required inevitably brings excessive normalisation (the result of a lower specificity), thus tipping the ‘collection *versus* collecteds’ balance firmly in favour of the collection. In this, it must be said, providing a document facsimile alongside

¹⁵⁶ This is where human transcription clearly has the edge on OCR: though it can only validate one ‘unit of transcription’ against itself, it at least allows a whole unit to be internally consistent.

¹⁵⁷ Transcription data gathered from *Transcribe Bentham*, while it is made available via the UCL digital repository (as unordered text–facsimile pairings), is also being used as the basis for further volumes of scholarly editions in the *Collected Works of Jeremy Bentham*. While it may *use* the data, this is really a matter of convenience. The editions are produced by experts, and the transcribed texts corrected (recursively) against the documents or facsimiles. Thus, while *Transcribe Bentham* might produce readable texts, producing self-consistent and valid texts involves a whole new effort: to wit, scholarly editing.

a transcribed text goes some way to providing a resolution, as the user is not restricted to reading a normalised (and possibly incorrect) text; the facsimile is there for re-interpretation if required.

Presenting data for analysis is also a source of potential pitfalls. In earlier chapters, I discussed the notion of the editor extending their ‘control’ into the user-phase as a kind of contract: use the output in the intended way, and the representation is a given. Much of this ‘contract’ is, indeed, not explicit but bound up in convention: read a book starting at the beginning; read the words in order. People know how books work. Even digital scholarly editions come with such conventions: modes of exploration, clear ordering where appropriate. At the end of the previous chapter, these structures and movements were described using Deleuze and Guattari’s notions of smooth and striated space. A striated space prescribes certain kinds of movement, which means closing down the scope for interpretation. Treating document-texts as data creates a predominately smooth space; boundaries between documents are effaced. Textual data therefore has very little restriction on its use, once it is available; there are fewer, if any, prescribed conventions. Data used arbitrarily is akin to reading a book by skipping randomly from one word to another. Moreover, the smooth space of textual data is not the heterogeneity of an unmediated text, but text that has been, firstly, broken down into operation units, transferred as such, and reassembled to create a corpus. It is thus a re-smoothed space, made subject to the regulative idea of the whole; but, again, not a total re-smoothing: the structures within the collection — discrete works, individual letters, and so forth — are still latent; as are, indeed, any structures created in the transmission process. Accurate metadata, which can be seen in these terms as recording traces of latent structures, is therefore essential.

Such analysis is also dependent on a completion. This is particularly pertinent — and, indeed, is the case in many such projects — when the output is disseminated incrementally; and even more so when a project ceases without having achieved completeness according to the initially delimited set. As Kretzschmar (2009) notes, ‘completion’ and ‘finishing’ can have a range of interpretations. The principle point, however, is that treating textual data as an object of study marks its limit at that point as its boundary — whether it is complete, incomplete, finished or unfinished. It is the analysis of the data itself that creates the ‘collection’.

4.4.4. Mass digitisation

This final section is concerned with the mass digitisation of text. Much of the ground concerning mass digitisation has already been covered in the above discussion of large-scale digitisation; the two are, indeed, alike, though primarily in methodology. However, the differences are manifest.

Coyle describes mass digitisation as “conversion on an industrial scale”: the “digitisation of whole libraries without making a selection of individual materials” (Coyle 2006, 621). Dahlström also highlights such an objective — to digitise “everything” — “with no particular means of discrimination.” (2012, 462) Hahn draws the same conclusion from Google’s work with the California digital library. (Hahn 2008, 21) It is this sense of scale and lack of selection that have led to, at times, unrealistic claims. Gooding (2013) cites Battles’ (2004) claim of mass digitised collections to be reconstructing a kind of digital “Library of Alexandria”, while terming the hubris of such endeavours as “the technological sublime” (Gooding *et al.* 2013) Further claims that large corpora allow objective analysis, free from

interpretation, have also been made (see most notably, Moretti 2005; Jockers 2013), refuted in turn by Bode (2017).

In one sense, mass digitisation is determined by scale; indeed, as Dahlström notes, its value resides principally in scale over depth. (Dahlström 2012, 462) Google, for instance, has digitised upwards of 20 million volumes. (Juola & Ramsay 2017, 1) However, I would suggest that it can also be articulated as an *absence of scope*.¹⁵⁸ Large-scale digitisation, as argued in the previous section, can be distinguished from mass digitisation by its restricting itself to digitising a particular, discrete (though large) set of entities. However, the two notions, *scale* and *scope*, are really two sides of the same coin. It is not as if, at least conceptually, such notions as ‘all printed books’ does not form a discrete set. It is as much a disjunctive delimitation as any other: entities selected by their relation to conceptual-book and having-been-printed. The difference is, rather, that such a large domain is created by such an operation that it can never actually be instantiated: *a* scope exists, but it is never a scope that can hope to be realised.

There is, I think, another way of looking at this in terms of process. In this conception, the first step is not the delimitation of a domain, but a kind of preliminary indication that one might exist — more a pre-emptive assertion. Obviously, it is not wild speculation that such things as printed books exist, but there is a limited preliminary engagement with that domain; no attempt to determine its boundaries. The domain, such that it is, is therefore not a pool for selection; rather a speculative gesture towards the vastness of the ocean, whose far coastlines have not even been sighted. Selection is, therefore, not really *selection of* something insofar as it is part of a consistent whole; rather an algorithm: ‘Go forth into the world, find

¹⁵⁸ For one thing, this allows us to sidestep a kind of Sorites paradox: that of identifying a number of digitised entities beyond which ‘large’ tips over into ‘mass’.

a thing that is a book and is printed, digitise it, continue searching.’ We may conceive of selection in terms of two computer programming paradigms: *procedural* and *declarative*. The procedural is an algorithm: how to produce a result. A declarative language defines what things *are*.¹⁵⁹ The two are, of course, equivalent, but only if the procedure is complete. Mass digitisation is never complete. And not simply because its objective is large (there are lots of printed books), nor because it is far from static (new books are printed all the time), but because no limit has been posited: it must go on looking, go on digitising.

Mass digitisation may therefore be defined purely in terms of process. As such, the pragmatics of the process — what actually happens — are essential. The first stage to consider is therefore selection. The above ‘algorithm’ is, of course, far from the reality of a real operation. We cannot trawl the whole sea at once; but only in portions. The selection process for mass digitisation is highly dependent on where one looks — typically, this means partnering with libraries and other institutions. Selection is, therefore, a two-stage operation: one selects libraries, and then books. ‘Selecting libraries’ is not a trivial matter: it requires funding, willingness to collaborate, and consideration of myriad other factors (licensing, copyright, other issues social and political — it is not my intention to explore these in detail.) Moreover, all of these ‘discriminating’ elements are, it seems, at best tangential to the object of the exercise: selecting books.¹⁶⁰ There are, accordingly, undeniable selection factors at play.

¹⁵⁹ A classic example is the generation of numbers in the Fibonacci sequence. A procedural approach says: iterate over numbers, add first and second numbers, store result, add third number... A declarative approach provides a definition: the *n*th number in the Fibonacci sequence is the sum of the previous two numbers in the sequence, and we merely define the first (in this case two) numbers.

¹⁶⁰ This is, granted, a *mass digitiser’s* perspective: a one (digitiser)-to-many (institutions) relation. From the perspective of the institution whose collection is being digitised, the relation is one-to-one. However, it seems naïve to imagine that a mass digitiser is just at ‘your institution’ to digitise ‘your collection’.

The industrial scale of mass digitisation also entails industrial practices, which means an almost exclusively linear workflow; a production-line. Even selection, as we have seen, is linear. Moreover, each book (and, to a lesser degree, page) invokes its own, isolated transmission path. It is not my intention here to provide a comprehensive description of the mechanics of a ‘typical’ mass digitisation workflow, as it is practiced by any of the big players in the field (Google, most notably, but also Microsoft, and other organisations such as the Open Content Alliance and the HathiTrust¹⁶¹), but to broadly outline the stages involved. The first involves the capture of a page image, using either cameras or scanners. (Sutherland 2008, 17) The system implemented by Google uses automated machinery to rapidly turn pages and photograph each. Sutherland notes a range of possible qualities of image: it may be assumed that quality is sufficiently high for reading by both humans and OCR software (Google Books, for instance, presents books as legible page facsimiles, while searching the collection requires that the words be made machine-readable.) The mechanics of OCR software vary, and are described at length elsewhere (see Holley 2009; Tanner *et al.* 2009). Generally, they work by identifying zones of text, followed by individual printed characters, and attempting to pattern-match them against a database of letter forms. (Tanner 2009) The accuracy of OCR software is limited by a number of factors, not least the quality of the images captured and the typeface of the source material. Coyle cites the accuracy of the OCR used by Google Books at 98–99%, observing that this amounts to a considerable number of errors when applied to whole volumes. (Coyle 2006)¹⁶²

¹⁶¹ Many of these organisations work in concert with technical providers, such as the aforementioned.

¹⁶² These studies cited are not particularly recent, and thus it is likely that accuracy has improved; whether improved OCR software has been retrospectively applied to previously-digitised content is another matter.

We have previously discussed OCR, particularly insofar as it applies a necessarily pre-determined interpretative model. OCR has a relatively high degree of specificity in relation to print (an archetype-model trained specifically for print). It also exhibits high criticality (a high degree of formalisation, again within its proscribed domain), unlike a page facsimile. However, this high specificity and criticality undoubtedly come with a cost, when applied outside its rather narrow domain. In particular, it is prone to over-interpreting — imposing its own normative model. A specific illustration of this is the apparent inability of some OCR systems to adapt to the format of the inscribed text. Older books, for instance, may employ a typeface that is more difficult for a computer to recognise. Mandell provides a case-in-point: she identifies (using Google N-Grams) a sudden rise in the use of the word ‘presumption’ from around 1800; but, she notes, this is as a result not of actual rise in the word usage, but of an OCR error which encoded ‘prefumption’ (i.e. with a long-s) as ‘prefumption’ (with an f). The decline in ‘prefumption’ correlates perfectly with the rise of ‘presumption’, corresponding in fact with a change in printing practice to eliminate the long-s. (Mandell 2013) At first glance, it is natural to classify such an occurrence as an error, which, of course, it is: however, it is also indicative of a system of assumptions (or prefumptions) that is based on a pre-established, highly formalising model.¹⁶³

Such a high specificity–high criticality model also brings another possibility: that of rejection. It acts as a kind of selective model, rejecting entities that it cannot interpret. This means, for instance, that it is limited to printed books that are amenable to OCR, ruling out books with non-standard typefaces (particularly older books). This selection by interpretative failure has, indeed, already been seen earlier in the process. The utilisation of high-volume scanning machines which automatically flip pages is also, in this sense, a specific and critical

¹⁶³ Google has now ‘corrected’ the *presumption/prefumption* error; though this is not necessarily the same as getting it right the first time.

interpretative model — and its failure to interpret is an act of selection. This rules out fragile books, and, as Coyle observes, atypically-proportioned books and those with foldout pages. (Coyle 2006, 43) By the same logic, the *presumption* problem could have been entirely circumvented by excluding any book containing (or likely to contain) a long-s from being digitised in the first place. Errors can be avoided by not doing anything difficult, but this must be seen as an act of selection.

Thus, while the stated aim of mass digitisation projects may well be to digitise ‘everything’ — that is, to forgo selection and to strive for an objective totality —, as Dahlström observes, this “objectivity is no more than a chimera.” (Dahlström 2010, 92):

We need to remember that the products of mass digitisation can well be thought of as dependent on interpretation and selection — but that these are in effect ignored and silenced, which leave the user helplessly dependent on the unknown choices that praxis forced upon the mass digitising institution. (Dahlström 2010, 93)

In other words, there *is* selection, even if it is by negation (*not* the things that we cannot digitise) rather than affirmation. Moreover, it is difficult to provide an account for such non-selection. It certainly does not conform to any self-consistent rationale, only a tautology: mass digitisation digitises that which mass digitisation can digitise.

Mass digitisation’s inherent normalisation extends to the *work*. Not only does it treat all books as ‘generic books’, all editions or reprints of a single work are taken to be identical: a naïve assumption that all exemplars of a book instantiate the same work identically. As Dahlström observes:

Given that the many copies of a published book are normally thought of as identical, digitization projects [...] are usually more or less indifferent as to whether the one copy or the other is picked (rather than selected) as source document, i.e., as ‘ideal copy’. (Dahlström 2012, 460)

To say that, for instance, Google Books transmits *the book*, or *works*, is rather overstating the case. What it transmits is the *Google-Booking-of-the-book* (or the *-of-the-work*). Of course, as has elsewhere been argued in this thesis, no transmission ever transmits ‘the book’ or ‘the work’ in its totality, or objectively; the text of a scholarly edition is always the editor’s text. But it is the editor’s text created out of, and from thorough analysis of, the source material, not a pre-determined, highly normalised interpretation. It is in one sense a difference of degrees, but as suggested in Chapter Two, it is a difference that crosses a threshold: in this case, into *naivety*. Google Books’ representation of a book is naïve in its representation (in much the same way as a cheap, knock-off edition of a work).

Let us now turn to the output of a mass digitisation transmission. Here we must be careful not to assume that Google Books, which has been taken as an exemplar hitherto, is the only output. It might, therefore, be useful to distinguish three kinds of output. The first is the digitised book presented for reading (of which Google Books’ interface is indeed an example); pre-built, online tools for analysis of the underlying data (such as Google’s N-Gram search); and the underlying data made available directly to researchers. As Gooding argued (see the previous section), the exigencies of these three requirements are often at odds with each other.

Take first the ‘reading mode’. Google Books enables this through the presentation of facsimiles. This alone distinguishes it from most editions and other digital transmissions, whose primary representation of the text is as a transcription. That page facsimiles serve such a purpose — being trivially easy to read (versus, say, a manuscript) — can be seen as yet

another rationale for selecting printed books. With regards to the type of space created, and possibilities for movement therein, we may contrast it with a physical library. As Dahlström argues, “[Google Book Search] is not being modelled on the hierarchical library but on the heterarchical index of the web.” (Dahlström 2012, 463) Here I am not altogether in agreement. A library — with its catalogue and organisational systems (collections, categorisation, dates, alphabetical order) — can be seen as having hierarchical organisations. Moreover, the content of a book is subordinated to its physical container; one cannot read a book without first taking it from the shelf and reading it. However, movement between books is non-hierarchical: a visitor to a library can move from book to book along many lines (linear, random). A library is, therefore, a relative smooth space; a kind of patchwork. Google Books presents an alternative kind of hierarchy. Firstly, the content–container hierarchy of the book is inverted: one finds a book from its content (by searching for the content). Secondly, the system is, if anything, more hierarchical. To move from book to book requires returning to the ‘search’ page and entering a new search. It is not like Wikipedia, where hyperlinks create a vast network and multitudinous avenues for movement. Indeed, the dependence on this hierarchical mechanism is almost total. If one removes all the catalogues, front desk and librarians from a library, the books would still be accessible, albeit at some inconvenience; lose the ‘search’ from Google Books and digital entities are left floating in a void. It is, therefore, an even more totalising hierarchy; not just one that reduces books to a generic interpretation, but which subjects them to a single mode of access: a highly striated space.

Set against this are other modes of utilisation. Google N-Gram search can be seen, as in examples in the previous section, as a re-smoothing of space, removing the boundaries between books and creating a heterogeneous set of words. But, of course, the question is: what is this newly-smoothed space? It is a merging operation, but one invoked by the user;

what it represents is therefore the state of the collection at a given time (not only prone to change, but necessarily incomplete in its representation of books-over-time). It is true that sheer scale can, to some extent, create an approximation of representation. However, even with millions of digitised books, evident bias emerges: as Mandell (2016) observes (via a study of female authors) a sparsity of data at given points exaggerates results (or so it is presumed; there is no way to distinguish latent factors from actual trends without full access to the collection, which is really the problem). Inadequate metadata¹⁶⁴ is also a problem here, as is treating any copy of a work — not necessarily the first edition — as representative (evidently, use of a particular word in the year a book happened to be *re*printed is not what is intended).

A final use case is that of data made available as data. For instance, the agreement between the Austrian National Library and Google allows the library to keep and utilise the created data. (Kaiser 2012) Other projects such as the HathiTrust also make such data available, either in raw form or via Web-based analysis algorithms. Again, the same problem as presented by N-Gram searching presents itself: namely, the representativeness of the collection. This is an epistemological problem, and one that is central to many areas of Digital Humanities scholarship. If it can be so crudely summarised, the point of contention is this: can very large data sets be large enough to overcome the biases in their lack of representativeness? Possibly; though such a claim is normally followed by the maxim: *know your data*. But approached from the perspective of textual scholarship, this hardly feels a sufficiently rigorous basis,¹⁶⁵ even leaving aside the question of how one might ‘know’ several million books.

¹⁶⁴ Nunberg’s famous blog post (2009), describing Google Books as a “metadata train wreck”, catalogues many, and several quite absurd, failings on this score.

¹⁶⁵ I am constantly reminded of this, from Bowers: “How many conventional readings in the text of *Hamlet* — one, two, five, ten, twenty, fifty, a hundred, two hundred? — must be

In a recent article, Bode raises similar questions, again from the viewpoint of the textual scholar. She is particularly withering regarding the more hyperbolic claims of the proponents of ‘distant reading’, notably Moretti and Jockers, to be dealing with “facts [...] independent of interpretations” (Moretti 2005, 3–9):

More particularly, while claiming direct and objective access to “everything,” these authors represent and explore only a very limited proportion of the literary system, and do so in an abstract and ahistorical way. (Bode 2017, 1)

Her solution, however, is not a return to more traditional approaches, most obviously close reading, but by reimagining scholarly editions, with all their rigour, on a wider scale: “a scholarly edition of a literary system” (Bode 2017, 16) Implicit in this proposition is, I think, the very evident fact that mass digitisation does not produce a scholarly edition, and for a very simple reason: it does not represent anything. At the global level, this is evident; and at the local level, the requirements of digitising *en masse* according to pre-established, globally-applied models produces, at best, a naïve representation.¹⁶⁶

proved unsound before the ‘total values’ of the play *are* affected and the literary critic should begin to grow uneasy about the evidence on which he is formulating his hypothesis for the whole?” (Bowers 1959, 2)

¹⁶⁶ In this discussion, I have ignored the political implications of putting the determination, creation, remediation, storage and presentation of a canon of knowledge in the hands of an, if not *the*, Internet giant. Doing so has its utilities, but also poses many risks, as van Pearsen argues: “Google Book Search opens up new horizons for the accessibility of scholarly literature. But what will be the effects? Will digitised sources that can be easily accessed by a simple search command be quoted more frequently than sources that exist only in printed editions and for which someone needs to go the shelves of a library?” (van Pearsen 2010, 11) This is, however, beyond the scope of this thesis, and will undoubtedly be more competently discussed elsewhere.

4.5. Conclusion

In this chapter, we have traced a path through many kinds of scholarly textual transmission. The arc of this trajectory may be mapped on to various continua outlined in the previous chapter: from the product-oriented scholarly editions at the beginning to the firmly process-determined entities produced by mass digitisation; from a total, single representation, to various pluralities, and to only contingently representative entities; from strongly delimited conceptual delimitations to open-ended (non-)selections. (At the same time, there have been departures from this smooth trajectory: various degrees of relation and subordination between whole entities and component parts.) Accordingly, the categories presented here — insofar as they are categories — are not intended to be normative, or be clearly divided; rather, they are intended to be illustrative of the conceptual interaction of process, outcome, and configurations of some of the concepts suggested in the previous chapter.

In truth, all transmissions — scholarly editions, editorial projects, digitisation projects — are unique combinations of elements. Labels are for convenience. When Price (2009) asks, “What’s in a name?”, the implicit answer is: not much, beyond Price’s well-made argument that it is useful to give whatever it is that we do a name. At the same time, in describing something as having “the care of treatment and annotation of an edition and the inclusiveness of an archive” (Price 2009), it is evident that he recognises the ease with which hybrids arise.

The final chapters of this thesis will focus in depth on the *Letters of 1916*, a project perhaps unrivalled in its hybridisation of various methodologies: the open-ended selection of mass digitisation, the crowd-sourcing of transcriptions of Transcribe Bentham, the textual encoding of a documentary edition, and the output of a digital archive.

5. *Letters of 1916*

By its structures alone, without recourse to its content, historical discourse is essentially the product of ideology, or rather of imagination.

— Roland Barthes

Literature is an assemblage. It has nothing to do with ideology. There is no ideology and never has been.

— Deleuze and Guattari

5.1. Introduction

The principal focus of this chapter is the *Letters of 1916* project at Maynooth University, Ireland. Its aim is to analyse the process of transmission, starting from the selection of material and ending with the user output, according to the approach developed in Chapter Three and applied more generally in Chapter Four. I worked on various stages of the *Letters* project concurrently with my PhD. A much more detailed analysis of process is therefore possible. *Letters* is particularly interesting, as it combines many of the different paradigms seen in the last chapter, and hence such analysis is necessary for considering the nature of the final product.

It should be noted that the project received additional funding in 2017, allowing it to expand its scope beyond 1916, the year of the Easter Rising, to cover the Irish ‘revolutionary period’ from 1916 to 1923. In this analysis, I am considering the project in its original scope. Moreover, as part of this expansion, much of the methodology, and many of the systems used, are being re-evaluated, with the ultimate goal of producing an end-to-end editorial platform. Again, the description here is based on the project as it was set up during my time at Maynooth.

There may be some question as to whether it is fair to subject a particular project — especially its ‘inner workings’ — to such scrutiny, especially a project whose aims are quite diverse. *Letters* is intended to be a public humanities project, and uses crowdsourcing on this basis. Moreover, it is explicitly *not* a scholarly edition; so to consider it alongside things that are scholarly editions may appear unreasonable. However, as this thesis has argued, ‘scholarly edition’ is more of a label. *Functioning as* a scholarly edition, by contrast, is the result of a particular kind of process, one that provides a self-validating outcome. The intention here is to look at *Letters* as an example of how different processes produce different results. We have already seen that scholarly projects, such as critical digitisation, produce not scholarly editions, but an entirely different outcome.

Letters of 1916 is, as suggested above, a complex project: a hybrid of methodologies. How these methodologies work with (and against) each other is significant. For instance, if something is done by a non-expert, what amount of expertise is required to validate it (if this is indeed possible); or will some residual aspect of not being expert-driven throughout remain, and have a significant effect? Crucial in this analysis is, therefore, the identification of points in the workflow in which decisions become ‘locked in’ and unalterable. Mak (2013) describes these points in her analysis of Early English Books Online (an “archaeology of a digitisation,” as she calls it, employing Foucault’s term). The movement from bibliographic lists to microfilm to digitisation represent points at which decisions are locked in, and their effects cannot be entirely expunged. The aim here is to do the same for *Letters* — though, of course, compared to EEBO it is a recent, and self-contained, project (which perhaps makes things easier).

This chapter comprises three principle sections. The first describes selection of material, from both conceptual and practical angles; the second describes the transmission process; and the third looks at the spatiality of the output and the way these reinforce, or countermand, the decisions taken in the transmission process.

5.2. Project overview

The *Letters of 1916* project was conceived as “Ireland’s first public humanities project” (Schreibman 2016). The project began in 2013 at Trinity College Dublin, under the directorship of Susan Schreibman, before moving to Maynooth University in 2014. (The digital infrastructure of the project was not fully migrated to Maynooth until 2015.) The team at Maynooth has comprised, since the project’s inception, some 30 members, including historians, digital humanities specialists and postgraduate research assistants. The first stage of the project (a web-based crowdsourced-transcription interface) was launched in 2013, followed in 2016 (to coincide with the centenary of the Easter Rising) by a searchable, digital collection of transcribed and edited letters.

The project’s principal objective is the creation of a historical correspondence from the period surrounding the Easter Rising of 1916, a seminal event in the history of Ireland, and often viewed as a pivotal moment in Ireland’s struggle for independence.¹⁶⁷ Though the Rising itself only lasted a week (beginning on Easter Monday, 24th April 1916), the project’s collection of correspondence is more wide-ranging, spanning six months either side of the event (from 1st November 1915 to 31st October 1916). The subject matter of the collected letters is

¹⁶⁷ The scale of scholarly literature on the Easter Rising is enormous. Townshend (2006) McGarry (2010) provide good general references, and were the basis of my reading on this area. (It is, I think, fair to say that no historians seriously doubt the significance of the Rising in Irish history.)

deliberately open: the collection aims to contextualise the Rising in terms of life in Ireland at the time, and against the backdrop of the First World War. According to Schreibman:

[By a]llowing letters from personal collections to be read alongside official letters and letters contributed by institutions [the collection] will add new perspectives to the events of the period and allow us understand what it was like to live an ordinary life through what were extraordinary times. (Letters of 1916 press release, Schreibman 2013)

Accordingly, though the Easter Rising can be regarded the basis for the collection's rationale, the event itself is more of a central 'pivot' around which heterogeneous material is collected and organised. In other respects, the Rising is more significant: the commemorations surrounding the centenary of the Rising in 2016 led to much renewed public interest, upon which *Letters* as a "public humanities project" aimed to capitalise.

From the beginning, public engagement and participation in the creation of the digital collection were as important as the digital collection itself. This engagement primarily takes two forms: members of the public are able to contribute letters in their possession to the collection; and the project adopts a crowd-sourcing approach to transcription. The project is also notable for organising a large number of public outreach events, lectures, transcription workshops, and classes.

Much has been written about the application of crowd-sourcing methodologies to various text-digitisation projects, including their efficacy and accuracy. Ridge (2013) notes the role of crowdsourcing in public engagement and in the completion of time-consuming work. Causer *et al.* discuss many aspects of crowdsourcing, based primarily on their experience with the *Transcribe Bentham*, with particular reference to the costs of carrying out such a project (Causer *et al.* 2012), and the additional work required to verify and correct crowdsourced

transcriptions (Causser *et al.* 2018). Siemens *et al.*'s 'social edition' of the Devonshire Manuscript may be seen as an attempt to involve 'the crowd' in a more broad editorial role, over and above the relatively mechanistic tasks of transcription and encoding. (Siemens *et al.* 2012)

It is not the intention of this chapter to further explore the efficacy of crowdsourcing in such terms as those considered by Causser. As already suggested, crowdsourcing formed an intrinsic part of the *Letters* project for very different reasons — most notably, engagement, in which respect it has undoubtedly been a success. For the purposes of this chapter, then, crowdsourcing is a distinct set of functions in an overall transmission process. The question, therefore, is not whether it is 'better' according to any particular metric than some hypothetical alternative; rather, what effect does crowdsourcing — and not crowdsourcing in general, but specifically how it was carried out in this case — have in shaping the end product? To what extent does it determine, or even constrain, later functions in the transmission process? What aspects of the textual assemblage passed from function to function does it 'lock in'? In answering these questions, of course, the 'crowdsourcing function' must be considered in the context of those functions that preceded it, and those that follow.

5.3. Selecting source material

5.3.1. Conceptual delimitation

The first step in describing and analysing a textual transmission process is to account for its input. As Chapter 3 argues, selection of material is an interpretative act. It was also argued that selection, considered itself as a process, operates by transforming some document or entity into a selective model, whose selections accumulate as a part of this model and form

the basis for further selection. We identified three ‘base’ forms: *sequential* delimitation (when a text, for instance, continues across multiple documents, forming the documents into a logical sequence); *parallel* delimitation (when one document or sequential set selects another on the basis of a textual identity — e.g. versions of a work); and *disjunctive* delimitation (where documents are different, but selected on the basis of one or more ‘connecting entity’). It was also argued that selection operates on a continuum between a conceptual orientation and a pragmatic one: a strong concept operates to limit the pragmatics of selection, while a weaker concept allows more latitude for the pragmatic.

Letters of 1916 is particularly interesting in these regards. At a global level, it is evidently disjunctive. It selects individual letters that are not connectable on the basis of a text or other signifying system. It is similar in this regard to documentary editions (of more than one single document or document-sequence) and all kinds of digitisation that operate on discrete entities. At a local level, it is equally evident that letters are sequences of documents (or single documents: sequences of length one). *Letters* is perhaps unique — certainly atypical — in that it also selects copies of some letters (especially circulated copies of the Rising leaders’ ‘last letters’), which comprise parallel delimitations. This is interesting, in that it gives rise to the potential for these letters to be considered almost as works; moreover, in a purely numerical sense, it increases the number of letters-as-objects, without increasing the number of letters-as-unique-texts.¹⁶⁸

The *Letters* selection is unique in several other respects. Firstly, it most clearly resembles the kind of selection carried out by mass digitisation projects (see previous chapter). Though, at

¹⁶⁸ This makes, to give a simple example, naïvely carrying out topic models prone to certain pitfalls: once reduced to bags-of-words, some letters are essentially counted twice (notwithstanding the minor differences between the versions).

face value, its selection criteria appear to delimit some kind of domain (a ‘pool for selection’), it can be seen rather as articulating a kind of algorithm for selection. This is, in part, due to scale; but also dispersion: it is impossible to analyse where the limits of selection might lie. (In the previous chapter, I described mass digitisation as ‘gesticulating towards an ocean’ from which to select; the same is true here.) Such a situation is markedly different to large-scale digitisation projects, whose selection is determined by the boundaries of an existing collection or collections (e.g. *Transcribe Bentham*’s boundaries are the Bentham collections of UCL and the British Library). Nobody knows, in other words, how many ‘letters of 1916’ there are;¹⁶⁹ there is no real notion of completeness, or even of what completeness would really mean in this context.

The situation is, however, arguably even more complex. Considered as a disjunctive delimitation, and articulated as an algorithm for selection, the *Letters* collection criteria could be seen as comprising necessary connections to three conceptual entities: ‘conceptual-letters’, a time-period (November 1915 to October 1916), and a connection to ‘Ireland’. The first two, despite some potential difficulties, are relatively straightforward. The latter, as we might indeed expect of an entity that comprises a geographical area, a political construct, a social construct, and a national identity (amongst other things), is more complex. To make any headway, let us first start by adopting Schreibman’s more precise version of the selection policy: “epistolary documents,” “[written between] 1 November 1915 and 31 October 1916,” “with a tangible connection to Ireland.” (Schreibman *et al.* 2017, 1–2)¹⁷⁰

¹⁶⁹ Novick (1999) gives some sense of the scales involved: “In 1914-1915, the last fiscal year during which records of letters posted were kept, approximately 192 million letters were mailed within Ireland, which works out at roughly forty-four letters per person.” Even assuming a similar number for 1915–16, it is difficult to know how many of those letters have survived; moreover, as we shall see, *Letters* does not restrict itself to letters “mailed within Ireland”.

¹⁷⁰ The earlier “letters about Ireland” was taken from the public-facing project site, and thus (quite reasonably) steered clear of an overly academic language.

“Epistolary documents” comprises, evidently, a more open category than ‘letters’ — it includes postcards, telegrams, and even notes passed by hand. (For this reason alone, determining scale from the characteristic of having passed through the postal service is potentially inaccurate.) Of more importance than what is classed as a letter, however, is the rationale behind how any given letter is delimited: in other words, where its boundaries lie. Above, I suggested that individual letters can typically be considered sequential delimitations (pages connected and ordered by their text). However, things are not quite this simple. Does, for instance, an envelope count as part of the letter? What about another document included with the letter? Or — even more nefariously — one letter forwarded in an envelope with another letter? Or two letters on the same page? The potential complexities are endless. The problem is not a failure to adequately define ‘letter’ (or “epistolary document”). Rather, it is an inevitable consequence of selecting a heterogeneous assemblage and giving it boundaries, i.e. hierarchising the strata. (Indeed, in the previous chapter, we saw the same problem with *Godwin’s Diary*: is the diary the notebooks, or the text?) Decisions have to be made, and they are always, to a degree, arbitrary.¹⁷¹ The question is, what effect does this necessary normalisation have on later processes, and the output? Some of these points will be raised again in later discussion.

¹⁷¹ In part, such decisions are determined by the requirements of later processes and the output. A more nuanced and adaptable (i.e. recursively adapting) approach to the materiality of the letters could assess each case individually. However, this creates its own problems: particularly, how to we ‘normalise’ letters enough to present them via a relatively generic interface, and how to make the text amenable to full-text searches. The decision taken by *Letters* can, I think, be seen as a sensible compromise: to relegate ‘additional material’ to a secondary status, including the facsimiles so that a user can see them, but not transcribing the text or presenting it as a part of the letter. See, for instance, O’Rahilly: “Letter from The O’Rahilly to Hely’s Limited, 1915/1916”. (Schreibman, *ed.* 2016)

Now let us turn to the second ‘connected entity’, the date-range. There are, of course, the difficulties one would expect of accurately dating undated letters. In general, historical expertise is the guide here (elevating the project above mass digitisation approaches, which simply sees no date at all). On a technical level, the TEI encoding uses the *notBefore* and *notAfter* attributes to designate a range; however, this raises difficulties for searching and categorisation by date (as we shall see in later sections). A more fundamental question at this stage is, what is signified by the date range for selection? In general, I think it is reasonable to argue that a smaller date range is inevitably more signifying than a larger; it entails a more restrictive conception (assuming the date range was not chosen entirely at random). Moreover, it is impossible to consider a date-range in isolation when there are other factors at play: 1916 in Ireland is quite a different prospect to, say, 1916 in the United States. It is the year of the Easter Rising, whether the subject matter of any selected letter is explicitly concerned with it or not. Furthermore, by extending the date-range six months either side of the Rising, as opposed to the calendar year 1916, its importance as a ‘pivotal’ moment is further enhanced. In opposition to this is both the arbitrariness of the cut-off points (a very pertinent letter with respect to the Rising might have been written a day after the cut-off) and the heterogeneity of the material. In this way, the selection criteria points to the creation of a collection that is ‘about’ the Rising, considered as an “inclusive and complicated narrative” (Schreibman *et al.* 2017, 2), as well as simultaneously being about anything and everything else — a mixed regime of signs, as Deleuze and Guattari would say.

The final selection parameter, a “tangible connection to Ireland”, is more complicated. Indeed, properly considered I do not think it can be reduced to a single parameter. To my mind, there are at least three interpretations, which do not necessarily overlap, or overlap in complex ways. Moreover, they have the potential to select very different kinds of letter, and of varying degrees of relation to the other parameters. And, perhaps most significantly, any

one of them is sufficient for selection. One possibility is the criterion ‘about Ireland’ — that is, Ireland as the subject of the letter. This, however, only need be a partial reference, and even then takes in myriad possibilities: reference to a place in Ireland, reference to someone in Ireland, or someone Irish, or some ‘Irish topic’ (however that might be defined). In general, such a parameter is likely to overlap with the second: being sent *from* Ireland, either internally or externally. In this case, the letter is almost *de facto* ‘about Ireland’, regardless of its actual subject. By contrast, letters sent *to* Ireland from abroad are almost the opposite: they are not necessarily about Ireland at all (letters sent home from the First World War fit most obviously into this category). However, the ‘about Ireland’ category could also include, for instance, letters sent within the British government *about* Ireland (most obviously the political situation in and around the Rising). This raises a curious possibility: the collected letters that most overtly ‘discuss’ Ireland are the ones *not* written there. They have to be more overt in subject-matter, not being otherwise assured of selection via the more evidently ‘tangible’ categories.

These parameters, considered together, delimit a highly complex space — indeed, an order of magnitude more complex than the ‘all printed books’ of Google Books presented in the previous chapter. It comprises not only a heterogeneous collection of material, but a heterogeneous set of *spaces* from which to select, the limit of any one of which is difficult, if not impossible, to determine. This, I think, not only confirms Dahlström’s argument that inclusive approaches tend to merely occlude their selections (Dahlström 2012), but that the more open-ended a selection policy, the greater the potential for surreptitious selection criteria to have an effect. Deleuze and Guattari’s account of re-smoothing space is relevant here (see the final section of Chapter 3). A discrete collection is necessarily restrictive, but in closing off a space it makes it smoother: a patchwork of documents, in which all are, in a sense, equalised. The openness of an undetermined space, by contrast, encompasses

numerous striations, with not all documents accordingly ‘equal’. These effects are, as we shall see, further heightened by the actual, real-world operation of selection.

5.3.2. Pragmatics of selection

This is where (and it is not really a pun) things get real. As argued in Chapter Three, in the absence of a tightly bound conceptual delimitation, much is left to the pragmatics of how selection is actually carried out. In this regard, the parallels with mass digitisation are also clear: it *really* depends where you look. As seen with Google Books, the project does not, in the first instance, ‘select books’; rather it partners with institutions (i.e. selecting them). So too with *Letters of 1916*, whose collection is most obviously shaped by the institutions (predominately archival collection) that are ‘selected’ as territory in which to search for letters. Before continuing it is worth re-iterating a point made earlier: that selecting from archives is not the only method by which *Letters* gathers its material. Letters are also contributed by members of the public.¹⁷² This mode of selection is difficult to assess, its most important determining factors lying to a much greater degree outside the control of the project, and being therefore less amenable to being described in terms of a ‘rationale’; rather, we are looking at a good dose of serendipity. Briefly, then, this mode of selection is predicated on numerous factors: a member of the public must have some letters (which means the letters having been preserved — almost an accident in itself); the member of the public must have heard of the *Letters* project and be willing to contribute their letters; they must have the means and ability to photograph and upload those letters, or to bring them to an outreach

¹⁷² It would be a worthwhile endeavour (though one beyond the scope of this thesis) to determine whether these crowd-contributed letters differ in any significant way from letters taken from archives. The relative small number of such letters presents a challenge in this regard. I would guess, also, that determining such differences — that is, extracting differences in some qualitative or quantitative way — would prove to be difficult as crowd-contributed letters are likely to be even more diverse in subject-matter than a similarly-sized set gathered from a particular archive. Put another way, they are likely to illustrate the heterogeneity of letters from a disparate set of archives in microcosm.

event. All of this stands quite apart from any aspect of the letter itself, in terms of content, authorship or historical significance. We could, of course, give an account of any particular selection: make a rhizomatic map of the myriad entities and actors that brought about that selection.

The same is evidently also true of letters selected by team members visiting archives. Selection is ‘achieved’ thanks to myriad, contingent, and above all different, historical processes. However, an archive in this respect acts as a kind of stop-gap: much of this work has already been done and, significantly, rationalised to some degree by the archive’s approach and policies. None of this makes an archive a neutral space, nor an objective one. As Derrida, most notably, observes, archives represent both the physical collocation of documentary evidence (the physical space), and a site of power (the *nomological* sense). (Derrida 1995, 9) “Consigning” material to an archive is thus not a neutral act:

By consignation, we do not only mean, in the ordinary sense of the word, the act of assigning residence or of entrusting so as to put into reserve (to consign, to deposit), in a place and on a substrate, but here the act of consigning through gathering together signs. (Derrida 1995, 10)

Likewise, in the terms of Deleuze and Guattari, it may be said that an archive — although open and ‘non-selective’ in a conventional sense — at a given moment represents a closed, bounded space, and thus signifies to some degree (though we do not necessarily know what). That said, it represents an altogether more mixed space — or mixed regime of signs; it is not a purely signifying system that overcodes or determines its component parts. (This is, in part, due to its storing the physical documents, rather than mediating them by, for instance, digitising; such processes, as we have seen, necessarily involve the application of an interpretative model, which overcodes to some degree.) Thus an archive can be considered

heterogeneous — a patchwork assemblage;¹⁷³ though, as Derrida recognises, it is never absolute heterogeneity. (Derrida 1995, 10) Most significantly, therefore, we can appeal to the archive's own rationale to give at least a partial account. In selecting an archive, therefore, one selects not just that archive's collection, but its selective rationale.

Thus, while the fullest account of *why?* in any selection is no less than the full history of a letter up to and including its moment of selection, we do not need to dig quite as deep. Rather, at this point, we may consider some general tendencies. The first is a predominance in archives of pre-existing collections — that is, where such a collection has been established and 'imported' wholesale. Into this category come letters written to figures of authority, or with some official status. These may even be preserved as distinct sub-collections. The *National Archives of Ireland*, for instance, contains a multitude of official letters written to figures such as Sir Matthew Nathan (the Under-Secretary for Ireland, 1914–16). Others are centred on important individuals, whose letters form part of their estate and are bequeathed to some repository. Letters written to Lady Clonbrock, which are held by the Clonbrock estate, fall into this category. The letters of Sir Wilfred Spender have been left wholesale to the *Public Record Office of Northern Ireland* (PRONI), again forming a discrete collection. Mary Martin's correspondence and letters from Sir Roger Casement may be similarly considered. Other archives are more 'topically' centred: for instance, various military archives deal predominately with military correspondence, while St Patrick's in Maynooth focuses on religious matters.

¹⁷³ Archives are, to a degree, also a striated space: they are full of other 'striating' structures, such as boxes, collections, sub-collections; but also smoothing structures, such as catalogues.

Another tendency that can be observed is for archives to contain letters on the basis of receivership rather than authorship. This should not be particularly surprising: letters are sent, and thus considerably more likely to reside with their recipient.¹⁷⁴ There are, of course, exceptions to this general rule, particularly administrative letters, of which it was common to retain a copy. (Kline, 1995, outlines such arguments relating to correspondence in great depth; moreover, as the next chapter shows, these tendencies can be retroactively inferred by analysing the *Letters* collection.) Another notable exception is the ‘last letters’ of the Rising leaders: these were collected, copied and passed around by supporters. Nevertheless, there is a perceptible orientation towards receivership.

It is interesting to see how these tendencies interact with the various interpretations of “a tangible connection to Ireland”. Selecting letters from archives in Ireland (as *Letters* predominately does) naturally favours letters sent to Ireland, and particularly to Dublin (the most obvious centre of power, and — not coincidentally — the site of many archives), over letters sent either overseas or away from power-centres within Ireland. The latter are much less likely to have been preserved by an institution, or (especially if abroad) relative inaccessible to the *Letters* project. In particular, this points to letters being written to important figures, rather than by them.¹⁷⁵ It also explains the myriad cases in which, of what was obviously a two-way correspondence, only one half survives (or is easily accessible). Sir Wilfred Spender’s letters to his wife Lillian (from France, where he was serving in the British Army, to Ireland) are a case-in-point: we have his letters, but none of his wife’s. This points,

¹⁷⁴ As a general example, Kline notes that much of the *outgoing* correspondence of George Washington is only available due it having been copied into a letter-book. (Kline 1998)

¹⁷⁵ Letters held by the Clonbrock estate exemplify this: there are a large number of letters written to Lady Clonbrock from the First World War (particularly thanking her for food and cigarette parcels) but relatively few letters *from* Lady Clonbrock. In this case, a not insubstantial proportion of the letters written ‘about’ the First World War are relatively homogeneous — precisely because they were written to a single person, with a single purpose.

at the level of the collection, to a greater heterogeneity of authorship than of receivership. It also points to a relative sparsity of letters from what might be termed ‘significant’ figures — except when one such figure writes to another. (Indeed, as we shall see in the next chapter, it is these particular correspondences that transcend the boundaries of individual archival collections and give the impression of a more fully interconnected network.)

Before concluding this section, it is also necessary to consider other factors underlying the selection of specific archives. As seen with Google Books, an element of this is determined by a willingness of the institution to participate. This is dependent on myriad factors, not least archival policy with regards to remediating letters and making them public. A further factor is the proximity of an archive to Maynooth, which increases the already high predominance of Dublin-based archives (Maynooth is 25km from Dublin) compared with smaller archives in other parts of Ireland or elsewhere. It results in, for instance, an abundance of official correspondence that was sent to Dublin, but not that sent from Dublin (now held principally by the British National Archives at Kew — obviously geographically less accessible).

It is, perhaps, a matter of debate whether it is fair to critique a project in such a way, especially a relatively small project with a limited budget such as *Letters*. From one perspective, it is obvious that all projects have to make compromises somewhere, and operate in the real world. Selecting letters ‘closer to home’ is simply making an (eminently sensible) pragmatic decision, and thus not one that can be reasonably criticised. However, this does not mean that such decisions do not have an effect. For *Letters*, these effects are all the more significant because of the open-endedness of selection. By contrast, such concerns are marginalised when

editing of, say, a literary work.¹⁷⁶ One either has the relevant manuscripts or not: a binary state of affairs, in which not having the manuscripts means no project at all. *Letters*, however, is not in the same all-or-nothing situation. It operates on the letters that it finds. As such, where it looks for letters is significant, and a whole host of real-world factors enter into this consideration. The question is, what is the effect of such considerations, particularly when mediated by further stages of the transmission process?

5.4. The transmission process

This section will look at the process of digitising the letters collected by the project. As with all transmissions, the entire process can be conceived of as comprising a whole — a single interpretative act — but can also be decomposed into interpretative functions. What produces a particular outcome, is, as we have seen, the units (assemblages) on which a given function operates, the point at which an interpretative function is created (and from what), its relative isolation from other functions, and the organisation of those functions into linear or recursive arrangements. In this analysis of the *Letters* transmission process, we shall see that — in common with most digitisation projects — it is largely linear in its organisation; however, there are also points of recursion. What is fundamental to understanding the outcome is to see what scope these recursive functions have (how far they can reach backwards or forwards to alter the outcomes of previous functions) and how much is fixed, unalterably, by other functions. In conducting this analysis, the type of diagram developed in Chapter 3 will prove useful.

¹⁷⁶ Even then, not *totally* marginalised. Eggert, for example, notes that some mistakes made by Gabler in editing *Ulysses* arose from his necessarily having to work at times from reproductions rather than the original manuscript. (Eggert 2009, 176)

In this discussion, I will leave aside discussion of letters contributed by members of the public. As noted, these form only a limited fraction of all the letters collected; moreover, once uploaded to the transcription site, these letters are treated the same as any other.

5.4.1. Photographing and collecting preliminary metadata

The beginning of the *Letters* digitisation process is actually concurrent with selection of material. That is to say, once an archive or other repository is chosen as a source of letters, once a letter matching the criteria is found, it is immediately photographed and its related metadata logged. In this respect, the process is different to other projects. As Kline notes, the establishment of a ‘control file’ (a database of possible documents for consideration) is often a precursor to actual selection. (Kline 1998, 60–61) This enables an analysis of the domain and a determination of its boundaries. As the selective approach of *Letters* suggests (see above), this is not possible. Metadata is, of course, logged (usually in an Excel spreadsheet); but this is more of a running accumulation, carried out concurrently with selection. As such, selection itself is a linear function. This is not to say that there is no quality control (especially with regard to image quality; we always check!), but that there is, by and large, no possibility of returning to the archive in question to re-examine or re-photograph a particular letter at a later point in the digitisation process. What is done *in situ* is therefore a fixed point of departure.¹⁷⁷ This is illustrated in Figure 39 (below).

¹⁷⁷ Of course, the degree to which a particular aspect is absolutely fixed depends on the nature of that aspect. Evidently, some things are more easily corrected or normalised than others. The author of a letter, if logged as ‘Augusta Dillon’ can later be normalised to ‘Lady Clonbrock’. The photograph of a page, however, can be cropped and colour-corrected, but not rephotographed without returning to the archive and the original document.

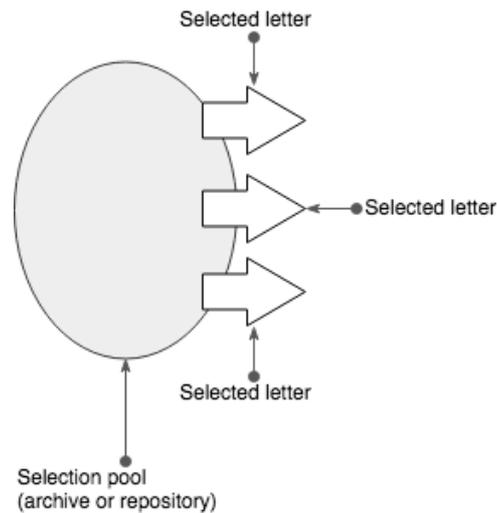


Figure 39: Selection of a letter initiates a linear transmission path

It is, therefore, at this point that the extensity of a letter in a physical (and also textual) dimension is established, at least as a limit. The pages that comprise the letter are thereby determined, as is the association with that letter of additional material and envelopes. Other interpretative components can of course be added to the letter-assembly, but no further text or material can be added, only subtracted (for instance, if a page was later felt to be not part of the letter, or comprised ‘additional material’; see Footnote 171).¹⁷⁸ The selection of a letter is, therefore, a function that invokes a linear path. The act of photographing the letter also establishes each page (or, more correctly, page-surface) as an extensity, to the extent that later actions (such as using image-editing software) can only subtract or divide. Each photographed page-surface invokes its own transmission path. (See Figure 40, below.) This is particularly significant, as it has a number of effects that are felt later in the transmission process.

¹⁷⁸ This is, I think, obvious enough. Adding more pages to a letter (and thereby more text) would involve a return to the original documents. There is a possibility that two letters that were initially considered separate turned out on closer inspection to be a single letter, and were merged accordingly. Or, indeed, that a single letter turned out to be two. Such instances, in my experience of the project, are rare, if they occur at all.

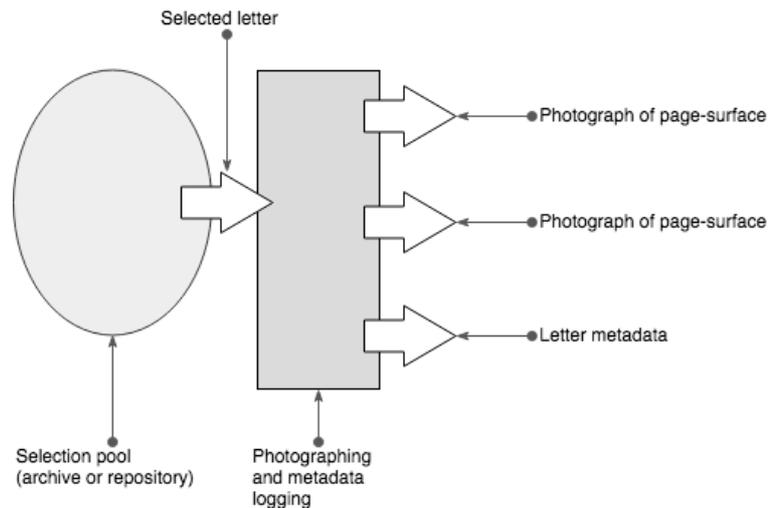


Figure 40: Photographing and metadata logging is a dividing function, producing multiple outputs on discrete paths

The most obvious is that it ‘locks in’ an interpretation of *what is a page-surface*. We see this most readily in the case of letters written on a piece of paper folded in half like a book. The outside pages (1 and 4) are typically photographed separately (even though they comprise one surface of the page when it is unfolded). Pages 2 and 3, however, are very often photographed together as a single surface. In this, we see the interpretation of the text already overriding the letter as a material document — the establishment of a sequential delimitation of the text over the top of a material delimitation. (Pages 1 and 4 are photographed separately as they do not produce a textual ‘flow’.)¹⁷⁹ There are other cases — a kind of logical inverse — where the letter’s text concludes in the blank space at the top of the first page; in this instance, the letter text is necessarily subordinated to the physical unity of the page. None of this is to say that these decisions are flawed; indeed, they represent pragmatic solutions to the evident fact that a letter is a rhizomatic assemblage, which must be interpreted (some strata

¹⁷⁹ See, for instance: Amelia Jephson. “Letter from Amelia Jephson to Nancy O’Rahilly, 14 June 1916.” *Letters of 1916*. Schreibman, Susan, Ed. Maynooth University: 2016. Website. <http://letters1916.maynoothuniversity.ie/explore/letters/1055>.

elevated, some subordinated). However, this does not mean that these interpretations do not have an effect, especially as the ‘page surface’ is the default unit for the transcription site.

5.4.2. Uploading letters

The next stage of the transmission process involves the uploading of page-images and letter metadata to the public-facing transcription site (now a subsection of the main *Letters* site behind the ‘Contribute’ link). The transcription site comprises three parts: a customised Omeka platform using the Scripto plug-in and a transcription interface modified from that of *Transcribe Bentham*; a custom-built Ruby On Rails upload tool; and a MediaWiki installation to handle version control. Though registered members of the public do have access to the upload page (see above), in general the task of uploading is carried out by project team members.¹⁸⁰

It is at the upload stage that what we might call a digital ‘letter-object’ is created in the Omeka database: this comprises, initially, the uploaded page-images for that letter, preliminary metadata entered into fields corresponding to the Dublin Core standard, and a unique identifier that is persisted throughout the transmission process. Of most interest here is the metadata standard. Dublin Core comprises a relatively small vocabulary (fifteen fields) designed to describe physical objects (such as artworks) as well as digitised entities. As such, it is a relatively blunt instrument for describing the potential complexities of correspondence; on the spectrum outlined in Chapter Three, it can be regarded as a low-specificity model.¹⁸¹

¹⁸⁰ Or under their supervision by team members at public outreach events.

¹⁸¹ Much of this, it should be noted, is the result of its implementation in Omeka. For example, while Dublin Core allows multiple authors, Omeka presents a single database field. In a paper (appropriately titled “Full Dublin-Core Jacket”) at the DiXiT Convention at the Huygens Institute, Roşu describes the particular difficulties in using Omeka as the primary infrastructure for a digital scholarly edition. Many of the problems she encountered resulted from attempting to match Dublin Core fields to the manuscripts she

To partially mitigate against these restrictions, further fields have been added by the *Letters* project: most obviously, *recipient*, *recipient's location*, and *sender's location* (sender being equated, in this case, with the Dublin Core 'author'). Though this model can be seen as quite restrictive, it does not reject any letter on the basis of non-conformity (quite unlike mass digitisation in this respect) — even if some temporary 'massaging' of the data is required to make it fit. This is, in part, aided by a lack of normalisation at this stage (if it were indeed possible; the linearity of selection goes some way to preventing this, as there is nothing established — such as a control file — against which to normalise).

Three other important fields are completed at this stage. The first is a description of the letter, which eventually becomes the letter summary in the final output. The second is an arbitrary number of 'keyword' descriptive tags, drawn from a pre-established, but not explicitly controlled, vocabulary. The third is a 'subject' field. This is, essentially, an 'elevated' descriptive tag, corresponding to the most prominent theme of the letter. In the TEI document created later in the process, this subject field is relegated back to the same level as the other tags, and hence is not permanent. However, it forms the key mechanism for organising letters within the transcription site (letters are organised on a single page corresponding to this 'subject'). As such, as Schreibman suggests, the assignment of a particular topic can be instrumental in encouraging transcription.¹⁸² As such, a degree of latitude is (quite justifiably) shown in the assignment of a subject. This, of course, would be a problem if the subject assigned were permanent.

was trying to encode. (Roşu 2015) Such problems for *Letters* are fleeting, however, as letters are later encoded using TEI, which affords much greater flexibility.

¹⁸² As Schreibman *et al.*'s research into crowd-transcription practices has shown, certain subjects are substantially more popular than others, with strong preferences shown by various demographic groups. (Schreibman *et al.* 2017)

The same applies, I think, to many of the other metadata fields completed at this stage. The requirements of Dublin Core and, particularly, Omeka are too restrictive to describe letters, beyond their more generic document characteristics. Later conversion to TEI, as part of a highly recursive and validated process, is therefore fundamental in overcoming such restrictions. This also raises the question of whether it is preferable to utilise off-the-shelf tools (which are inevitably more generic in the interpretative models they impose) or to build a highly tailored tool from scratch. The decision taken by *Letters*, to use Omeka but sufficiently modify it, is, I would suggest, justified from a number of angles. As a public humanities project, generating public engagement from the outset was vital — which meant having a viable, working transcription site immediately was preferable to having ‘the perfect’ transcription site several years down the line. Secondly, practical experience — especially when it comes to building a tool for non-experts — can only be gathered from having a tool and having it used; a ‘perfect’ tool built without this is unlikely to be perfect.¹⁸³

5.4.3. Crowd transcription and mark-up

Having been uploaded, letters are then made available for transcription. The transcription interface comprises a facsimile of a letter page-surface (with zooming capability) along with a free-entry text box for transcription. The transcription box features a number of buttons for automatically inserting prescribed TEI tags. These elements are, in general, concerned with visual or textual features, and with structural elements specific to correspondence.¹⁸⁴ The choice of elements creates an interpretative model, and, significantly, one that is necessarily pre-determined (that is, it takes as its basis ‘typical’ features of letters, rather than the recursive

¹⁸³ Such experience is now being brought to bear on the design of a single, end-to-end transcription platform, which aims to overcome some of the problems associated with Omeka.

¹⁸⁴ Into the former category fall *hi* elements (with a *rend* attribute to denote underlining and superscript), *gap*, *line-break*, *page-break*, *addition*, *deletion*, *paragraph*, and *note*. The latter category comprises *address*, *date*, and *salute*.

interaction with each specific letter). Of course, this is a pragmatic decision for a crowd-sourced project. Members of the public can hardly be expected to have the knowledge of the TEI. Moreover, creating a large collection to be presented in a single interface demands this kind of uniformity. This does, however, involve a degree of over-determination: letters must be made to fit the model, not the other way round.

Also important is the ‘unit of transcription’ (to borrow Causer’s term; Causer 2012) — in this case, a single page-image, rather than the whole letter. This is typical practice for crowdsourcing, as Rockwell (2012, 147) suggests: dividing a complex task (transcribing a letter) into smaller, autonomous tasks (transcribing a page). Single pages are, however, only *relatively* autonomous: each transcription-page is accessed via links from a page representing a single letter, which are in turn grouped into subject pages. (Figure 41, below, illustrates the transcription process.)

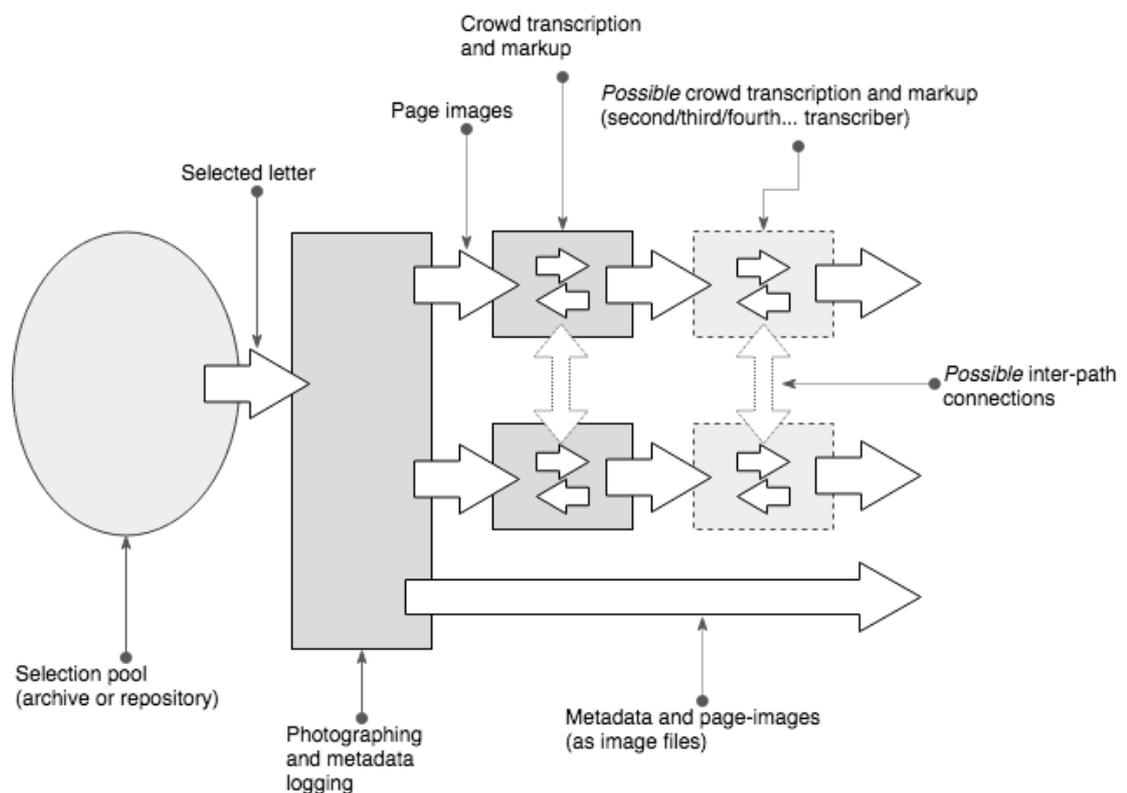


Figure 41: Simplified diagram of the transcription process for a single, two-page letter. The number of transcription functions is not determined; multiple transcribers may return to a given page. Inter-path connections (considering two or more pages together) are also a possibility.

This creates a degree of uncertainty with regards to modelling page-transcriptions as transmission paths. Are they genuinely isolated tasks? Or, as when a single transcriber transcribes *every* page of a letter, is there some inter-path connection? And, in the latter case, does a transcriber behave in a ‘fully recursive’ manner (i.e. returning to alter a previously transcribed page on the basis of knowledge from transcribing a later page)? Or is knowledge only accumulated, only propagated forward to later page transcriptions? Determining an answer to such questions evidently requires particularly detailed user interaction studies, which are beyond the scope of this thesis. Whatever the answer, it is unlikely to be simple, or, indeed, singular. Into this equation we must also add the possibility of multiple transcribers transcribing individual pages of a single letter; or (even more difficult to model) one transcriber partially completing a transcription of a page, which is then completed by another transcriber. There are also a not-insubstantial subset of transcribers who dedicate themselves to proofreading transcriptions. In short, there is simply too much scope for complex user interactions; too many permutations to consider.

That said, there are a number of points that can be raised in this regard. Firstly, a limited study of the revisions stored in the MediaWiki backend that I carried out with my colleague Roman Bleier,¹⁸⁵ suggested that the average letter-page required a mean of 4.2 revisions to reach the stage of ‘completion’ possible within the transcription site (75% of pages required fewer than five revisions). Perhaps more significant is the wide variation in the amount of

¹⁸⁵ Presented as a paper at the 2016 TEI Conference, Vienna: “Capturing the crowd-sourcing process: revisionDesc and storing the stages of crowd-sourced transcriptions in TEI” (Bleier & Hadden 2016) The research was carried out on 873 letters that were completed at the time.

work done towards the completed page transcription by a single transcriber.¹⁸⁶ In our limited analysis, many pages were largely (more than 90%) completed by the first transcriber, while in other cases, the first transcriber completed as little as 30% of the transcription.¹⁸⁷ As such, decomposing the transcription process into discrete, mutually recursive functions is evidently highly complex — and more complex than analysing the thought processes of a single individual.

Though, as suggested above, we can never be sure of the degree of isolation of each page-unit transcription path, there are some clues. Most notable is the case of a paragraph that crosses a page boundary being encoded as two separate paragraphs. In such cases, a ‘correct’ TEI element would cross the page boundary (i.e. place the opening tag on the first page and the closing tag on the second). However, carrying this out evidently requires both pages to be considered simultaneously. That (almost universally) single paragraphs in this situation are encoded as two separate paragraphs points to each ‘unit of transcription’ being carried out in relative isolation.¹⁸⁸ A similar conclusion may be drawn from the relative absence of *page break* elements at the end of a page (despite this element being available in the toolbar): transcribers view the page as a discrete entity, not as part of a single letter.

In general, therefore, we may consider crowd-sourced transcription to be a highly recursive process, with implicit ‘does this make sense?’ validation carried out by a single transcriber,

¹⁸⁶ Calculated by diffing each saved revision from the final version.

¹⁸⁷ Interestingly, for many letters, when the total amount of work done was summed, it came to more than 100%. This suggests that an initial decision taken by one transcriber was altered by a second transcriber, before being changed back by a third, pointing to the transcription stage being not only locally recursive (see Chapter 3) but subject to multiple ‘validation passes’.

¹⁸⁸ That said, it should also be noted that encoding-as-a-single-paragraph also requires an understanding of the semantics of XML — which is not necessarily a requirement when a transcriber only has to click a button to insert tags.

and also re-validated by later transcribers. It is certainly not a one-shot, linear process as carried out by OCR engines. Certainly, human transcription (even not-expert transcription) is more adaptable, generating a high-specificity model out of interaction — for instance, of previously unseen handwriting. Moreover, unlike mass digitisation processes, it is a permissive model: letters are not rejected wholesale due to a failure of interpretation at any given point (compare this to the rejection of books for not fitting into a scanner).

That said, members of the public are not experts either. This is most readily seen in the utilisation of *gap* and *unclear* elements. In the hands of an expert, a gap caused by illegible text, or an unclear tag denoting an uncertain interpretation, can legitimately be seen as actual gaps or lack-of-clarity. Of course, this is not an absolute; but, I would suggest, still a legitimate argument. From a public transcriber, however, gaps or unclear text might be instead be interpreted as ‘I cannot read it.’ This is borne out by an analysis of a sample of 873 completed letters (the same sample used by Bleier and Hadden 2016), which showed a 24.5% decrease in the number of gaps and unclears from the initial transcription to the final, *Letters*-team validated version. Expert interaction, therefore, can be seen as having a higher *criticality* — the ability to interpret complexity (see discussion in Chapter Three).

The question, as elsewhere, is to what degree this matters: what becomes fixed at a given point, and what is open to revision. As well shall see, the *Letters* process comprises many further stages that are essential creating the end product. In part, these are due to the inability of the crowd — or, rather, the crowd and the transcription tool in concert — to produce valid XML documents. This is not to say that public transcribers cannot, on the whole, accurately add TEI tags to the text. Rather, it is the requirement of crowdsourcing in general to present discrete tasks — a page as a unit of transcription — that precludes the necessary considerations for producing a ‘holistically’ valid document. Simply concatenating the page

transcriptions in order does not achieve this. (See, for instance, page breaks and page-crossing paragraphs discussed above.) The requirement of further stages therefore opens the possibility for (in theory at least) completely disregarding the results of the transcription process. Of course, this does not happen — but the question then becomes, what *is* modified? And what are the effects of *not* modifying other things?

5.4.4. Automated processing and creating XML documents

The third stage of the *Letters* workflow extracts the transcribed letter-data and metadata from Omeka. This is carried out on letters verified in Omeka by a project team-member (this final Omeka-based validation achieves what might be called a perfect imperfection: transcription and, specifically, TEI tagging perfected within the constraints of Omeka). Data is extracted from Omeka/MediaWiki as a single data-dump, which is then fed into a series of Python scripts that transform the data for each letter into (again imperfect) TEI documents. This is the area of the project in which my direct involvement was greatest: both the Python scripts and the later editing platform were built in collaboration with my colleagues Linda Spinazzè and, later, Roman Bleier.

As with many aspects of the project, the decisions taken at this stage were inherently pragmatic solutions. The chief difficulty in automating any process is to be assured of the validity of the inputs, something evidently problematic when crowdsourced, and when the well-formedness of the transcribed XML could not be relied on.¹⁸⁹ The approach taken was therefore to produce a highly fault-tolerant set of tools that would automate as many tasks as possible, while relying on a further editorial phase to ensure the validity of the result. This

¹⁸⁹ Indeed, as noted above, the page-at-a-time approach sensibly taken to facilitate crowdsourcing prevents the simple concatenation of marked-up page texts into a well-formed XML document.

may be considered a kind of Unix philosophy (normally phrased as ‘Worse is better’, or ‘Ninety per-cent of the work done for ten per-cent of the effort’).¹⁹⁰ Another approach taken was to design the processes as modular, and with the intention of being run on batches.

This latter decision was again a pragmatic decision. At the point at which it was implemented, the Omeka installation was still run from servers at Trinity College Dublin (where the project was initiated) and had not yet been migrated to servers at Maynooth. As such, extracting data could not be done easily or as-required; instead, it was sent as a dump of the entire database. The processes also required the pulling of data from lists of normalised names and places, as well as other control files, which (to be utilisable by a large team) were stored as shared spreadsheets. As processing even a single letter required the assembly of up-to-date versions of all these documents, it made sense to process batches. This, of course, has a downside in that it creates a bottleneck, which was often problematic. (The solution to such problems is the creation of a more comprehensive end-to-end system, such as currently in development at Maynooth.)

The Python scripts start by extracting letter data from the data dump, creating a letter object in Python, identifying the most recent revision, and storing previous versions for injection into the *revisionDesc* element of the TEI document. These letter objects are then operated on by further processes. One process attempts to ‘correct’ XML tag errors (transcribers often accidentally delete angle-brackets). This does nothing (except accidentally) to tackle problems of well-formedness; however, run on a test-batch of 1953 letters showed that 40% (788) had at least one tag-error of this nature, which this script fixed. Even with sophisticated XML editing software, manually tracking down such errors would take a considerable length

¹⁹⁰ Or, the other way round, making a system perfect for all contingencies takes up most of the time, which could be spent doing more in a not necessarily perfect manner.

of time. Further scripts automated the normalisation of named entities (senders, recipients, and their locations) using fuzzy matching.¹⁹¹ Another process attempts to determine page-types (envelopes, for one; but also the sides of postcards) by identifying a significant portion of the transcribed text as an address; yet another tries to identify pre-determined structural elements of a letter (headers, closers) by looking for collocations of crowd-added tags (specifically addresses, dates and salutes). The final script in the toolchain combines the modified transcription and metadata together into TEI file (one per letter), with appropriate page-breaks (including references to the image files). The scope of these operations is shown in Figure 42 (below).

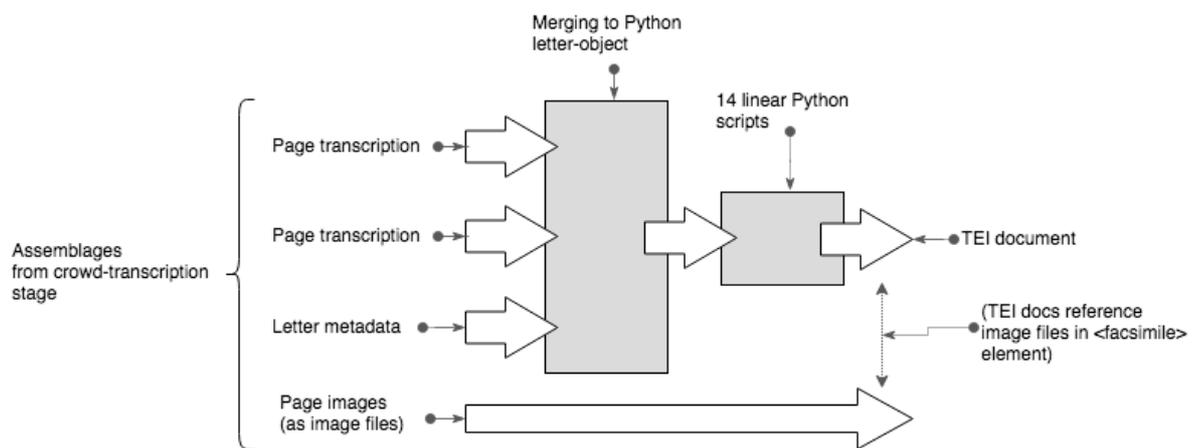


Figure 42: Simplified diagram of Python scripts (for a two-page letter)

These processes are, perhaps, the closest *Letters* comes to mass digitisation-style automation. They are evidently linear, apply very restrictive (and necessarily pre-determined) models — especially in their imposition of a very generic letter structure —, and produce output that are somewhat dubious in their quality. Very seldom were valid, or even well-formed, TEI

¹⁹¹ This, while largely effective, was *never* trusted to produce the correct result. Where such entities occurred in the TEI documents produced, they were accompanied by a (non-TEI) <VERIFY> element, which a later reviewer needed to remove in order to validate the document.

documents produced. However, there is, I think, a fundamental distinction to be drawn: these processes were intended to speed up otherwise time-consuming work (for instance, manually copying and pasting, or re-typing, metadata into a *teiHeader* would undoubtedly be an extraordinarily longwinded process); their intention is not to produce a final result.

5.4.5. Final editing and validation

In this final stage, we see a move away from the mass collection of data, public-crowdsourcing, and automated processing, towards a more traditional approach (in line, to some degree, with a critical digitisation project). At this stage, project team members have access to the entire TEI document produced by the Python scripts, and are thus able — for really the first time since the initiation of the digitisation process — to deal with a letter as a distinct entity, rather than as disparate pages. This means that, not only can valid and well-formed TEI be produced, but the text be validated in the context of the letter as a whole.¹⁹² Moreover, it is carried out by trained project team-members, and thus achieves a high level of accuracy.

However, the fact cannot be ignored that this is a function that is situated linearly *after* the previous stages, not a fully recursive return to the beginning. It is a unique, contingent interaction, in different circumstances, and the effects of previous stages cannot be entirely erased. In particular, it does not try to substantially reconfigure the XML structure previously imposed, nor add a range of new TEI elements that might be suitable for a given letter. Rather, it is a human validation of the structures and elements already added — whose own validity is determined by a series of automated tests, including validation against a highly restrictive TEI schema — as well as a proofreading (rather than a transcription from scratch)

¹⁹² For example, it is the first time that the text of a letter that ‘wraps round’ (to conclude in the blank space at the top of the first page) can be properly treated as a linear text.

of the letter text. Shillingsburg —undoubtedly bringing his knowledge of publishing houses and typesetters to the fore — describes the differences between transcribing and proofreading, and the perils of the latter with respect to accuracy: a proofreader “risks dizziness and a crick in the neck”, as he eloquently puts it. (Shillingsburg 1996, 135-6) This is not intended to cast aspersions on the diligence of the *Letters* team members — but no process is error-proof, and proofreading large batches of letters (while at the same time checking tags and getting XML to validate) is a tiring business.¹⁹³

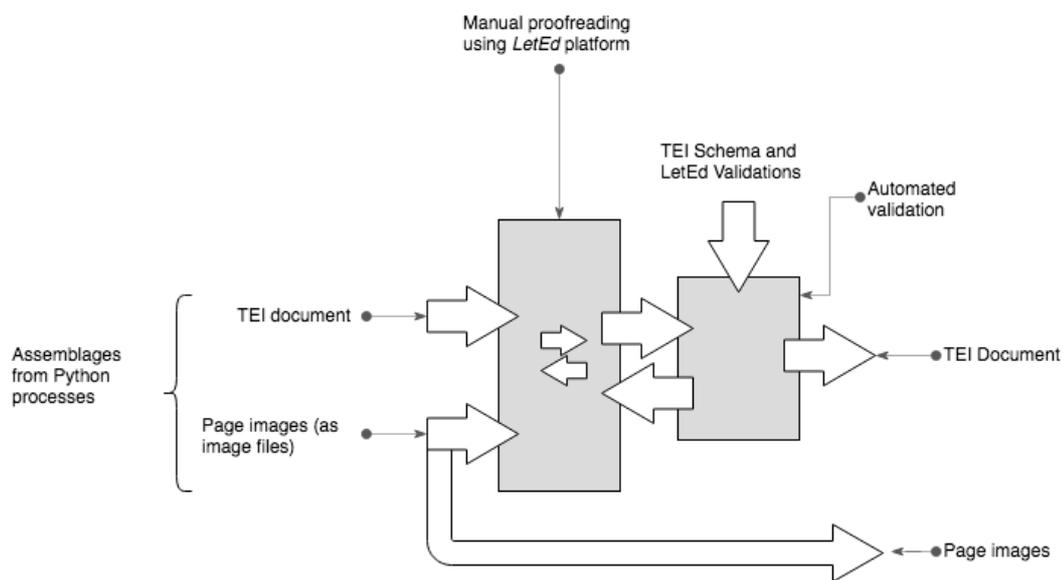


Figure 43: Diagram of the final proofreading stage (for a single letter)

The verification of letters is predominately carried out using a tool called *LetEd* (the building of which was my other significant contribution to the project: *LetEd* — “Letter Editor” — because I have no imagination). Initially intended as a server-based Git repository to allow collaborative work on the letters collection, it spawned a web-based front-end (managed using a Python Flask backend) that allowed editing of the complete TEI document, as well as dynamically loading the page images and providing an on-the-fly ‘rendered view’ of the

¹⁹³ *Letters*, moreover, admits as much and makes it very easy to report mistakes. The project also has two advantages: it is digital, so (the relatively few) mistakes can be easily rectified; and a large number of users, many of whom are also engaged in transcription, and are highly attuned to error-spotting!

letter text to aid proofreading.¹⁹⁴ It is notable — I think, even, unique — for using the TEI *revisionDesc* to log all changes to the document *in the document itself* (alongside previous activity carried out on the letter in Omeka and by the Python processes). Accordingly, it is possible to access the complete history of a letter as a digital object.¹⁹⁵ The *LetEd* platform also incorporates a highly restrictive TEI schema, against which a letter needed to validate in order to be declared complete. Its suitability for more complex XML editing tasks (particularly where a letter's pages required radical restructuring) was more questionable; in these instances, team members often resorted to more powerful tools such as oXygen.¹⁹⁶

At this point, we have reached the end of the *Letters* transmission process — or, more precisely, the editorial phase. Letters are then uploaded to the Explore section of the site. Throughout the transmission process, individual letters (though added in batches, a batch is entirely arbitrary) are operated on in distinct, parallel paths, in almost perfect isolation, and with little synchronicity. This influences, for instance, the writing of summaries; it is difficult to refer to another letter in a series of correspondence when that letter is at a completely

¹⁹⁴ This view uses background colours to designate TEI elements, so that they can be checked at a glance, as well as aiding proof-reading by freeing the editor from the intrusion of the actual angle-brackets.

¹⁹⁵ This data was used by Roman Bleier and me in our analysis of crowdsourced transcriptions (see above).

¹⁹⁶ *LetEd.*, it must be noted, was far from perfect — again, it represented a pragmatic solution to a particular problem. For one thing, it inherited the batch-of-letters approach from earlier processes (necessarily, as the earlier processes produced batches), which created another bottleneck. Moreover, adding new letters to the system had to be done manually, by uploading them to the server using secure-copy and committing them to the Git repository. The use of Git to manage versions was also questionable. While it served the purpose of version control for XML files, where a commit (triggered automatically by a user saving a document) failed — an occasional occurrence — the system would jam until the commit could be manually resolved (again involving logging into the server). The use of a restrictive TEI schema, while it again served its purpose, also created a bottleneck: by ensuring that normalised named entities were correct, it needed to be updated in order to validate a named entity that had not already been added to the list. Not, therefore, an ideal design; but one whose chief problems can really only be solved by incorporation into the kind of streamlined, end-to-end platform that is now in development.

different stage of the process, and is, anyway, a challenge to identify. It also precludes the use of formal encoding in the same vein (for instance, employing the *correspContext* element). From a purely editorial standpoint, this might be seen as less than ideal: as Kline suggests, starting with a rigorous selection, carefully logged in a control file, allows the determination of context — in particular, the relation of letters to each other.

However, it is also a question of priorities and pragmatism, a perspective from which the approach taken by *Letters* is perfectly justifiable. As a public humanities project, the crowdsourcing of transcriptions means having a constant throughput of new letters — to ‘keep the crowd fed’, as it were. It is also doubtful whether a project of this scale (and deliberate inclusivity) could be achieved without adopting a more production-line approach. A documentary edition can function in a more artisanal manner, precisely because its endpoint — a discrete, well-defined set (i.e. determined by a regulative idea) — is known. The aim of *Letters* is, rather, to digitise, and to continue digitising, letters. The collection as a whole is, therefore, more accidental than substantive: not only is it not static, it can only really be considered as a set of independent letters that happen to have been digitised in the same way and end up in the same digital space: a thousand cohabiting ‘micro-editions’.¹⁹⁷ Returning to the validation thresholds explored in Chapter Two, we may here see two distinct regimes of representation. At the level of the individual letter, there is a one-to-one (letter-to-digitised letter) relationship that is, by and large, *reflexively* valid. At the level of the collection as a whole, the regime is one of *articulated* transmission: it can articulate what was done, but not what the result *is*.

¹⁹⁷ This is not to say that a single, cohesive platform is not *useful* — quite the reverse.

In this regard, it is similar to a critical digitisation project (see Chapter Four), with one possible exception: we must ask to what extent a letter is over-determined by carrying out a reflexive transmission under the auspices of a single project — which requires, at least to a degree, a generic approach to transmission — and the requirement to publish that letter in a largely homogeneous platform. With regards to the former, we might say that it has some effect. Evidently, the models applied are pre-determined, and hence normalising. However, by taking a view of letters as text-bearing objects (rather than explicitly physical documents), such normalising is — to a much greater degree — built in as a rationale. Text is, after all, a highly abstracted representation of the document. The nature of the output will be explored in the final section of this chapter.

5.5. Outputs and spatiality

In this section, I will look at the project's 'official' output:¹⁹⁸ the *Explore* section the *Letters* site. The key questions to be answered are — given the discreteness of, and independence with which, each letter is transmitted — to what extent is this discreteness perpetuated in the *Letters* output; and, simultaneously, to what degree do the mechanisms of a single homogeneous platform conjure the idea of signifying collection? Much of the theoretical basis for answering these questions is discussed in the final section of Chapter Three.

To begin with, we can point to the Web as a medium. As argued previously, the boundaries of a website are considerably less distinct than that of a book, which mitigates against the perception of a signifying whole. (A book, as discussed, presents a notion of unity between its covers.) Likewise, a website does not, by default, impose an order on its component parts.

¹⁹⁸ In addition to this Web interface, complete TEI documents for each letter are also created, which can — and have been — employed for other kinds of research. These will be considered at the end of this chapter, and more fully in the next.

The *Letters* site is organised into several, fairly rigid hierarchical arrangements, accessible from various index-like pages: a *month* index; a *category* index; and a full-text search that builds an index page on the fly. Access to letters is only possible by descending from these index pages; there are no other direct lines of movement (for instance, between letters in a correspondence, or via some connecting entity such as a person or place). Each letter is present in a silo of its own web-page¹⁹⁹—continuing the parallel transmission paths into the user phase. (See Figure 44, below.)

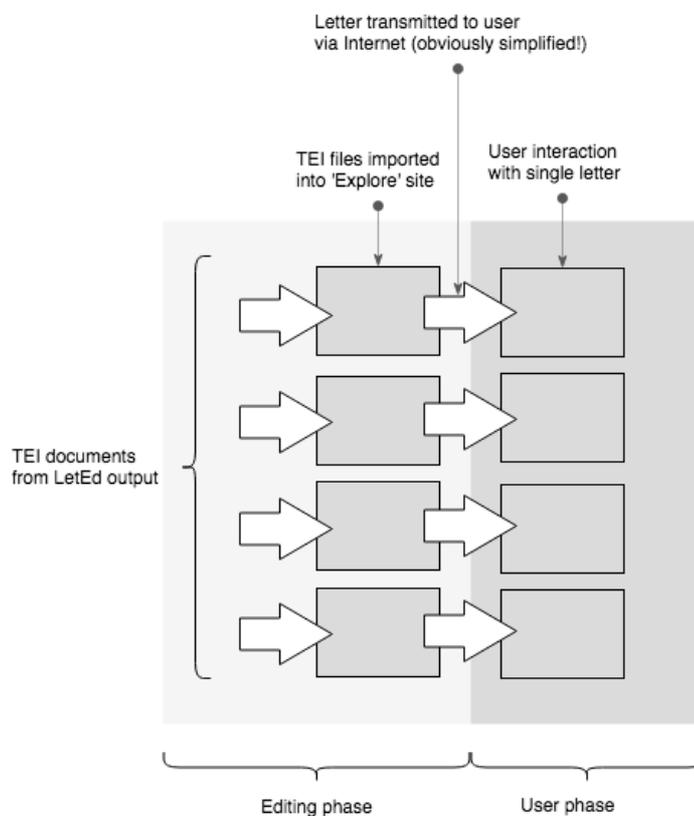


Figure 44: Letters uploaded as discrete entities (and not synchronously)

It is important to recognise what the effect of other kinds of movement would be. Take, for instance, an arbitrarily-introduced linearity of a correspondence between two individuals.

¹⁹⁹ These are referenced by a single URI, using a REST-like naming convention, i.e. `/letter/1769`.

This has the potential to be misleading — for certain, the validity of such movement could not have been established. If, for instance, one could move in order through the correspondence between Sir Wilfred Spender and his wife, the impression would be created that he wrote all the letters and she never replied. Of course, this might be the case, though more probably her letters were simply lost. Treating the letters as self-contained entities, avoiding the imposition of un-validated lines of movement, mitigates against the creation of such impressions. The metadata is there, which would allow a user to discern such connections if they wished; but the onus in this case is on the user — it would not be presented as a ‘fact’ of the collection.

This page presents what might be called a ‘pluralistic’ view of a single letter: there are three distinct representations. The first is a representation of the text of the whole letter, produced by concatenating the pages and removing some TEI elements (most notably, line breaks and deletions)²⁰⁰ to achieve a flattened reading text. The second and third views are side-by-side views of the text of a page (this time with line-breaks and deletions maintained) and the page facsimile. These views are available from a single web page. Each letter page also includes the letter metadata, and the letter summary. Being presented in such a way subsumes the metadata elements as *properties* of a letter — even when, as in the case of people and places, they could conceivably stand as distinct entities. The space created thus far can, therefore, be seen as a mixture of smooth and striated: not even, at this point, a patchwork; but patches, with a potential to be joined together. The space is organ-ised (to use Deleuze’s term), with

²⁰⁰ This approach did, initially, produce a problem for lines ending in a hyphenated word: removing the line-break would leave the hyphen (seemingly inexplicably) in the middle of a word presented on the same line. This was solved by introducing a small post-proofreading script that converted literal hyphens preceding line-break into a line-break with *rend*=“hyphen” attribute.

each letter a clearly distinct entity with few lines of movement; but no totalising organisational model.

Further striations are, undoubtedly, overlaid by the various ‘index’ pages, which organise the letters into categories or date ranges. Such pages inevitably give greater prominence to the collection as a whole — though the degree to which this is the case is variable. We might, for instance, contrast a contents page and the index of a book. The former more obviously gives rise to the impression of a cohesive whole (it evidently marks the extensity of the contained material), while an index is more of a finding-aid. The ‘index’ pages (categories, date-ranges) in the *Letters* site fall, I think, somewhere between the two. As a contents-page, they are only ever partial, being restricted to displaying a subset of the collection. Indeed, there is no complete content page, except by performing an empty full-text search. Even then, the results are so large that the extensity of the collection is hard to grasp.

That said, it is also important to consider where the overlaid striations imposed by these index pages originates. In the case of *category* index pages, these are derived from the letters themselves, though normalised to an extent. Categorisation is a necessarily reductive operation. Indeed, that is the whole point. However, as a letter can be assigned to multiple categories, the relative independence or overlap of these categories must be considered. For example, such categories as “Irish Question” and “Easter Rising Ireland 1916” both appear to contain a lot of letters — but, given the obvious overlap in such categories, this is misleading: they both contain a lot of the same letters. Moreover, in such overlapping cases, there is the possibility that one category serves as a kind of proxy for another. A letter that was included in (say) the categories “Art and Literature” and “Medicine” can be more readily assumed to talk about both these things. A lot of this depends, clearly, on the interpretation the user places on the category. That is, it is a useful finding aid, or actually some signifying

entity? Are these striations possible lines of movement, or are they walls that divide the collection and to which might be assigned various properties?

Of course, we can never know what a user thinks or does, or build in absolute restrictions. This problem points back to the notion of a ‘contract’ between the editor and the user, much of which is determined by convention. For example, the footnotes of a critical edition are, by convention, the critical apparatus, not some funky footnotes produced by the author of the work. Web-based textual projects, needless to say, are far from establishing such things as conventions. However, with regard to categories, *Letters* also presents them as additive filters, rather than exclusive categories, providing strong hints to the user that they are really finding aids as opposed to something signifying.

Date-range index pages (representing months from November 1915 to October 1916) are, in some sense, more arbitrary striations. This might appear strange: categories are, after all, based on considerably more interpretation than the date of a letter. However, we are not really talking about the date of a given letter, but the *organ-isation* of smooth time into months — which are arbitrary creations, and, significantly, little to do with letters.²⁰¹ (Days of the week — especially with no post delivered on Sundays — are of much more importance.) At the same time, month-long date ranges are evidently exclusive categories. As such, and all other things being equal, it is largely a matter of chance that a letter gets assigned to a certain date-range. This is not to say that organising by month is not a sensible approach, provided again that it functions as a finding-aid rather than a kind of categorisation. However, there is, I think, a case may be made for it being (to a greater degree than the categories seen above) the latter. As noted earlier in this chapter, while the beginning and

²⁰¹ The absurdity (*one* of the absurdities) of astrology is its arbitrary divisions of time: be born one day either side of the *Virgo–Libra* threshold and your prognosis is radically different.

end dates of the *Letters* collection are relatively arbitrary, the central pivot — the Rising itself — is not. This attaches a degree of importance to the central months: more likely to be about the Rising.²⁰² Again, so much of this depends on the user’s interpretation: what is required, if such a thing were possible, is a mode of conveying, ‘This is for convenience, not meaningful interpretation.’

These striations, while they apply a kind of global model to the collection as a whole, can nevertheless be seen principally as finding-aids: tools for whittling away a large number of letters to find relevant ones. They open, therefore, potential avenues for movement within the collection, but without completely taking over the collection and re-smoothinging it ‘in the name of’ the collection as a whole.

The final mode of accessing letters is the on-the-fly creation of ‘index’ pages using full-text search (which can, as with other index pages, be filtered by various categories). As discussed in the previous chapter (in relation to Google Books), full-text search necessarily operates on a closed search space. As such, it can be considered as a user-invoked merging of hitherto discrete letters into a collection, and at a specific point in time. (As the site itself notes: “New letters are constantly being added to this database, so search results may differ from time to time.”) To what extent this constitutes a takeover and smoothinging of the collection space depends, to a very great degree, on the search algorithms used.

²⁰² Though, given the date of the Rising (towards the end of April) and the disruption to the postal service that it caused, its greatest effects are felt in May — simply because they cross an arbitrary threshold. Does this mean that more letters were *written* in May because of the Rising, or people just put writing their letters on hold, knowing they were unlikely to be delivered — or even that more letters written at this point were likely to survive and be collected?

Many of the problems typically associated with full-text searches are discussed elsewhere. Mann observes the problem with synonyms: that a user may search for a given term (and find matches) but not find others that evidently carry the same meaning. (Mann 2005) For example, searching for ‘Dublin’ will find matches of Dublin, but not of, for instance, ‘the city’. Beall (2008) identifies other potential difficulties. Of these, variant spellings (particularly of Irish names of people and places: *Patrick Pearse* does not match *Pádraig Mac Piarais*) are most pertinent for the *Letters* collection. A user not knowing an appropriate search term may also be problematic. Lastly, and most problematically, is the absence of a match for a seemingly obvious search term. In user testing for *Letters*, it was found that users were surprised that searching for ‘Easter Rising’ produced no results. This is, of course, because no one writing at the time used the term in a letter (or, indeed, at all).

Partly to mitigate against such confusions, the search space was extended to include metadata (‘Easter Rising’ would therefore match the category label). This, in turn, had a knock-on effect when combined with the search algorithm used. As originally implemented, this returned only matches with a high frequency in specific letters but relatively low frequency in the collection as a whole. Adding metadata to the search space elevated the presence of certain terms (particularly the category labels), even for letters whose text did not actually contain that term, resulting again in no results for common searches.²⁰³

This kind of relevance-based searching is particularly pertinent here, especially when compared to a more simple presence-based searching. In the latter case, the search space of the collection is not arbitrarily closed off: rather, indeed, like the *Letters* selection strategy, it simply involves going through every letter and checking whether it matches. Searching by

²⁰³ This was eventually solved by using, firstly, the relevance-based search; but, if no results were returned, searching again by simple presence of the term.

relevance, however, requires a scope — a closed space — against which relevance can be calculated. In other words, it elevates the collection-as-a-whole to a quasi-signifying space. There is no knowing what it signifies in itself, but it produces meaning by operating as a basis against which undoubtedly more signifying elements (i.e. search results) are determined. A search term matching or not is determined by not just the letter that does match, but by a sufficient quantity of everything else not matching. This is problematic when the collection is inherently highly contingent, and far from static.

A final problem with full-text search (to which I shall return in the next chapter) is the ordering of results. (Beall 2008, 441) Matching a search term is binary, so there is no self-evident way of ordering the results. Hence the success of Google's Page-Rank algorithm (pages being *linked to by other pages* serves as a good proxy for relevance). However, this depends on knowing in advance what you are looking for. Searching for something on the web and exploring a historical collection are two very different tasks. A user wants to find matches to a term as a basis for exploration, not find a specific letter (unless they already know it is there). The particular case in which this might be problematic arises when there are several pages of results returned. A naïve user, accustomed to Google, might well just click the top result. But even when a user is more diligent, a result that appears on the first page is more likely to be accessed than one that appears on the tenth. This is (by and large) fine when the letters are, firstly, organised by some relevance metric, and secondly, when 'the most relevant thing' (according to said metric) is what is sought. Neither of these things are true in the case of *Letters*. What is created is, therefore, an arbitrary striation of the collection space, one determined by a likelihood to be accessed.

To be clear, this is not a problem unique to *Letters*. Such problems are inherent in all full-text search systems, especially those for humanities research collections. Against this, again,

must be weighed the evident utility of being able to carry out full-text searches — and most projects, quite naturally, tend to the utility. Again, the problem can be best considered as a project–user ‘contract’: a search tool should be clear about what it does, how it does it, and what the user should be wary of (over)interpreting from the results.

5.6. Conclusion

In this chapter, we have put to use many of the concepts developed earlier in this thesis to describe the *Letters of 1916* transmission process and its outcomes. In modelling the process as such, it has been possible to consider each interpretative step in the process, the local conditions in which they were carried out, the levels of criticality brought to bear in each instance. *Letters*, as suggested earlier, can be seen as a combination of many hitherto existing paradigms: mass digitisation, crowdsourcing, documentary editing, critical digitisation. Should we, therefore, attach a label to it? It might be useful to do so, but it would not truly reflect the project. *Letters* is, really, what it does.

As such, studying the process of transmission is the only way to consider what it is. In answer, I think, it is a largely heterogeneous assemblage of epistolary documents, taking as its centre a historical event, but in no way restricted by it. It has a rationale of openness, which precludes saying what the collection as a whole *is* — which is to say, what it represents. It does not represent an entity — physical or conceptual — except ‘openly’, contingently. This does not stop each digitised letter *representing*, and representing well, its physical counterpart. Moreover, as we have seen, the final output goes a long way to replicating and enforcing this sense of ‘letters collected’, as opposed to a signifying ‘collection of letters’.

That said, these restrictions, though necessary, are restrictions nonetheless. *Letters* is a collection that can only be *read*, and read at the level of the individual letter. As noted earlier,

the project also produces fully marked-up TEI documents, complete with correspondence descriptions and unique identifiers. The potential is there for the data, more holistically, to be studied *en masse*. Set against this is a lack of certainty of the effects of contingency: one can never be sure what is being studied in terms of what it represents beyond itself. This is what distinguishes a scholarly edition, as Bode (2017) argues. There are, I think, two solutions to this. One is to introduce certainty: analyse the source material, painstakingly construct a regulative idea. But, as experience with *Letters* shows, such a task is practically impossible to achieve at scale; and actually impossible to achieve while also producing an inclusive and heterogeneous collection — the two are anathema. Another approach, which will be more fully considered in the final chapter, is to not create a regulative idea; not to attempt to factor out contingency, but instead build it in. The result is, of course, passing the buck to the user — but at the same time giving them the tools and information to deal with it.

6. Towards a Rhizomatic Edition

6.1. Introduction

This chapter will discuss the possibilities of a different conception of the *Letters of 1916* output. It builds on the theoretical model of textual material, in particular *accumulated collections*, as an assemblage (as well as re-using much of the work carried out by the *Letters* project team). It suggests that the collection be considered as a rhizomatic assemblage, rather than a hierarchical signifying assemblage; and that by mapping such a system as an interactive network on which to experiment, an alternative conception of a scholarly edition can be created.

I will deal firstly with an outline of the theoretical conception for this output before moving on to discuss the technical challenges of building a system capable of representing a large collection as a network. The chief difficulty encountered is designing a system capable of both representing the entire network in a useable manner, and at the same time allowing access to the text of the letters as documents. Typically in Digital Humanities projects, as this chapter will discuss, the two are regarded as separate activities. The digitisation of text (whether it is considered scholarly editing or not) is normally treated as a preliminary phase, with its own particular output: the presentation *as text* of the text of the documents. Visualisation of the entire corpus, whether as a network, or via some other processes such as topic modelling, is considered a secondary activity, under the remit of what might be loosely described as ‘further research’. The conception presented here is different: it suggests that visualisation of the entire collection is not just a *post hoc* activity, but the primary mode of representing the collection and a gateway for accessing the text of individual documents.

In this respect, the approach advocated here is not startlingly different from the notion of the *work site* (Eggert 2005; Gabler 2010) or a *knowledge site* (Shillingsburg 2009). Both can be seen as conceiving of the digital scholarly edition as a “web of discourses”. (Gabler 2010, 44) It differs, however, in a number of critical ways, which really stem from the nature of the *Letters* collection. The above notions of the work site are explicitly concerned with a more clearly delimited system, organised around the literary work. Therefore, they have a ‘centre’ — the conceptual work itself — around which to organise the material documents. Moreover, this suggests a clearer hierarchy: material versions of the work (authorial manuscripts, fair copies, amended proofs, various print editions) are closer to the centre, while at the periphery one may find other pertinent (but implicitly less so) documents, such as the author’s letters.

This is not the case with *Letters of 1916*. It is not a literary work; nor is it the correspondence of a single individual or set of individuals. Rather, there is no conceptual centre — no regulative idea that determines the content. It is open in too many dimensions for this to be the case. Accordingly, the *Letters of 1916* project itself is the only thing that, in any complete sense, ties together the collected letters into a single system. Its boundaries are, if not arbitrarily set, then arbitrarily realised. As such, each letter in the collection is of ostensibly equal standing, and the connections between them heterogeneous; they do not ‘point back’, as it were, to the central notion of the work. As a “web of discourses”, therefore, it is of quite a different nature.

The second point of distinction, I think, is in its realisation. These models envisage, I think, a very similar sort of realisation to that of, for example, the *Rossetti Archive*. While conceptually underpinned by rhizomatic assemblage-like conceptions of the work (the ‘social text’ for instance), in their actualisation there is still a marked document-prioritising

approach: each document corresponds to a single web page, to a single URL. The ‘network’ aspect is realised as a kind of emergent phenomenon, the result of connecting together pages with hyperlinks: a very Landow-like notion of hypertext. The realisation intended here, by contrast, is one in which the network itself is foremost: presented as a network, on a single page.

There are, I think, real benefits to such a conception, particularly in the case of the *Letters* collection. (Its possible relevance to other projects will be discussed later in this chapter.) As argued in the previous chapter, the processes of selection and transmission mean that only the individual digitised letter, standing alone, is a reflexive transmission of its physical counterpart. At wider levels — of the collection as a whole, of any combination of letters — transmission is not reflexively valid: the extent of the collection (and, thereby, any subset extracted from it) is indeterminate. The *Explore* site mitigates against the consideration of these wider levels by treating the letters as highly discrete entities. This is not without its disadvantages, however. Firstly, it undermines the nature of letters themselves: letters are not just extensively delimited documents or texts, but specific communicative acts from one individual to another. Subsuming this communicative role under the letter-as-object is, in this regard, to dehistoricise them. Secondly, such a necessary focus largely undermines the possibility of using the collection as an entity, for any kind of computational analysis or representation, such as topic modelling (which means losing one of the benefits of a large textual collection).²⁰⁴

²⁰⁴ Topic modelling can of course be done, but only in a tentative, experimental manner, and any claims made on the basis of it must be underpinned by a deep understanding of the collection itself. To give an example, an early topic model made of the *Letters* collection produced one topic consisting of keywords such as ‘cigarettes’, ‘food’ and so forth. In one sense, this *is* a topic: it comprised at least a tenth of the collection. A deeper knowledge of the collection, however, made it plain that such letters were all written to Lady Clonbrock (thanking her for the food and cigarettes sent to them as prisoners of war). It is not, therefore, a fact that a large number of letters written in 1916 were about food and cigarettes, but that

The reason for this, as Bode says of Moretti and Jockers' use of literary datasets, is that such a dataset — where it is arbitrarily assembled — cannot be simply assumed to represent something. (Bode 2017) 'A lot of books from the nineteenth century' is not a representation of nineteenth century literature.²⁰⁵ (Not, of course, a problem if you just read the books; but highly problematic if used as a basis for making arguments about nineteenth century literature.) The same logic, I think, applies very strongly to *Letters*.

Bode's proposal is, accordingly, to create a 'scholarly edition' of such things as literary systems, that would validly represent that actual literary system. (Bode 2017, 13–7; see also McGann, 2014, who argues for grounding of digitisation practices in a philological model.) Such a model is predicated on a high level of curation, as well as the creation of explicit textual notes, articulated methodologies and rationales — the features, essentially, of any scholarly edition, only shifted up a level in their scope. The activity of creation is most significant, as it evidently gives rise to a *reflexive* transmission. Curation in this sense can be considered as, to coin a phrase, 'selection with due diligence': that is, analysing a domain for selection, selecting a thing from it, returning for more analysis... with the goal of carefully (reflexively) establishing the extent of the domain. Which is to say, developing a regulative idea.

a large number of the letters collected were from the Clonbrock archive. Of course, one should know one's data — though this creates a kind of hermeneutic circle: how is one supposed to know the data when there is so much of it (being a reason to topic model in the first place?)

²⁰⁵ "The meaning derived from a literary historical dataset – like the interpretation of an individual literary work – is shaped, profoundly, by the theoretical framework through which it is approached, and by the selections and amplifications that that framework produces." (Bode 2017, 7)

There are, I think, several problems with this approach if it were applied to *Letters*. These stem, fundamentally, from it being a very different kind of project. Firstly, selection-ascuration involves a comprehensive, and pre-emptive, analysis of the whole domain. Secondly — and logically flowing from the first — it must be, at some point, actually complete according to the limits determined for the domain. Thirdly, without the other two, it is impossible to articulate methodologies or rationales in a way that is reducible to something other than a complete history of the project (or, rather, a complete account of how each letter made its way into the collection). The approach of *Letters*, incrementally (and asynchronously) building up a large collection of letters, cannot be made to fit into such a paradigm. Hence the approach taken here: to articulate — by modelling — the contingent decisions that produce effects as part of the edition.

My argument is that such a representation is a ‘kind’ of scholarly edition — that is to say, in the hands of the user, it can function as a scholarly representation. By explicitly providing and mapping enough information, contingencies are not naively marginalised. As such, it has a different epistemological value to the ‘reflexive transmission scholarly edition’: it does not make the case (to put it crudely) that the buck stops with the editor. There is not, for one thing, a single ‘buck’ but multiple — some of which are passed to the user. Put in these terms, a rhizomatic map of the collection articulates the relevant contingencies — i.e. *which* bucks are passed to the user — so that these can be considered or built-in to any knowledge claims.

The result is, to borrow Dahlström’s description, a scientific instrument, but of a very different order: in the terms of Deleuze and Guattari, it is an instrument of ‘nomadic science’ rather than ‘Royal science’. The latter is concerned with an objective conception of reality,

achieved by rigorously marginalising contingent factors in order to make universal claims.²⁰⁶ Nomadic science deals with actual systems, where contingent factors cannot be marginalised. It is, as Deleuze and Guattari suggest, “[E]ntirely oriented toward an experimentation with the real.” (Deleuze & Guattari 1987, 12)

Before continuing, I should say a few words about this interface and infrastructure (dubbed, by me, *LettersGraph*, as I needed a name for the Git repository). It is an experimental project I have been working on in conjunction with this thesis, intended to illustrate the potential uses of the kind of analysis presented in previous chapters. As such, it should be very much regarded as an ‘unofficial’ project, quite apart from the official *Letters of 1916* workflow and output. That said, it utilises much of the data accumulated in the production of *Letters*, particularly the intermediary TEI documents and spreadsheets of normalised names (along with the data formalised therein). For access to this, I must thank, especially, the project coordinator, Susan Schreibman, but also all those members of the project team whose efforts in preparing the data used by this ‘experiment’ has been essential. A large dataset — particularly a well-encoded one that required (on my part) little to no modification — does not come about by accident, but extremely hard work.²⁰⁷

6.2. Rhizomatic rupture and the internalisation of connections

In this section, I wish to outline a case for considering a kind of documentary edition that moves beyond the traditional encapsulation of information within the parameters set by the

²⁰⁶ For instance, in a scientific experiment where all other variables are maintained, in order to study a particular phenomenon.

²⁰⁷ A prototype of this interface was also the subject of a paper I presented at the DiXiT *Scholarly Editions as Interfaces* symposium in Graz in September 2016. The arguments presented in this chapter draw, to some extent, on those presented.

physical artefacts that it represents. As Schloen and Schloen argue, this document-centricity is prevalent in Digital Humanities, and, they argue, can be seen as closing down avenues of enquiry. (Schloen & Schloen 2014) The first aspect of this discussion is therefore an analysis — stemming from the conceptions outlined in earlier chapters — of the reasons behind this: in brief, I would suggest, it is a necessary requirement of reflexive transmission (self-validating transmission that creates a signifying system) when applied at the document level. Secondly, I will discuss an alternative paradigm, based on a previously explored concept, Deleuze and Guattari's notion of the rhizome. This, I would suggest, can be represented, albeit imperfectly, via graph- or network-like models — a form of what has been described as “nodalism”. (Gochenour 2011):

[I]t [a network-perspective] opens up new connections between and understanding of objects in the material world, and does so through processes of analysis and re-combination that display systemic relations.

(Gochenour 2011)

Similar ideas have been explored elsewhere. Ciula, Spence and Vieira's work on the *Henry III Fine Rolls Project* takes the digitisation of the textual content of the rolls and attempts to extract and derive associations between ‘factoids’ (identifiable named entities), which can be modelled as a graph (RDF in their case). (Ciula *et al.* 2008) Romanello explores citation networks in classical scholarship. (Romanello 2016) Elsewhere, Sukovic points to the rhizome as a mode of knowledge organisation that enables (to cite the title of the paper), “information discovery in ambiguous zones of research”. (Sukovic 2008) Robinson and Maguire likewise envisage a hyperlinked information system modelled rhizomatically (rather than hierarchically). (Robinson & Maguire 2010)

What distinguishes this approach is, I think, the extension of the rhizomatic conception to cover the transmission process as well as the output. In previous chapters, the entire process

of editing was described as a set of interactions between assemblages. This has no organising logic, merely a set of connections and interactions, and may thus be considered a rhizomatic system. However, the gap in space and time between the editing stage and the user-interaction stage causes a *signifying* rupture:²⁰⁸ it results in the delimitation of the output, simply by its disconnection from its means of production. I will firstly discuss the result of this, before exploring how it might be represented.

The single-letters-only restriction referred to above (that letters can only be treated as discrete, extensive *collecteds*, not a *collection*) can be seen as a direct result of being reflexively transmitted. As Chapter Two argued, reflexive transmission involves both a complete delimitation of the source-as-assemblage, making it an extensive entity, and thereby the internalisation of heterogeneous, rhizomatic connections. This is the only way it can be transmitted as a self-contained and self-validating representation. It necessarily produces hierarchical systems, with connections subsumed as *properties* of the root (in this case, the single letter).

This is particularly problematic for properties that are not readily ‘internalisable’: that is, which should be considered entities in their own right. The most obvious examples in the case of *Letters* (and letters in general) are the sender and recipient. They are, as Latour would suggest, actors in a network; not properties of a letter. Of course, a letter can ‘contain’ authorship or receivership as attributes (this is the extensive view). However, it is not the only view, and not one that, I would suggest, should be given pre-eminence. This is especially true

²⁰⁸ As Deleuze and Guattari state, one of the properties of a rhizome is *asignifying rupture* (see Chapter 3) — the breaking of any connection without destroying the rhizome (and thus creating a signifying system). However, sending the output out into the world completely destroys the rhizomatic system of its creation, resulting in a delimited, *signifying* system — even though we cannot say what it signifies (insofar as it represents something else).

when considering a collection of documents, some of which are authored by the same person.²⁰⁹ By internalising authorship of each discrete document as a property or attribute is to say, conceptually at least, that it does not matter that the same person was responsible for another document (and, therefore, the documents have no connection to each other). This might be fine when the object being transmitted is a literary work — it matters little to *The Brothers Karamazov* as a novel that Dostoevsky also wrote *The Idiot* — but for a multiple historical documents (especially letters) it is essentially dehistoricising. In the case of a letter, a communicative act between two people, who wrote it and who received it matter immensely — and even more so because they may have written more letters to each other, or to someone else. What this requires in terms of transmission is avoiding the rupture of a rhizomatic assemblage that is caused by extracting only the letters and internalising the people, and instead maintaining the rhizomatic nature of the assemblage by transmitting the letters *and* the people as distinct entities with external connections.

We see this in some cases with *Letters of 1916*. Firstly, the names of sender and recipient are normalised in a list of people. Moreover, the abstracts written for each letter can be seen as comprising three distinct parts: a description of the letter, followed by brief biographies of the sender and recipient (indeed, it is typical in the project workflow to reuse the biographical sections across multiple letters). For the consistency of the abstract's prose, it is useful to have the biographies 'tailored' to fit with the letter description — but, properly considered, biographies are really properties of the sender and recipient, not of the letter:

²⁰⁹ In the case of a scholarly edition of a single, standalone work, this distinction matters little: it is not a difficult piece of mental gymnastics to flip from authorship as a property of the work to viewing authorship as a connection to an external actor. We might wish, as literary critics, to externalise (or, in Barthes's case, completely marginalise) the author to detach meaning from authorship — but this is based on the capacity of a literary work to stand alone, in a way that a piece of correspondence from one person to another cannot.

again a case (and not just for avoiding repetitious work) for considering the sender and recipient as external to the letter.²¹⁰

The same case can be made for treating the repository from which the letter was selected as a separate, external entity. In one sense — and this is especially true when dealing with any single letter — it does not matter where the original document is held; legal acknowledgements aside, while it is certainly nice to know where the original letter is held, the letter itself is not affected by happening to be in the *National Archives of Ireland* or the *Irish Jesuit Archives*. Neither, at the level of the letter itself, does it really matter how it got there. Whether it is treated as an external connection or internal attribute is, again, unimportant in the singular case. But, when considering multiple letters as a collection, where each letter comes from is particularly important, as it fundamentally shapes the entire collection.

In this case, the repository itself is — and, importantly, functions as — an external and independent actor. Typically, we are even less likely to see the repository as an external factor than the sender and recipient, as it ostensibly plays little part in the ‘production’ of the letter as a historic artefact (beyond, of course, its survival: but then, if it did not survive, it would not be being digitised...). A further case for it being considered external is the pragmatic connection between particular archives and individuals,²¹¹ which themselves, along with the letters, form a complex, rhizomatic assemblage. To the collection as a whole, therefore, selection of a repository selects letters in turn; but it also selects senders and recipients. By

²¹⁰ As for the section of the abstract describing the letter itself, this, too, could be considered as external to the letter. In this suggestion there are shades of Gabler’s desire to see critical commentaries and other components of a scholarly edition elevated to a status equal to that of the edited text. (Gabler 2012)

²¹¹ See the previous chapter for discussion on this point.

internalising these variables as properties of individual letters, such connections are masked from the user.

A final aspect, which is naturally taken to be an internal property, is the presence of a full-text search term in the document. To begin with, it is important to recognise that the text of a document is not, necessarily, an internal property of that document. Considering the document as an intensive assemblage suggests only a *connection* between the text-as-language and the material surface. (In Chapter 2, I suggested that Sahle's 'Wheel of Text' could be turned inside out, making 'text' not a set of dimensions but connections between heterogeneous entities.) Indeed, the separation of linguistic text (or any inscribed sign) from the material document is predicated on it being considered an intensive, rhizomatic assemblage. It is, again, the transmission of the source-as-assemblage that internalises the linguistic codes, making them hierarchically subordinate to the extensive document.²¹² Since — having been formalised for transmission — there is evidently a one-to-one connection between the document and its text, it again matters little whether we consider it an internal property or externally connected. Letters are relatively physically discrete objects, and thus 'naturally' and unproblematically internalise the text. (In the same way, it does not seem strange to have an envelope hierarchically subsumed into the letter: it does not typically have an independent life outside this letter-context.)

However, this is not true to the same extent with the presence of a search-term. Firstly, it is obvious that a full-text search operates across the collection, not individually on discrete letters. Moreover, though it appears to connect letters on the basis of an internal property

²¹² This is not a necessity. A letter, for example, could be conceived of as a material entity connected to a particular linear path through an external, abstract lexicon. Indeed, Dekker's TAG system of encoding (see Chapter Two) operates on just this kind of principle. Needless to say, it does not make for easy reading.

(the presence or otherwise of the searched-for text string), this is only one way of considering the matter. Insofar as a full-text search *connects* letters (by, typically, returning the letters containing the searched-for string in a list) it does so by considering the search term as an external component: the search term, as a discrete entity, ‘points to’ the relevant letters and forms an external connection. (In this respect, we are treating the text of the letter as potentially ‘externalisable’.) Considering text, and especially a user-input search term, in this way is a difficult conception, and challenging to argue for. As the previous chapter suggested, though, full-text search amounts to a complete reconfiguration of the collection — insofar as it stratifies it into ‘relevant’ and ‘irrelevant’, and imposes some kind of new ordering on the documents — on a basis that is at least somewhat arbitrary. Therefore, considering it as an external component connecting documents, rather than a reconfiguration based on an internal property, has some merit, not least because a search-term itself may point to or correlate with other externally-connected entities.

The above has identified some properties of letters that, I think, have the greatest effect on the collection as a whole; or, to put it another way, where most information is masked by the requirements of a rigidly document-centric perspective. Of course, there are other properties that may be externalised: places and dates, most obviously. Indeed, any property of a document could be considered external and modelled as a discrete entity. This is what considering the document as an intensive assemblage entails, after all. The question is, what properties is it useful to treat in this way?

These ‘externalisable’ properties fall into two categories. The first are historical actors contemporaneous to the letters themselves (i.e. the sender and recipient). Internalising these as properties of discrete letters serves to hide historical connections, and thus decontextualises the letters. The second category (repositories and search-terms) are more to do with the

production of the collection. They lie at the root of the most significant latent contingencies in the determination of the collection. By internalising them, the nature of the collection itself is hidden. Modelling them as external entities, however, shows the connections between letters and their archival origins.

The approach suggested here by the prototype *LetterGraph* system is, therefore, in part an *articulation* of the process of creation as a whole, by treating such entities as independent and modelling the rhizomatic connections between them. Or, as Deleuze and Guattari put it, making a map and not a tracing: mapping the connections that exist, rather than subsuming them into a totalising, pre-determined model. As has been seen above, at the level of the individual letter, this kind of internalising model has a tendency to brush contingencies under the carpet. And, at the level of the collection as a whole, there can be no totalising model that accounts for the whole (see the previous chapter). *Letters* is, in this sense, really a rhizomatic system.

The result of mapping the collection in this way is not, therefore, a global account of the whole. It is not a reflexive transmission, where subjective decisions point back to a regulative idea and create a signifying system. Rather, it is a non-signifying system. “We don’t know what [it] entails when it is no longer attributed, that is, after it has been elevated to the status of a substantive.” (Deleuze & Guattari 1987, 4) As such, it functions more as a kind of holistic articulation — an articulation of the building-out to the boundary that is formed simply by calling a halt at a given point, rather than a specific notion of what that boundary represents. It does not represent anything.

6.3. Rhizomes as graphs

The principle problem with modelling a rhizomatic system is where to stop. Not just externally — the *Letters of 1916* has a boundary, albeit a non-signifying one — but internally. Where do we draw a line between making everything a large, heterogeneous assemblage — to the extent that there are no points at all, only pure connection — and where do we model components as molar extensities? To even talk of ‘modelling’ suggests some components are extensive. This is antithetical to a purely rhizomatic system. Considering a literal rhizome *par excellence*, ginger, suggests that there are no points: there are not multiple gingers, connected by some kind of ginger-tendrils, but only one ginger, the whole system. New shoots emerge, but not to encapsulate anything, rather to extend the rhizome. Thus, any kind of rhizomatic map we try to draw is going to have some points on it, and from these points emerge small tree-like structures. To take the most obvious example, individual letters, we may connect them heterogeneously with other entities, but they form miniature hierarchical systems that subsume the material document and the text. We could break apart the letter into its molecular components, but then modelling them becomes impossible without reconstituting the molecular components as molar entities (e.g. individual words, which then have to be connected together, somehow)²¹³. It is not that this cannot be done, but a pragmatic question of where it is useful. As Deleuze and Guattari observe, we cannot really avoid tree systems — “Is it not of the essence of the rhizome to intersect roots and sometimes merge with them?” (Deleuze & Guattari 1987, 13) — but it is a question of “method”, of connecting tree-entities

²¹³ Not only is this *connecting back together* the imposition of a model — which naturally tends in the case of text towards hierarchical conception — it is, I think, largely redundant in the case of letters: they are singular, discrete and, above all, unique objects. Unlike a text that exists in several manuscript versions, there is no benefit to externalising the text. It is, of course, difficult to talk of these issues in the abstract, and I will return to this point later in this chapter.

where we can: “the tracing should always be put back on the map” (Deleuze & Guattari 1987, 13).

Furthermore, it might be questioned whether modelling the whole system as a rhizome is, not, after all, the imposition of a kind of model. (“Does not a map contain phenomena of redundancy that are already like tracings of its own?”: Deleuze & Guattari 1987, 13) This is indeed the case; any kind of model assumes, pre-emptively as it were, the model over the system. In this case, we are making subjective decisions about what to connect (senders and recipients to letters; letters to repositories) — in other words, prioritising some connections over others according to what we (the modeller) consider important. But, as Deleuze and Guattari suggest, the order in which this is done is important. (Deleuze & Guattari 1987, 14) We are not beginning with a pre-determined structure, say a collection ‘containing’ a set of repositories, containing a set of letters, which restricts movement hierarchically (or even a set of overlapping, but nonetheless independently hierarchical, structures). Rather, the pre-emptive ‘model’ (really it is a kind of non-model) is the *map*, and it becomes a kind of tracing by limited reduction, rather than complete reconfiguration.

If a distinction may be drawn between a relatively hierarchical or relatively rhizomatic model, it is the number of dimensions involved: between a hierarchy that constantly adds new dimensions, or a flat ontology that proceeds by flattening dimensions to a plane. The conception here, therefore, is (broadly speaking) of a predominately flat ontology, connecting small hierarchical systems; as opposed to a predominately hierarchical system that may contain horizontal connections. We have already discussed the relationship between sender, letter and recipient, but this point can be illustrated by considering its encoding or representation. In the *Letters* TEI-XML documents (evidently hierarchical), the names of sender and recipient are encoded (in the <correspDesc> element) using unique identifiers

that are conveniently stored in a separate file. The hierarchical system is inherently document-centric — there is one TEI file per letter — though there are horizontal connections between the sender or recipient of one letter and those of another by virtue of normalisation. A more rhizomatically-inclined approach, such as that proposed here, would place the letter and the persons on the same footing — as separated, though connected, actors on a single plain (what Deleuze and Guattari refer to as a “plane of consistency”).

In picturing this rhizomatic system in such a way, it is easy to see how what is being suggested is really a kind of network or graph, at least between the elements that we have placed on one plain. Above (or below — in supplementary dimension) the network-plain, in particular in the case of letters, is the subordinated content of that letter. These local hierarchies will be considered later.

A network, or graph (the term I shall use here, though they are essentially interchangeable in this instance) comprises a set of *nodes* (or vertices), and a set of *edges* designated by two nodes that are joined by that edge.²¹⁴ To take Ruohonen’s formal definition:

Formally, a graph is a pair of sets (V, E) , where V is the set of vertices and E is the set of edges, formed by pairs of vertices. E is a multiset, in other words, its elements can occur more than once so that every element has a multiplicity. (Ruohonen 2013, 1)

²¹⁴ In writing this section, I am particularly indebted to Scott Weingart’s (the Digital Humanities Specialist at Carnegie Mellon University) series of blog-posts, *Networks Demystified* (<http://www.scottbot.net/HIAL/index.html@tag=networks-demystified.html>). This excellent series of posts not only ‘demystifies’ the mathematics of dealing with graphs (without brushing it under the carpet), but uses examples that are directly relevant to digital humanities.

In this instance, then, letters, people and repositories can be connected and considered as a network; each entity is a network-node, and the connections — of various types — are edges. However, the connections in this case are not necessarily a simple edge, but possibly an *arc* or *directed edge* (Ruohonen 2013, 27); and thus the graph is a *directed* graph. A directed edge is non-reversible. (The edges are therefore considered an *ordered pair*, i.e. *from* the first node *to* the second node). In the *LettersGraph*, each edge is taken (as in RDF: see below) to represent some kind of logical predicate; these may or may not be reversible. As an example, the predicate “is a sibling of” is clearly reversible (by definition, X being a sibling of Y means that Y is a sibling of X), while “is the parent of” is non-reversible. In this conception of the *LettersGraph*, edges are taken to be directed, representing non-reversible predicates. This is, indeed, fairly obvious in all the cases considered here (a letter cannot be the ‘sender’ of the person, whatever that might mean). This is sufficient information to imply the reciprocal form, thus: “*Person X sent Letter Y to Recipient Z*” also implies “The sender of *Letter Y* is *Person X*”, and so forth. Given the very distinct types of nodes here, it also suggests that no two nodes of the same type are ever directly connected. (It might be true that *Person X* and *Person Y* were friends, but in our graph their connection is only ever via a letter.)

The *modality* of a graph describes the number of types of node contained in the graph. As Weingart suggests — and to extrapolate from its absence in Ruohonen’s very formal and mathematically-inclined book (Ruohonen 2013) — mathematicians and computer scientists most typically deal with unimodal graphs (one type of thing). By contrast, humanities disciplines often give rise to bimodal or even higher modality graphs. (Weingart gives the examples of books and publishers: bimodal.) For this *LettersGraph*, it is clear we have a potentially very high modality: persons, letters, repositories, and (eventually) search-terms. (Persons, incidentally, do not need to be divided into senders and recipients as this is

determined by the edges; nor is it useful to do so, as the same person can be the sender of one letter and the recipient of another.)

High-modality graphs are problematic from a computational perspective, as most algorithms are designed to work with uni- or at most bimodal graphs (the norm in computer science). Thus typical metrics, such as *degree centrality*, which identifies the most important nodes in a network by their being most connected, become increasingly meaningless. (What would it mean, for instance, for a repository, or a particular person, or particular letter, to be ‘most important’? A repository ‘more important’ than a letter?) Even more than this, such metrics are rather contrary to the purpose of this approach. Assuming a graph was constructed with just the people (the edges themselves designate the sending of a letter), the most important person would be the one whose letters — and most likely the letters they received, rather than sent — made it into the *Letters* collection. Thus it would tell us, in a highly reductive manner, only something about the sourcing of the collection materials.

The approach take here, therefore, is more qualitative, specifically using network visualisation as a way to ‘discover’ letters, within the overarching framework of both the historical networks they comprise (or instantiate) and the ‘networking’ carried out by the collection’s creation. That is to say, it is really a tool for knowledge organisation, rather than a computational search for significance. Moreover, it represents a particular kind of knowledge organisation, not bound by the restrictions of category-based lists (including date ranges and full-text searches).

6.3.1. Graph visualisation

To say graph visualisation is a ‘big deal’ in Digital Humanities is something of an understatement.²¹⁵ A panel dedicated to *Visualising Futures of Networks in Digital Humanities Research* at the 2017 *Digital Humanities* conference is a testament to this. Indeed, as the abstract for this panel notes, some ten percent of papers submitted to DH2017 were on the subject of — or utilised the methodologies of — network analysis and visualisation. (Kaufman *et al.* 2017) As Kaufman *et al.* observe, there generally exist two kinds of ‘network analysis’ in Digital Humanities: the representation of ‘real world’ networks, either historical (of which *Letters* is an example) or current (social media friendship networks, for instance); and the ‘extraction’ of networks from other sources. Of the latter, networks of characters in novels is an obvious example, and represents the kind of approach engendered by, in particular, Moretti (see Moretti, *Distant Reading*, 2013 and *Graphs, Maps and Trees*, 2005). This kind of network extends externally to connections between works or documents, and into computationally-generated networks such as topic models (which connect documents to computed topics by the importance of that topic in the document). Of the former — real-world network modelling — works such as that by Winterer (2012), who maps the correspondence network of Benjamin Franklin, are good examples. The benefits of conceptualising, and especially visualising, these entities as networks are enormous: it allows scholars to engage with complex relations in a way that is not possible from a document-centric perspective, or from any other form of data.

Even then, we can see two roughly divergent approaches. One, encapsulated most clearly by Jockers, is what might be described as network visualisations as *tests*. As an example, Jockers attempts to identify ‘influence’ in literary creation by, firstly, calculating a number of

²¹⁵ When the *logo* of a conference (Digital Humanities 2014) is a citation network, something is afoot!

distance-measures between literary works, and then visualising the results as a graph (colour-coding the nodes by date). The validity of the hypothesis being tested in this case — whether there was an identifiable notion of influence — was determined by the proximity of contemporaneous work and the rough chronological ordering of the nodes in the network from left to right. (See Jockers 2013, 154–68.) Whether or not the methodology used in calculating the distance between works can be seen as a valid proxy for ‘influence’, the point stands. *Assuming* the validity of the methodology, the network visualisation itself — like plotting the results of scientific experiment on a graph to see whether relationships emerge — can be seen as a *test* of a given hypothesis. It is a largely empirical approach.²¹⁶

The other approach, adopted by tools such as *Voyant Tools* (Sinclair & Rockwell 2017), but also more generally in humanities network visualisation, is a kind of ‘deformance’ (to borrow McGann’s term; see McGann 2001, 145). In my presentation of the *LettersGraph* prototype in Graz (see above), I described this as “serendipitous shuffling”: the idea that rearranging data in some subjective manner can stimulate critical thought by showing certain elements, which might otherwise be marginalised, in a new light. Into this category we may place such approaches as topic modelling. These can be seen as a more refined way of throwing everything in the air and seeing where it lands. Though it may be more refined, this does not make any particular method ‘most objective’ in any given case. As Sculley and Pasanek argue, the so-called ‘No Free Lunch Theorem’ (Wolpert & Macready 1995) means that there is no universally best algorithm for manipulating a given dataset; and, moreover, neither can there be a “*meta-learner*” that can determine this on the basis of the data alone. (Sculley & Pasanek 2008, 413) This is not, of course, to invalidate such approaches; rather, to assert that they

²¹⁶ Bode (2017) is particularly critical of this approach’s claim to a kind scientific empiricism.

make no claim to objective knowledge, acting instead as a spur to subjective reconsideration of the evidence, and then (possibly) to new knowledge.

What unites all the above approaches, however, is their position within a typical Digital Humanities workflow. They are typically associated with what might be called *post hoc* analysis. That is to say, they mark the culmination of work that began with collecting, and possibly digitising, artefacts (in this case, let us stick with text), the cleaning and encoding of data, and so forth. In this case, the product of the work carried out *is* the analysis; the data is “merely a byproduct” (as Faulhaber, 1991, puts it; see Chapter 1). In this regard, it is distinct from scholarly editing or even the digitisation of artefacts for their own sake: in these cases, the reproduced works or texts or documents are the product.²¹⁷ The user might take the rigorously-established texts from one or more editions and use them for computational analysis, including specifically the construction and visualisation of networks, but the possibility of doing this is really tangential to the interests in the editor in creating the edition.²¹⁸ (One possible exception is variant graphs, but I do not know of a digital scholarly edition that employs them as a user interface.)

What I am proposing here is, rather, that these two stages be merged. Or, put another way, that visualisation need not be the preserve of post-publication research into the edition’s text, but a just as legitimate — if not better — way of presenting the edition’s text(s). Or even, to take it a step further, that the visualisation is in part the ‘text of the edition’. This claim is, I

²¹⁷ If the analysis were the *product* (the text itself being merely a byproduct), the ‘edition’ (if we could even call it that) would comprise nothing but analysis, typically in the form of critical commentaries.

²¹⁸ Not, of course, that *no thought whatsoever* is given to potential future uses. The creation of TEI documents, especially, is a means to foster data reuse in other cases. The creation of a digital, rather than print, edition is also a nod towards facilitating reuse. My point is, rather, that it is not the editor’s responsibility to pre-emptively carry out such work.

think, especially valid in the case of *Letters*, especially considered as a rhizomatic assemblage. The 'text' is not simply the text internal to the letters (what is written on the page) but the connections between actors (re)instantiated by those letters. If we create a visualisation of the collection, therefore, the act of 'exploring the data' becomes at least partially synonymous with 'reading the text of the letters'. By the same token, the connections we might draw between letters and their repositories can be 'read' as a kind of apparatus. It can also represent (to a degree at least) the forms of commentary that, in print editions, could only be done in prose.

As such, the idea is to build a system to enable reading of the letters (which is half the objective of a text collection anyway), while at the same time 'reading' the collection as a whole (or, at least, reading the letters in the context of the collection as a whole). In fact, it could be said that these two operations are really mutually sustaining: to read in the visualisation the contingencies in the letter-texts (most notably, *which letter-texts?*) and via the letter-texts to read the contingencies in the visualisation. What is therefore required is a constructive mode of 'flitting between' these two types of reading.

One approach is to provide, on one hand, a static network visualisation of the *Letters* network (with the nodes labelled) and, on the other, the texts of the documents. Thus the visualisation can be read in concert with the letters themselves, the former function rather like an index (or index combined with contents page) in a book. For *post hoc* visualisation work, where one wishes to further investigate certain phenomena that emerge in a visualisation, this is exactly what must be done. However, interactive visualisation can combine the *index* and the *reading* properties in the same system. It is possible to 'flit between' the network-scale interface, and the local, letter-scale reading environment. This, I would argue, is an edition-like

representation; it is the primary interface for displaying and accessing the collection: not an adjunct to an edition, but integral to it.

6.4. Building a network interface

6.4.1. Preliminary considerations

Thus far in this chapter, I have articulated a rationale for, firstly, a scholarly edition to be conceived of as a rhizome; secondly, for a graph model to be a suitable approximation; and, thirdly, for any proposed system to allow both exploration of a graph and the reading of its constituent texts. In this section, I will discuss the preliminary considerations that underpinned the building of the *LettersGraph* prototype. More than anything, they comprise a set of desiderata.

If such a network-interface system is to function as a scholarly edition, it must (I would strongly argue) function as a relatively free-standing (digital) entity. This is, as much as anything, a pragmatic consideration, but one which is not without precedence in the history of scholarly editions, including their digital embodiments. In the most basic sense, a print scholarly edition does not need some extra, bolted-on infrastructure to function — at least, over and above any other book. As discussed in Chapter One of this thesis, after a brief flirtation with CD-ROM, the ‘digital scholarly editing community’ (if it can be reduced to such a generality) settled on the Web as the most suitable platform for dissemination. Of course, it requires a computer and a web browser; but it does not require anything over and above any other website. What I am arguing for is, therefore, a low-threshold and functioning-out-of-the-box approach. Just as a digital scholarly edition is not a set of transcriptions and the instruction to “Download CollateX and load them up”, I do not think the *Letters* collection as GraphML that relies on the user’s ability to use Gephi would count

as a scholarly edition. As I said, this is largely a pragmatic decision (there is no theoretical basis for saying a scholarly edition must use ‘Digital Format X’ any more than there is for arguing that a scholarly edition must be a book).

Therefore, one design criterion may be set: it must be Web-based. True, it is now possible with various high-level languages to produce platform-independent standalone applications.²¹⁹ However, these are beyond my expertise. Further, with modern server-side web frameworks, myriad excellent (if not superior to their standalone counterparts) front-end libraries, coupled with optimised JavaScript engines, and better Web standards (including for graphics), there are few benefits to doing so any more.

It must also — obviously — present a visual representation of the *Letters* collection as a network. It should not be, for instance, a SPARQL endpoint for querying the collection encoded in RDF. Moreover, neither should the ‘connections’ between letters, persons and other entities be HTML pages connected by hyperlinks. This latter fails to solve many of the problems hitherto discussed. While it makes the collection no longer obviously document-centric, it makes it instead entity-centric (that is, there is a page per entity). Furthermore, although it allows the user to move through the collection in a non-linear way, it masks not only the structure of the whole network, but anything that is more than one click away. Rather, it should show the entire collection, as both a way to contextualise letters in relation to each other, and to make visible to the user the ‘implicit biases’ in the collection.

The interface should also be interactive, dynamic, and reconfigurable by users. As the above suggests, putting everything into a single visualisation risks information overload, or at least

²¹⁹ For instance, there are several Python libraries for building GUI applications; mention should also be given to the (immensely cool) *Ruby Shoes* project.

creates difficulties for navigation. Therefore, a user should be able to filter the graph or highlight connections in which they are interested. It should also — this is imperative — allow direct access to the text of each letter, within or as close as possible to the visualisation of the network, ideally by clicking on it. A user should not have to identify a letter, then navigate to a different page to access the text of the letter.

Most importantly for a scholarly edition, it must be citable. This, as the later section on implementation further argues, is problematic. Typically, one cites a digital reproduction by providing a URI (Uniform Resource Identifier). But this assumes a direct mapping between the page at a given address and the represented document. With regards to the letter-text, this can indeed be the case (one letter, one URI). However, a purely rhizomatic system really cannot have any citability, as there are no *points* to cite. Even a network graph as an approximation — when visualised as a whole — has no fixed positions: one may zoom in on a particular section of the network, but this is not explicitly identifiable by reference to a particular letter or person or other entity. Finding a solution for this is, perhaps, the most conceptually difficult challenge to solve.

As a final point, the system conceived here should take as its input the TEI documents produced by the *Letters of 1916* project. Intermittent data formats are, naturally, required (data must be extracted from the TEI documents, after all) but should not require preliminary conversion. The advantage of starting from the TEI documents is the ease with which the text of the letter and its encoding can be transformed for output as HTML

6.5. Graphs, data models and data storage

There are many ways to digitally represent and encode graph data. TEI P5, indeed, has a module for describing graph structures.²²⁰ However, while this is potentially useful, it would involve re-encoding all the documents, or a redundant intermediary transformation as the format is not widely used by any software libraries. DOT, GML, GraphML and GEFX are more widely supported.²²¹ Moreover, the graph structure, while not explicitly encoded as such, is already implicit in the TEI documents. The use of the <correspDesc> element, with @ref attributes for the sender and recipient, provides what is essentially a graph-like model, that simply needs extracting as a list of nodes and edges. The name of the repository is also encoded and normalised, so can likewise be transformed into a node (and an edge, pointing to the letter in question).

Another possibility for encoding a graph is the use of RDF (Resource Description Framework), the principle framework for building the so-called ‘Semantic Web’. This has many advantages, being ubiquitous and, moreover, allowing the connection of our graph to external resources through existing ontologies (which provide unique URIs for every entity — where, of course, they exist). While certainly powerful, it has several drawbacks in this instance. Firstly, the basic model is *entirely* graph-like: everything is modelled as a subject—predicate—object *triple*. There is no simple mechanism for attaching an attribute to either a node (a *resource*) or an edge (a *predicate*); even a person’s name must be considered as an entity (a text string) that is connected to the person by a predicate (e.g. ‘has name’). While

²²⁰ <http://www.tei-c.org/release/doc/tei-p5-doc/en/html/ref-graph.html>

²²¹ Both DOT and GML are text-based formats, while GraphML and GEFX are XML-based, making the latter two more useful as target formats for TEI transformation. There is, broadly speaking, little to choose between the two: GraphML is a recognised standard, while GEFX was created by the makers of Gephi (though designed for exchangeability). The choice, given that the graph will be generated automatically from the source TEI documents, is therefore dictated entirely by the graph library selected (see the next section).

this aids in querying the data (there is never a question of whether something is an attribute or an externally-connected node: it is always the latter), this reduces the ease with which the graph may be manipulated. Secondly, while it is suitable for complex relation-based querying, turning the results of those queries into a useable format is another layer of transformation. Thirdly, the use-case for RDF is, at best, tangential to that proposed here.²²²

A final consideration is to store data using a database that natively supports graphs. Of these, the most commonly used appears to be Neo4j. Its particular strength, especially versus relational or other varieties of NoSQL databases, is its optimisation of graph-traversing queries (using its inbuilt query language, Cypher). However, this project does not require arbitrary querying of the graph. Moreover, a further layer of software will be required to render the graph visualisation and perform degrees of filtering, which makes such a solution somewhat redundant. Most significantly, while it allows an unlimited number of attributes on both nodes and edges, these are constrained to primitive types: the complex XML encoding of letter texts cannot be stored (except insofar as XML is a text string), which greatly lessens its potential.

The problem, therefore, can be seen as identical to that identified in the previous section, though in technological terms: how to combine network-level graph data and the hierarchical TEI model of the letter-texts. I therefore took the decision to retain the TEI documents themselves as the primary data to be stored, and settled on eXistDB as the most appropriate database software.²²³ The advantage of this is twofold. Firstly, extracting data as a graph from

²²² Frank Lynam (as part of his PhD project at Trinity College Dublin) did create an experimental RDF SPARQL endpoint for *Letters*, which was excellent for mining relations — but not oriented towards visualising them, or, more pertinently, reading the text.

²²³ This is, on the whole, an entirely subjective decision based on my personal familiarity with the platform (and acquaintance with eXist developers — always useful). Undoubtedly, other XML-native databases such as BaseX, could do the same job.

a set of TEI documents is relatively trivial using XQuery or XSLT — especially when the target format is itself XML-based. This can be compared with the much greater complexity of doing the reverse, i.e. modelling hierarchical TEI mark-up as a graph.²²⁴ Maintaining the letter-text as XML, moreover, allows its transformation into HTML to be carried out easily using standard approaches (again, XQuery and TEI). Furthermore, eXist-DB has inbuilt support and optimisation for full-text queries (including on specific elements of the TEI-encoded data). Using eXist-DB in conjunction with a graph library (which will ingest the graph model in an XML format) seemed a logical approach.

As the graph library chosen most readily supports the GraphML format (see the following section), this was chosen as a target format for extracting the graph data from the TEI documents. The relevant queries are carried out using stored XQueries in eXist.

6.6. Building the system architecture

As with the choice of data storage, the other components of the prototype *LettersGraph* system are largely pragmatic. Certainly, there is no single, off-the-shelf solution, and no naturally ideal approach; everything is, in some measure, a trade-off. The decisions amount to a choice of server-side language, graph library and web-framework, client-side graph rendering library, and, most significantly, at which point certain operations (graph querying, calculating layouts) should be carried out. The choice of Python for a server-side language was based entirely on my familiarity with the language and the (relative) simplicity and

²²⁴ As in, for example, Dekker's TAG model. The great advantage of this approach — multiple logical hierarchies — is irrelevant when dealing with data that has already been encoded as XML (the hierarchy problem has been resolved one way or the other).

agnosticism of the web-frameworks written in it.²²⁵ There are also a number of graph libraries written in the language, which will be discussed below.

The most difficult task for the system is the visualisation of the graph data. This can, more properly, be considered as two tasks: firstly, the calculation of the graph layout, and secondly the ‘visualisation proper’, the rendering of the graph in an interactive format on a web page. I will deal with the first of these here.

A visual representation of a graph can take many forms. All involve the flattening of the network onto a two-dimensional plane. Force-based visualisations are most common in representing the kind of networks seen here,²²⁶ and are default in software tools such as Gephi. The ForceAtlas algorithm is “focused on quality (meaning “being useful to explore real data”) to allow a rigorous interpretation of the graph [...] with the fewest biases possible.” (Gephi documentation²²⁷), and is therefore the algorithm used here. Other layouts, such as Circular and Radial Axis, correspond, on the whole, to different use-cases and data (particularly unimodal graphs; a Circular layout makes little sense for multimodal graphs).

²²⁵ Other potential options ranged relatively Node/JavaScript frameworks such as Express.js to more highly opinionated and more complex (and probably overkill) options such as Ruby on Rails (whose main benefit, close integration with relational database models, is redundant here). ‘Agnostic’ and ‘opinionated’ are commonly used to describe web frameworks: an opinionated framework has built-in default way of doing things, which makes it easy if your use-case matches the default. Agnostic frameworks, by comparison, provide a more bare-bones set of tools, which are, accordingly, easier to adapt to non-standard use-cases.

²²⁶ I have taken the work of Martin Grandjean, who specialises in humanities network analysis and visualisation, as something of a guide. (For a list of publications consulted, see <http://www.martingrandjean.ch/communications/>) In general, there is no ‘correct’ way to produce a visual representation of a graph: all emphasise certain features, such as the spacing out of communities, or connections over proximity.

²²⁷ <https://gephi.org/tutorials/gephi-tutorial-layouts.pdf>

Most force-based methods, including the ForceAtlas algorithm, are iterative processes, and computationally expensive, especially at the kind of scale required here (some 1600 nodes and 3000 edges). The practical decision of how and where (server or client side) to carry out the graph layout is therefore determined primarily by the speed at which they can be made to run.²²⁸ From a user perspective (though they would, hopefully, be more forgiving of an academic project than, say, Amazon) a loading time in the order of tens of seconds to a minute is problematic. The two classes of available option are: to calculate the layout on the server and render on the client, or to do both on the client side. Having tested the layout-calculating algorithms of various client-side libraries, none was found to be faster than around 40 seconds (running on my machine). I therefore decided to calculate the graph layout on the server, before pushing the results to the client-side to be visually rendered.

Of the various server-side libraries (written in Python), I experimented with three. NetworkX is, perhaps, the most widely used, and certainly the simplest. However, being written purely in Python, its performance in laying out large graphs is poor (towards a minute in repeated tests). iGraph (a library for both Python and R) was tested, but (probably due to my own failure) could not be made to import the GraphML data extracted by the eXistDB queries. The GraphTools library (which I eventually chose), has numerous advantages. Firstly, its layout algorithms are written in C++, making them faster to run than NetworkX (fifteen to twenty seconds). Secondly, it provides simple mechanisms for hashing a graph and storing it on the filesystem as a Pickle file;²²⁹ this allows graphs that have already been laid out by one user to be stored and loaded on further requests (taking two to three seconds).

²²⁸ While I did not perform any scientifically rigorous study, I tested all the options discussed in this section to see which ran at an acceptable speed (on my 2011 Dual-Core Intel i5 MacBook Pro with 16GB of RAM).

²²⁹ The hash provides a unique identifier for a particular graph; a Pickle file is a Python data object stored on the computer's filesystem.

Nevertheless, the requirement to spend at least twenty seconds laying out any previously unrequested graph is problematic for the design of the server. I had initially used the Flask web framework, but as a synchronous server (i.e. it can only handle one request at a time) it necessarily blocked simultaneous requests for long periods. Therefore, it was necessary to use an asynchronous web framework, for which I chose Tornado (its superficial similarity to Flask made it the obvious option), which allowed the running of the graph layout algorithm as a separate process, preventing requests from being blocked (including requests for cached graph layouts).²³⁰

²³⁰ This is especially important as not *all* user requests will be for a graph layout; simpler requests, for letter-texts, for example, are not blocked by more computationally intensive layout requests.

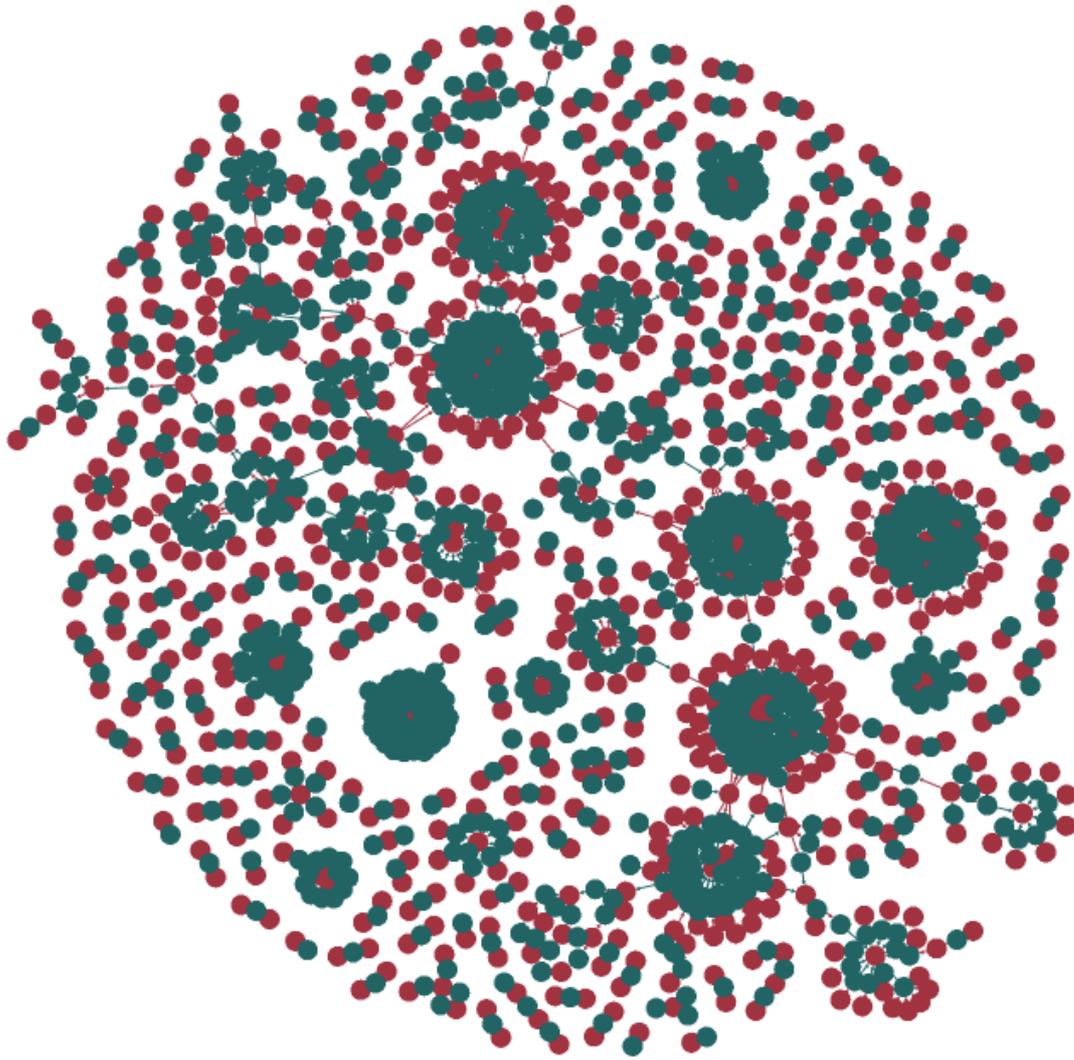


Figure 45: The default rendered graph (persons and letters). Letters are represented by blue nodes; persons by red. The sending of a letter is represented by a red edge, and its reception by a blue edge.

The client-side is designed as a single-page application. Rather than move from page to page (from one URL to another) by clicking hyperlinks, everything is rendered dynamically in a single page. This is necessary to maintain the user's context within the graph (i.e. the area they are looking at) while viewing a particular letter. The client-side functionality is, therefore, twofold: the visual rendering of graphs and the display of letter-texts, within the context of a single page. I will firstly deal with the graph rendering, which is undoubtedly the more complex problem to solve.

Having calculated the layout of the graph on the server, the layout and other graph data is sent to the client (via an Ajax request from the single page initially loaded). There are multiple browser-based graph rendering libraries. More general visualisation libraries with graph functionality include vis.js and d3.js. Other graph-specific options include Sigma.js, Springy, cytoscape.js and Vivagraph.²³¹ These libraries utilise, in general, one of three browser-based technologies to render graph visualisations (or provide multiple options). These come with specific advantages and disadvantages. SVG (Scalable Vector Graphics) is an XML syntax for describing vector graphics. In terms of browser-based interaction, it is excellent: each vector is a DOM object, and thus has support for event handlers (for user interaction events such as clicking). However — and as a direct result of this — it is notably slow for complex graphics such as a graph. Canvas, an HTML5 technology, by contrast, is bitmap-based (though based drawn by JavaScript-based vector descriptions); it has some level of interaction support, but is rendered as a bitmap, which limits this compared to SVG. WebGL (Web Graphics Library) represents the newest of these technologies; it uses the OpenGL standard for the drawing of two- or three-dimensional graphics (for this reason, it is commonly used for browser-based video games), and is, therefore, extremely fast compared to the other options. However, being a single entity in terms of the page's document object model (DOM), it lacks many interactive capabilities out-of-the-box.

Having experimented widely with a number of options, the two WebGL options (Vivagraph and Sigma.js) are clearly preferable — even compared to their own Canvas options (which are the default for both). Canvas-based graphs of this size suffered unacceptable performance issues, taking several seconds to redraw when scrolling around the graph. However,

²³¹ Anvaka, the creator of Vivagraph, has compiled a list and set of comparable demonstrations using these libraries (see <https://github.com/anvaka/graph-drawing-libraries>); this proved invaluable in deciding upon the best option.

interaction (especially clicking on nodes) remains a problem: Vivagraph’s support for events was inaccurate (it often failed to detect when the pointer was over a node) compared to Sigma.js, which also provides a much more intuitive, ‘native Javascript’-like API. For this reason it was chosen.

Upon receiving the graph data from the server, the graph library renders the graph (see Figure 45, above, for the default graph). Figure 46 shows a zoomed-in view of one letter, sent from Con Colbert to his sister, Lila.²³² This illustrates the general form of the graph: red edges indicate “the act of sending”, and blue edges “the act of receiving”; the arrows are intended as a visual indication of the direction of communication.

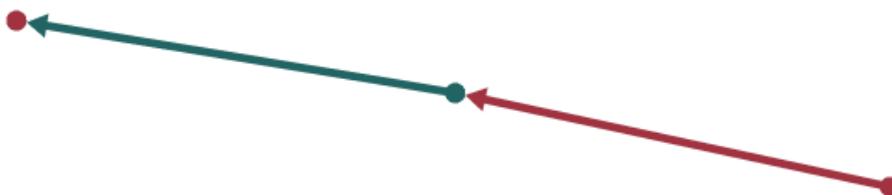


Figure 46: A zoomed-in view of the letter from Con Colbert to his sister, Lila. The red node on the right is Con (the sender, as designated by the red edge); the blue node in the centre is the letter itself; and the red node on the left is Lila Colbert (the blue edge represents receiving the letter). The arrows are intended to indicate the ‘direction’ of communication.

Each node can be clicked, to display the name of the individual or the title of the letter (these are taken from the <titleStmnt> <title> element in the letter’s TEI document). Figure 47 (below) shows the result.

²³² This example is an isolated set of nodes; no other letters sent by or to either person is in the collection.

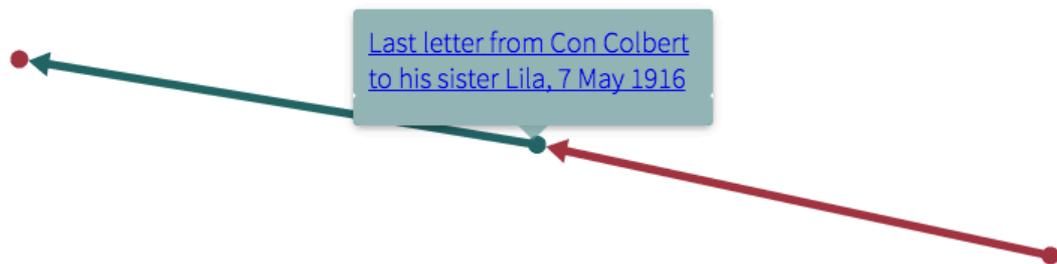


Figure 47: clicking the node reveals the letter title.

The operations involved in loading the graph data from the server and having it rendered by the Sigma library are relatively simple, and are achieved with a custom-built ‘framework’. There was no need, I felt, for a more powerful client-side framework, such as Angular, React, or Vue. The client-side code handles many operations. Manipulation of DOM elements uses JQuery. As a single-page application, further JavaScript is required to manipulate the browser address bar and history, and to make the ‘Back’ button function. Instead of returning to the previous page — as there is not one, by definition — clicking ‘Back’ triggers an event that reloads the previous state of the application.²³³

The client-side system also uses the (relatively recent) IndexedDB API, which allows data to be cached in a NoSQL-like document format in the browser (typically 50MB by default in most browsers).²³⁴ This allows both graph- and letter-data to be cached on the client and loaded locally when possible, rather than having to send a further (potentially lengthy) request to the server.

²³³ Historically, the “breaking of the back button” was seen as a major downside with Ajax-based applications; relatively recent developments, however, have provided JavaScript APIs to directly modify the browser history and address bar, allowing developers to re-implement this functionality. (See https://developer.mozilla.org/en-US/docs/Web/API/History_API

²³⁴ <https://developers.google.com/web/ilt/pwa/working-with-indexeddb> provides a detailed background of the API.

The other principal component of the client-side application is the interface for viewing letters. As discussed earlier, the switch between the graph view (and its inherent model) and linearly-read (hierarchically-encoded) text poses a particular problem. The challenge is to present the switch from graph- to letter-view in as seamless a manner as possible, without — permanently, at least — losing the context of the graph. The approach I decided upon was to overlay the letter-text, on top of the graph, which remains in position when the letter-view is closed. (See Figure 48, below.) The design of this view is intended to mirror the kind of network-relations present in the graph, borrowing explicitly the same colours and forms, with the letter node itself expanded to contain the letter text. Compared to other typical forms for displaying documents, it also intends to reflect the letter's status as a communicative act between two individuals: the letter-as-assemblage considered intensively, with the sender and recipient modelled as external relations.

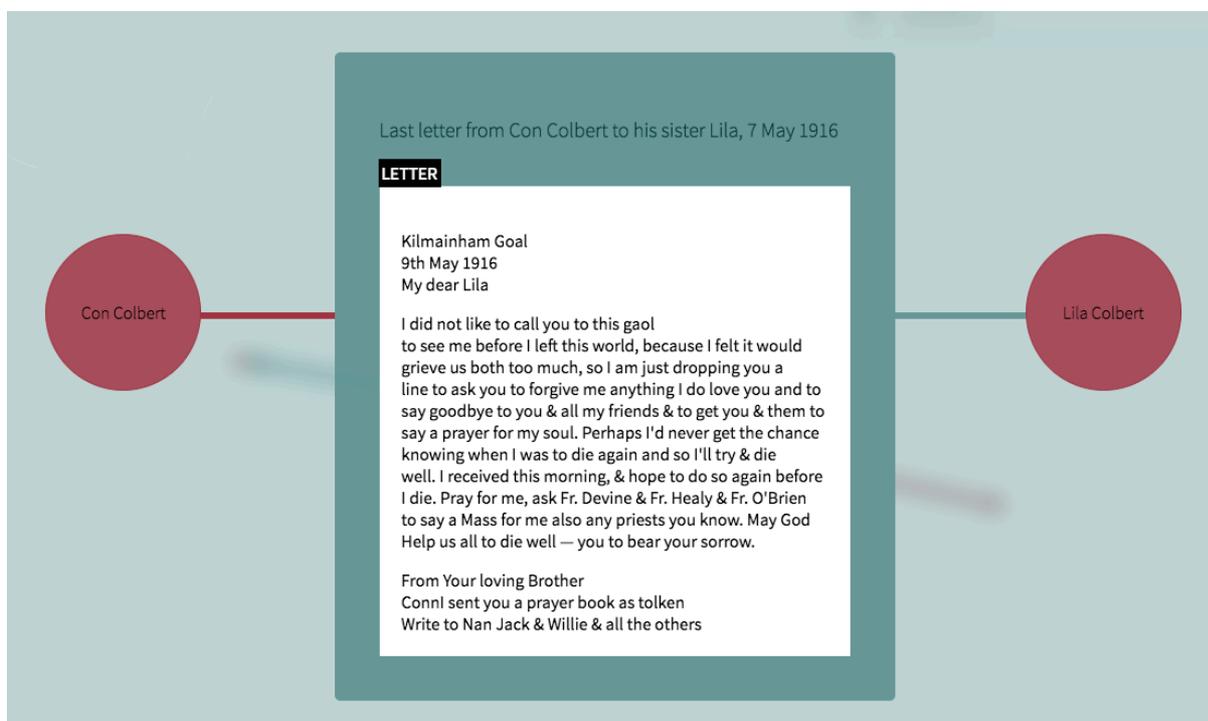


Figure 48: The letter view. The view overlays the graph view when a letter is selected (the graph may be seen blurred out in the background). The design aims to contextualise the letter as a correspondence between two individuals, preserving the visual metaphor of nodes from the graph view.

Clicking anywhere on the background closes the letter and returns to the unmodified graph underneath, providing as smooth a transition between reading environments as possible.

The text of the letter (as displayed) is generated directly from eXistDB, using a single XQuery script which invokes an XSLT stylesheet. The scope of what can be done with the text is constrained by the *Letters* encoding strategy and the textual theories instantiated by this. This is not intended as a criticism of the *Letters* encoding, rather an observation that reusing data unmodified is not the same as re-transcribing and re-encoding it: one is always bound by decisions already taken. The text is therefore transformed in a way that mixes the two formats of the official *Letters* site: a linear text is presented (as in the *Letters* reading version), though maintaining the line-breaks and page-breaks as in the side-by-side transcription–facsimile view. This is, in part, to sidestep problems caused by the removal of line-breaks, while at the same time adhering to the linear-text structure (rather than an explicitly document-based structure) that is inherent in the *Letters* TEI data. As the text view is generated directly, and on the fly, from the TEI documents, this does afford some scope for further modifications — for instance, to more clearly distinguish the pre-printed headers of pages.

6.7. Graph manipulation

As the interface for viewing the collection and accessing the letter-texts, the graph itself is designed to be a global mode for organising and re-organising knowledge. Once loaded, the interactive interface allows panning (a click-drag action) and zooming (scrolling upwards or downwards).²³⁵ Also, all nodes and edges are by default shown. However, holding the Control key (or Command key on a Mac) and clicking on a node will grey-out all other nodes and edges, leaving only that node and its connecting edges highlighted, which may be added to

²³⁵ These actions are widely utilised in such ways, Google Maps being an obvious example.

by clicking further nodes (other possible filters, such as the *union* of nodes connected to two other nodes — for instance to find all the letters between *A* and *B* and no other letters to or from *A* or *B* — have yet to be implemented in this version).

By making requests to the server-side graph-query and layout-calculating libraries (and the textual data in eXist), the user is able to dynamically reconfigure the graph to show different features. This reconfiguration is managed principally by using the “Console” in the upper-right corner of the screen. (See Figure 49, below.) Through this interface, the user is able to request new graphs, which are constructed on-the-fly from the TEI data and laid-out (or reloaded from the server-side layout cache).

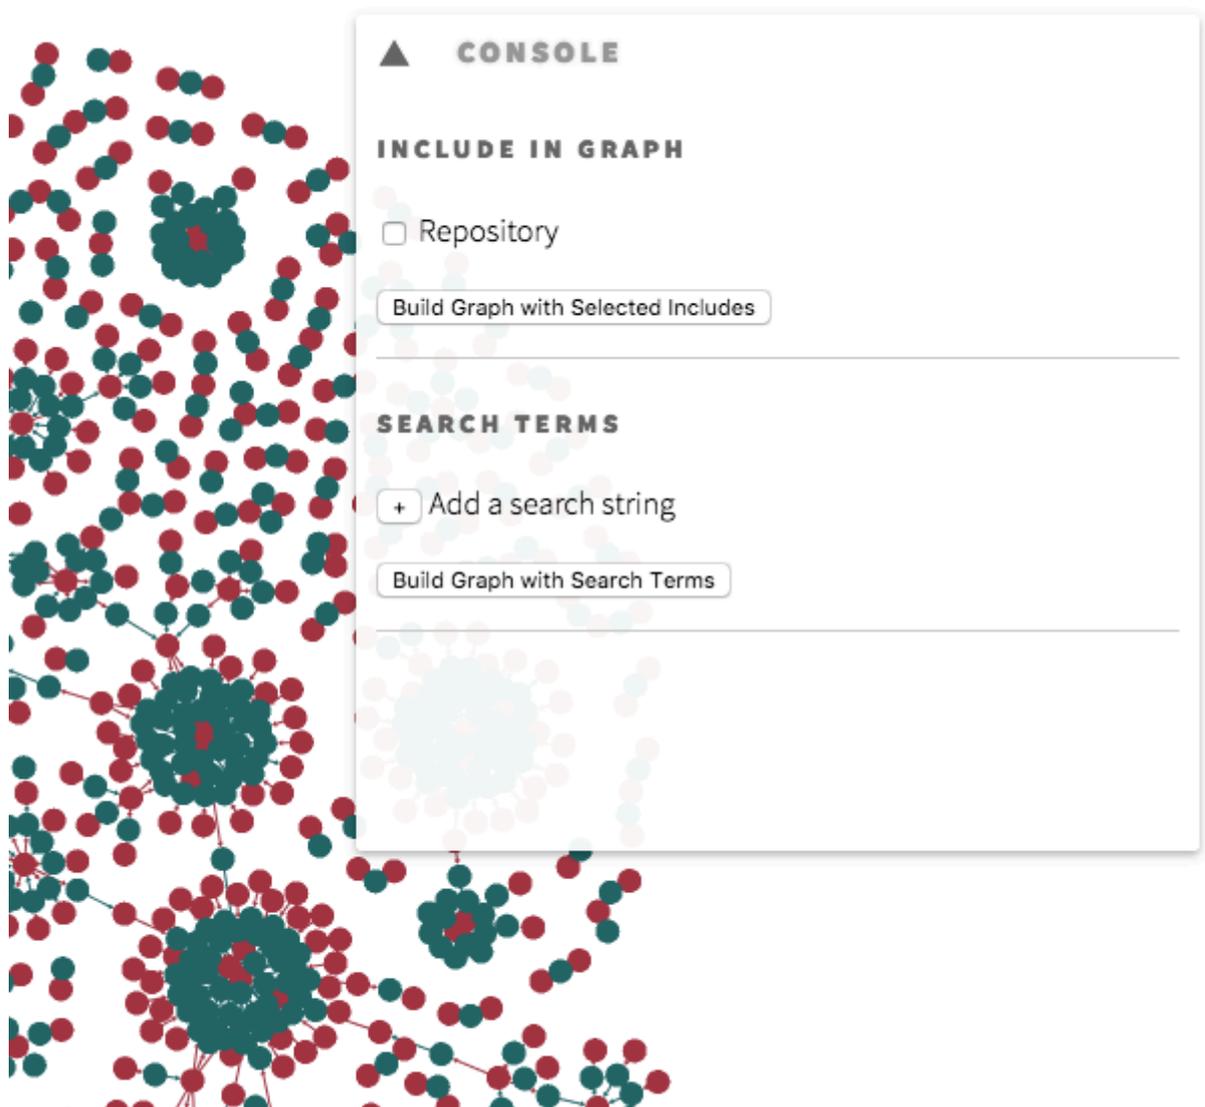


Figure 49: The drop-down “Console” panel. This allows the user to add the letter repositories into the graph, and to add search-terms as nodes to the graph.

The top section of the console allows the inclusion of additional node-types into the graph (at the moment, this is limited to the repository). On clicking the “Build Graph with Selected Includes” button, the graph is re-drawn. This, it should be noted, produces an entirely new graph, relocating nodes based on the ForceAtlas layout algorithm, which is dictated entirely by the nodes and edges that are present (there is, therefore, no guarantee that a particular node will be in the same place; rather it is extremely unlikely). This, in one way, may be problematic for the user, as a letter or person will shift position depending on the graph drawn. However, assigning to nodes (for instance letters) a fixed position makes it impossible

to gain an overall impression of clusters within the graph caused by, in this case, ‘important’ people and large repositories. To *fix* nodes in place is to undo this benefit. As noted earlier, there is no ideal graph layout; the decision to lay out the graph in this way was taken to emphasise, via proximity, particular communities (here considered clusters of people writing to each other; or people with especial importance in a given repository’s collection).

As Figure 50 shows, including repositories into the graph produces a vastly different picture of the collection. This will be discussed in more detail in the following section (especially with regards to its significance), though for now I think it is useful to note quite what an effect the larger repositories have in determining the shape of the overall collection, and how few connections there are between repository clusters. If, on the other hand (and speaking hypothetically), the letters were drawn from a more heterogeneous set of repositories the distortion of the graph in this view would be lessened (and at the limit, where each letter came from its own repository, there would be no distortion at all).

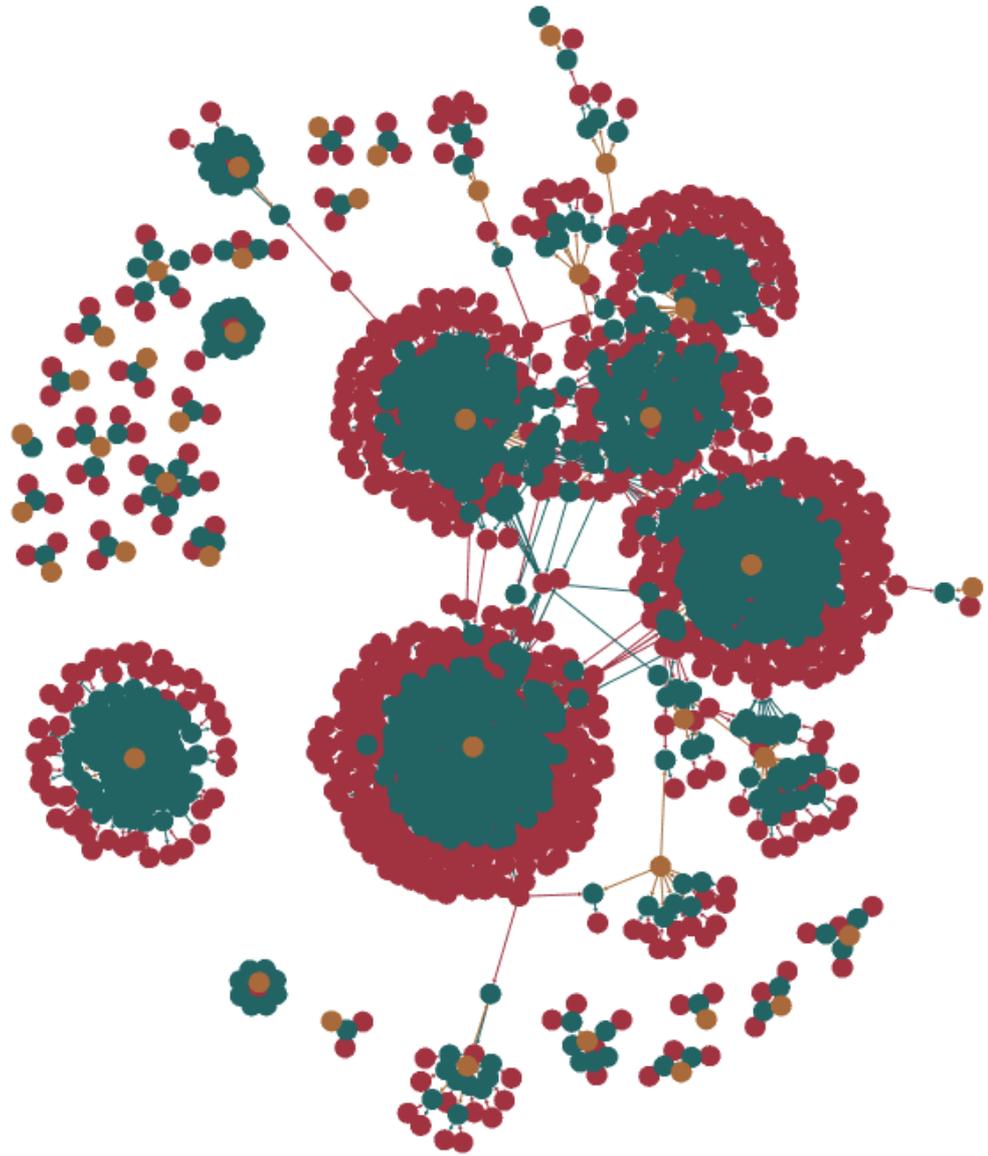


Figure 50: Graph laid out to include repositories (brown nodes); letters directly connected to repositories with a brown-coloured edge, representing “contains” (i.e. the arrow points from the repository to the letter).

Figure 51 (below) shows a closer view of the connections between a repository (in this case, the University College Dublin Archive) and the letters it contains. Note that by virtue of this *direct* connection, the layout algorithm naturally draws the letters closer to the repository, placing the senders and recipients towards the periphery of the cluster — with the exception of ‘significant’ persons in the context of the repository, who are dragged closer to the repository node by weight of letters to which they are, in turn, connected. This gives an

indication of the sub-collections contained in the repositories, which are often founded on personal collections. Again, this will be discussed at greater length in the coming sections.

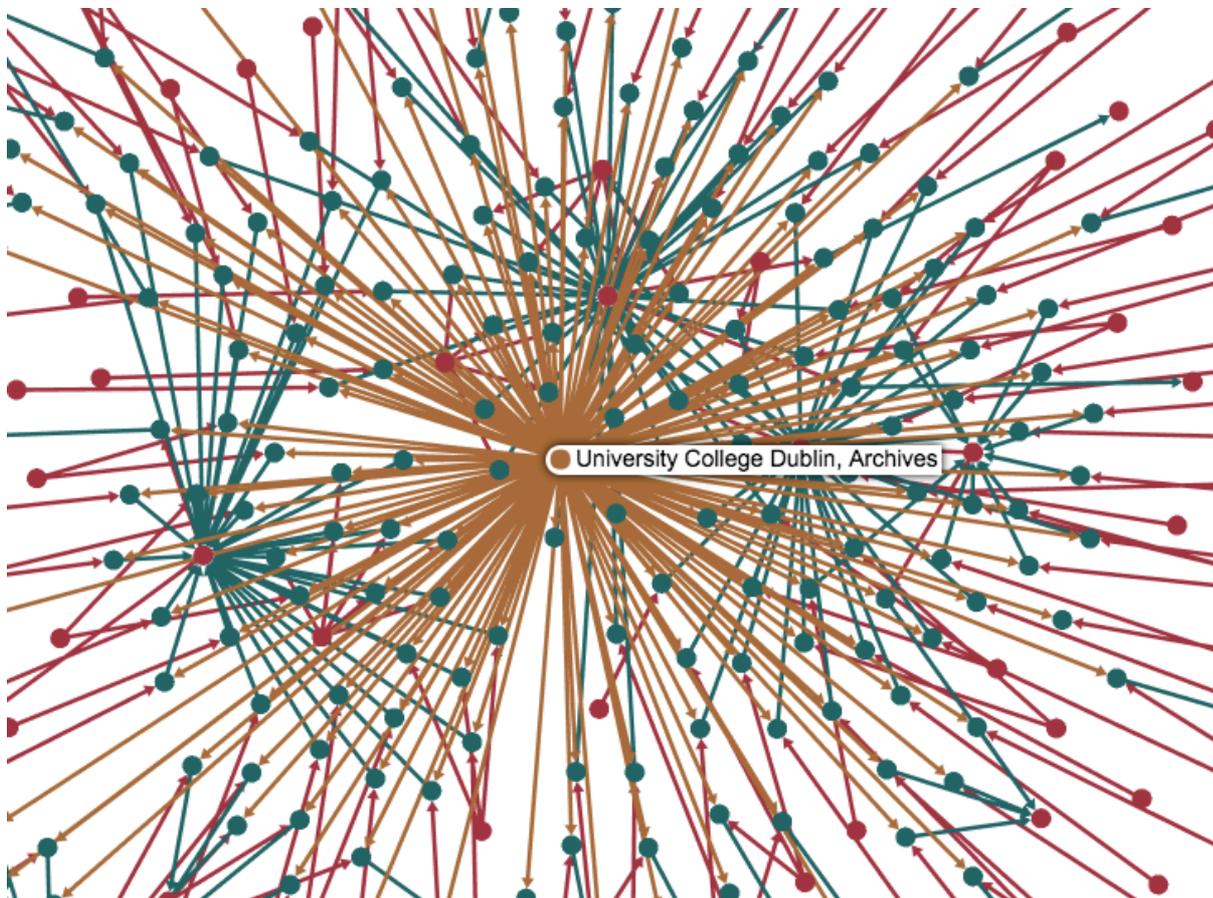


Figure 51: A zoomed-in view of the University College Dublin Archive. The brown edges connect it to letters drawn from that archive.

The final node-type that can be added to the graph is a search term (or rather, any number of search terms). As noted earlier, full-text search is implemented within eXistDB (using the Apache Lucene library), with additional useful features such as KWIC (key words in context: i.e. a snippet of the text surrounding the searched-for term). As eXist allows search indexes to be built on any XML element, it is possible to include or exclude certain fields (teiHeader metadata) as well as to restrict the search to certain sections of the document. These are

features that I hope to implement in the future; at present, searches operate on the complete text of the letters.²³⁶

The rationale for including search terms as nodes in the graph is twofold. On one hand, it is to address many of the problems of full-text search suggested in the previous chapter. One of these is the tendency of normal search pages to function, not as “search” engines but “hiding engines”.²³⁷ By default, nothing is on the page; and therefore the *negative* results are hidden. The effect of this, especially when dealing with a collection such as *Letters*, is to decontextualise; if there are search results, why? If the results are few, are they related in an obvious way (being written by the same person) or is the search term a more widespread phenomenon? Such things are difficult to determine from only *positive* results. Except in the case that a user wishes to relocate a specific letter, I would strongly argue that searching with the view to finding *relevant* items is alien to humanities research, in the sense that relevance cannot be solely determined by the presence of a particular word in a document. Rather, as Saracevic and Kantor suggest:

“[S]earching for information is far from being a science, and as yet the present heuristic rules or principles of searching as stated do not take into account some important aspects of what seems to be really going

²³⁶ It would, for instance, be especially helpful when searching for places, to *exclude* the address element: searching for “Dublin” is likely to find a lot of results that do not contain a reference to Dublin in the text, but have been sent to or from the city. Or, conversely, one may wish to include, or explicitly search, the address section for the reason that letters originating in Dublin were precisely what is sought.

²³⁷ I heard this latter phrase in a BBC radio interview with an academic studying Google’s effective monopoly on web searching (around 2012). Though I was unable to find any attribution for it later, the phrase struck me as particularly apt: one does not use Google to search *for* things, in the sense of *having things and searching through them*. Rather, it hides everything by default, bringing you the thing it thinks you want. Putting search terms “on the map”, I hope, solves some of these issues, particularly in humanities research where there is no sought-for thing, and full-text searching is another exploratory mode.

on. [...] Searching is still an art and a very imprecise art at that"

(Saracevic & Kantor 1988, 215)

In contrast to this scientific paradigm, the aim is to represent searching 'as an art'. To this end, placing the search terms onto the graph itself allows the viewing of the search term not just in positive or negative terms — finding something or not, or whether it is relevant or not — but in the context of the other connections represented in the collection.

Moreover, it addresses a problem that is specifically related to such a heterogeneous collection: there is no obvious way to sort the results. They can, for instance, be ordered by the sender's name, or the recipient's name; but this has the effect of separating letters that are part of a correspondence: letters sent from "Mr A" to "Mrs Z" would be separated — by pretty much the entire collection — from letters sent from "Mrs Z" to "Mr A" (the ideal is a combination of the two, but this is impossible in a list without duplicating results). Letters could also be ordered by repository, but this is of largely a tangential importance to the letters themselves. Or they could be sorted by date, which is useful in one contextual sense, but not in others. The essential problem is that, though *all* of these suggestions are useful ways to sort search results, they are mutually exclusive.

The reason for this is the heterogeneity (a high dimensionality) of the collection. In collections of lower dimensionality — say, letters written by a single individual, thereby fixing one dimension — the kinds of ordering suggested above make more sense. For example, if all the letters are written by one person, then sorting by the recipient, or by the date, is undoubtedly more useful. However, in higher-dimensional collections such as *Letters*, where sender and recipient are 'open' dimensions, such ordering takes on a degree of arbitrariness (one dimension or other is arbitrarily closed). The typical solution to this is to identify some measure of relevance with which to order the results.

One, rather naïve, approach is to order by the frequency of the searched-for term in the documents. This is not ideal because, for one thing, the frequency of a search term is likely to be relatively small, which means lots of letters would be equally ‘relevant’ (this is especially true for letters, which are relative short compared to, say, books). More sophisticated approaches could, of course, be envisaged that would rectify this problem. One could take the *importance* of the search term in the text (as, say, a percentage of all the words in the document), but this is easily distorted: words in short letters are automatically more ‘relevant’ than in longer letters. Other options, such as some kind of TF-IDF (term-frequency inverse-document-frequency) measure, show the importance of the search term in the context of the entire collection; but this serves to hide common terms.

Moreover, containing a higher frequency of the searched-for term is a particularly bad proxy for relevance. Algorithms such as PageRank could be adapted to the *Letters* collection (more ‘prolific’ senders or recipients would rank higher), but this suffers from the latent biases in the collection: archival collections are often founded on the correspondence of an individual, making them automatically more important. But, again, I think, the principle problem is that stated above: *relevance* in terms of humanities research is not, unlike Internet search engines, coterminous with *finding the right item*.

The graph-visualisation approach suggested here, however, affords another possibility: that the search term itself can be added as a node to the graph, with edges connecting it to letters in which it appears. (See Figure 52 and Figure 53.) Such an approach can be seen, as discussed earlier, as the *externalising* of certain textual elements (in this case, the search term) from the document. While searching a set of documents can be seen as an operation to *select documents* based on an intrinsic property, it can also be conceptualised from an intensive perspective:

that the search-act creates, on-the-fly, a rhizomatic connection between the search term and the document (or, better, the *conceptual* search-term points to *instantiations* that, themselves, are a molecular component of the document-as-semblage). Presenting search-terms on a graph is a visual representation of the latter.

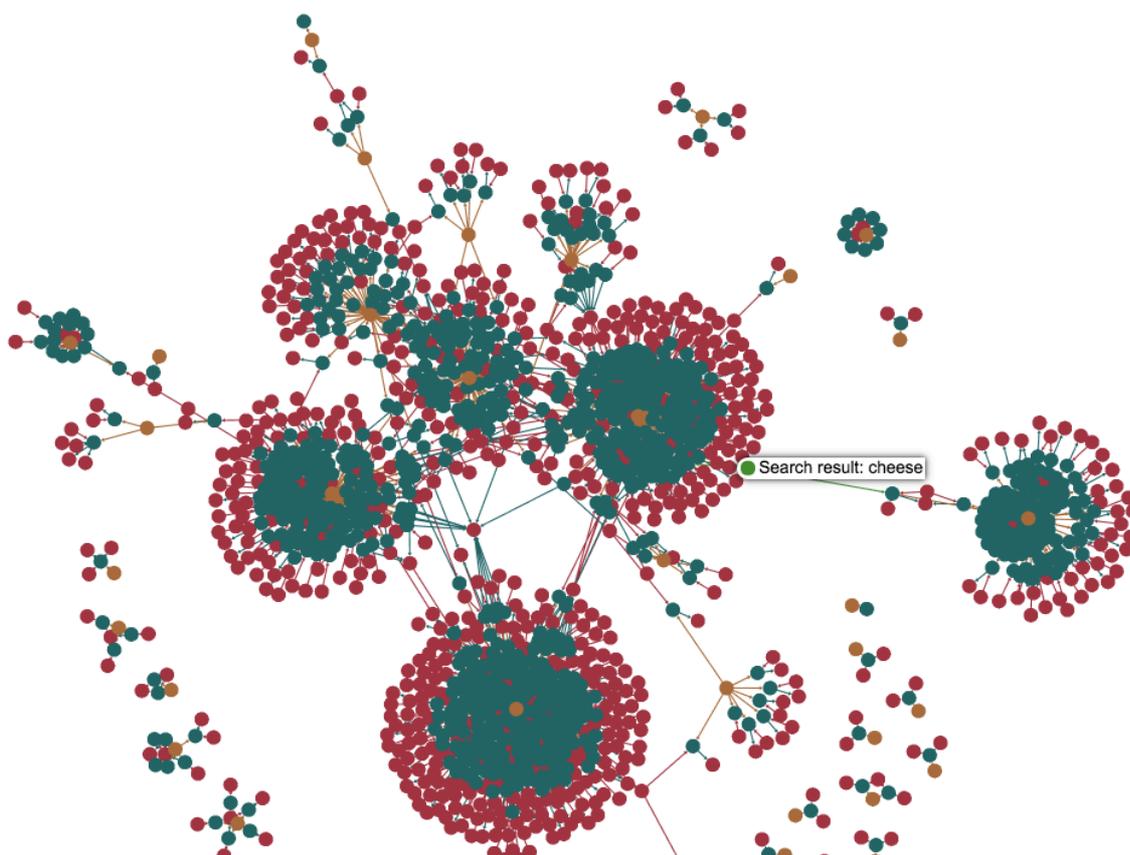


Figure 52: The search-term 'cheese' as a node on the graph. (The search term is a green node, and its connecting-edge to letters is also green.) This shows the (relative) unimportance of the term in the collection: it creates little distortion.

This approach provides, I think, two advantages. Firstly, it does not hide negative results; rather, the non-match is simply the absence of a connection. This allows the connection of the letter to the search-term to be contextualised in terms of the other variables. It also comprises a kind of multidimensional ordering. The graph visualisation's capacity to represent higher-dimension data (by treating dimensions as interconnections) means that no explicit ordering of results is required; essentially, all included dimensions are simultaneously 'orderings'.

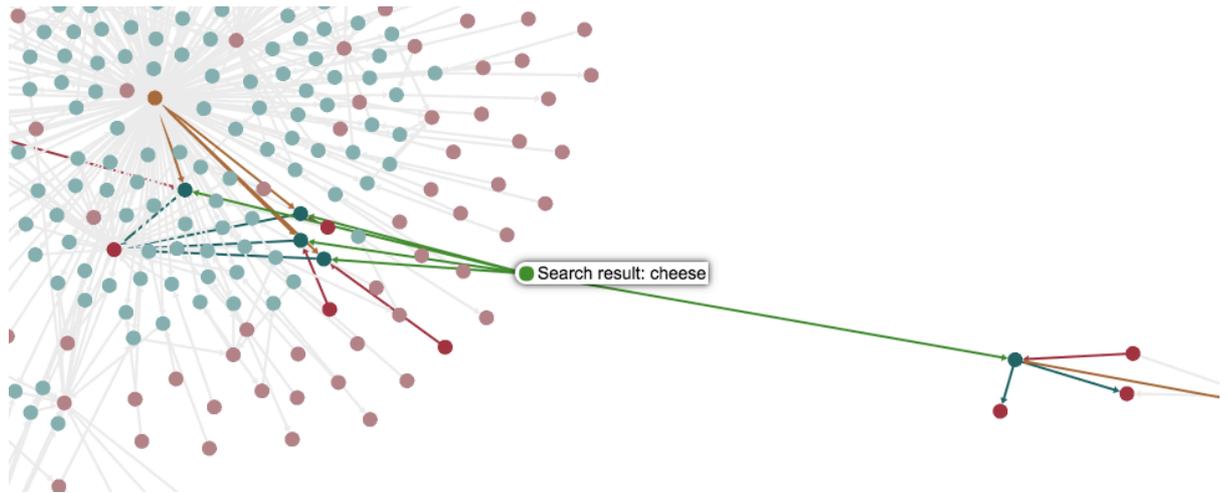


Figure 53: A zoomed-in view of the search term “cheese”, with the connections to letters, sender and recipient, and repository highlighted. (‘Cheese’, in all but one case, appears in letters to James Ryan —all from different senders.)

A further benefit is the possibility of adding *multiple* search terms to the same graph to see how they intersect or diverge. This can, conventionally, be accomplished by a Boolean search (x AND y ; x OR y ; x NOT y). Placing both search results on a graph presents these Boolean combinations simultaneously (and also *neither X nor Y*). This makes it very easy to see intersections — where the two terms co-occur — while also contextualising them in terms of other connections: *term X* and *term Y* appear together in letters from *repository Z*, and separately for other repositories (or not at all). Figure 54 (below) illustrates the overlap, in a single letter, of “cheese” and “meat” as search terms, while also showing *just* “cheese” and “meat” (the other connected letters) and all the other negative results. Moreover, this can be done with any number of search terms — an awful lot of Boolean combinations— including text-strings longer than a single word (treated as one ‘term’).²³⁸

²³⁸ To give an example, it is possible to see that the phrase “that kind of thing” is used twice in the collection, both times by Phyllis Ryan, writing to James Ryan, and both times following a list of food items that he might like to be sent (while imprisoned in Frongoch). This is, of course, a very trivial example, but one that illustrates the potential of overlaying various searches in this way.

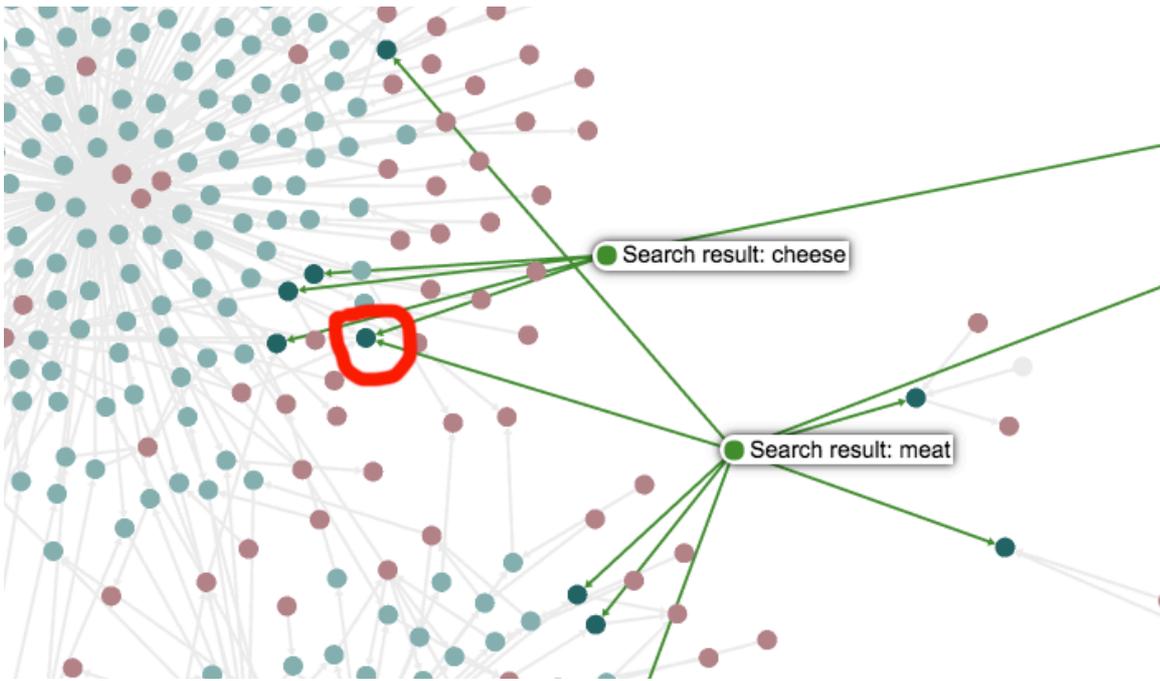


Figure 54: graph with two search terms, trivially “cheese” and “meat”; I have added the red circle to highlight the only letter in which they co-occur.

As with other aspects of this graph visualisation approach, the intent is not to give a kind of scientific or measurable validity (“cheese”–and–“meat” as an important, extractable phenomenon); rather to place text-searches within the overall paradigm of exploration, to draw attention to specific letters for further investigation. To aid this approach, it is also, therefore, necessary to contextualise the search from the perspective of the matching letters. This is achieved in two ways: firstly, when a search is carried out, to the caption of matching letters is added a brief contextualisation of the term (key-word-in-context); and secondly, the term itself is highlighted in the text of that letter when it is opened in the text-view. (See Figure 55 and Figure 56.)

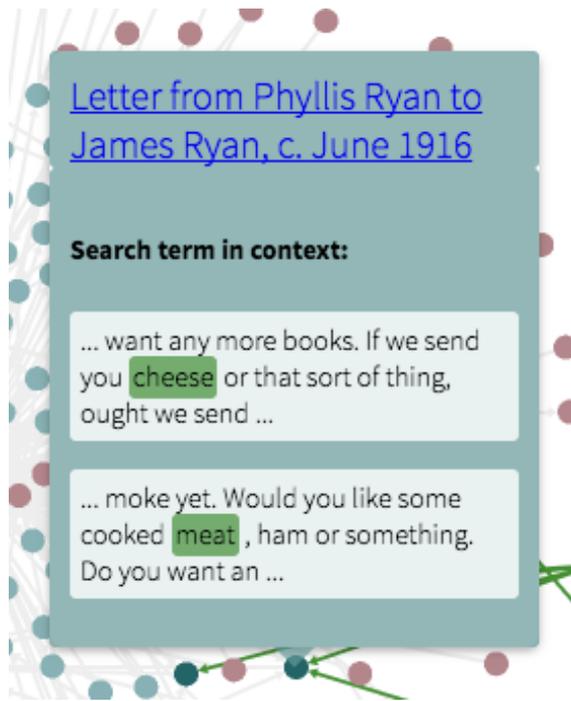


Figure 55: The search terms displayed as key-words-in-context in the letter label.

I sent you a parcel of food early in the week, and we had sent you a parcel of clothes earlier. Let us know next time you write whenever you write whether you got them or not. I shall send another parcel tomorrow of the rest of the things you asked for, and anything else I can pick up. Let us know when writing too what you would like particularly. Can you smoke yet. Would you like some cooked meat, ham or something. Do you want any more books. If we send you cheese or that sort of thing, ought we send you bread & butter too. Would you like bread & butter from home. Must everything be cooked & ready to eat.

Figure 56: The search terms highlighted in the letter text.

6.8. Citation mechanisms

It is a fundamental requirement of a scholarly edition that it be citable in future research. Approaches to this, in the case of books, are highly standardised conceptually, even where they differ in terms of form. In general, one cites the title of the book, the date, publisher and so forth, and the *relevant pages*. In the case of a print scholarly edition, one also cites the editor. In a digital edition, the same information is cited, along with the URL of the relevant

page. In a digital collection, such as the *Letters of 1916*, of course, the author may vary from entity to entity. The recommended citation (in terms of informational content, rather than necessarily citation style) used by the *Letters* cite is thus:

James Davidson. "Letter from James Davidson to his mother, Clara Mary Davidson, 3 November 1915". *Letters of 1916*. Schreibman, Susan, Ed. Maynooth University: 2016. Website. <http://letters1916.maynoothuniversity.ie/explore/letters/2168>.

In all of this, I would suggest that the URL is the most vital aspect, insofar as, in the digital realm, it is the only requirement for someone else to *relocate* the referenced digitised entity. This, I think, is a key underpinning of scholarship in general: that it should be reproducible, in either a scientific sense (replicating experiments), or the somewhat more vague humanities sense, of allowing future readers to re-trace your activities and possibly disagree with your interpretation of the sources used. (This is, by the way, not to suggest that authorship or the intellectual effort in creating the digital resource, or its institutional affiliations, are not important for very different reasons.²³⁹) Indeed, the RIDE criteria for reviewing scholarly editions has the following as a criteria:

Are there persistent identifiers for the objects of the SDE? Which level of the content structure do they address? Which resolving mechanisms and naming systems are used? Does the SDE supply citation guidelines?

(Sahle et al. 2014)

The persistence of identifiers is obviously vital; if the underlying system is changed (say, a migration to a new database system) it is important that a URL points to the same document. (Hakala 2010 provides a comprehensive overview of various PI systems and their

²³⁹ Indeed, the author (or editor)'s name is vital in this regard in print, for locating items in a library (even with such tools as ISBN numbers); this is less useful on the web, which is not rigorously organised in any way — your website does not have next-door neighbours — and thus requires a *direct* route of access.

development.) Of more concern here is the level of content addressability. When dealing with a collection of digitised artefacts in particular, it is important that the specific artefact be directly addressable with a URL; unlike, say, a digitised book, there is no ‘built-in’ addressability at lower levels (say, page numbers). The instruction, *Navigate to <http://letters1916.ie>, click on the Explore panel, then click on the Search box at the top and type in the title of the letter*, does not, in any conventional terms, constitute a citation. For this to be possible there needs to be a direct, one-to-one relationship between the URL and the unit of content — in this case, a single letter.²⁴⁰

This direct mapping of address to digital object is more challenging as the website becomes more dynamic. In some cases, clearly, a degree of self-evidence is justifiably relied upon: the *Letters of 1916* site does not, quite reasonably, provide a URL for the side-by-side text/facsimile pop-ups; addressing the letter itself is a sufficient level of granularity. (Analogously, no one cites which word or paragraph on a page — by number — they are referring to in a book; the page number is sufficient.) However, as the degree to which a single web page may be dynamically modified increases, the self-evident relation between the URL and the content referred to is lessened. Many digital editions provide the ability to dynamically turn on and off annotations, or highlight certain readings of a critical text; but this creates a problem if one is quoting the annotation or certain reading specifically. Another user arriving at the page, even if they use the browser’s built-in search functionality, will not find the quoted annotation or reading as it is not there. The general question, therefore, is whether, in a sufficiently dynamic web page, the page itself is a sufficient citation, or whether the state of the page is an intrinsic aspect.

²⁴⁰ For simplification, it is also common for database-backed web frameworks to simplify what are essentially dynamic database requests — to take a query, <http://website.com/artices/?id=1>, and map it to a more convenient external address, <http://website.com/articles/1>.

For *LettersGraph* this is especially problematic: there is, intrinsically, only one URL: *the* page of the single-page application. As suggested earlier, this problem is not as constraining as it might appear, as the browser's address can be programmatically altered (and, moreover, any change made — for instance, by clicking the Back button — captured as an event that triggers an update to the state of the page to reflect the changed address). The question still remains, however, of which states need to be encoded in the URL, and what they represent. Moreover, there still remains a distinction that must be made in this case, between *the data on the page* and *perspectives on that data*.

Ignoring the underlying graph for a moment, it is easy to encode the letter, if one is loaded for viewing (see Figure 48): each letter has a unique identifier (the one used globally by the *Letters* project). Thus a parameter can be added to the URL (i.e. `www.lettergraph.com/?letter=854`)²⁴¹ when the letter is opened, and removed when closed. Moreover, though technically a state of the page rather than a new page, this URL (specifically the letter-identifying parameter) corresponds directly to the entity (the letter) that the user wishes to cite.

By the same token, each possible version of the graph itself can be seen as a distinct entity, albeit one determined by multiple parameters. One of these is the type of node to include. For practical reasons (it is easier to parse, especially with an eye to future developments incorporating other node-types²⁴²) this parameter does not include search terms. Thus the only value for this parameter is the repository nodes (`graphIncludes=repo`) — or its absence,

²⁴¹ The domain name is purely illustrative!

²⁴² If search terms were included in the same parameter, it would be difficult to determine what was a type of node and what was a search term (i.e. does that mean *include repository nodes* or include the search-term node 'repo'?)

which returns the default graph with only letters and persons. Search terms are treated as a separate parameter, with each term separated by a vertical bar (searchTerm=cheese|meat). As things currently stand (I propose later to add different layout options, which will need a third parameter), these two parameters describe the graph that is generated on the server side. There is, therefore, again a one-to-one mapping between the encoded URL and the data that is loaded.

More challenging, however, is to encode the state of the page with regards to scrolling (three-dimensionally) around the graph and highlighting certain nodes. These are, properly considered, *perspectives on the data* rather than discrete, identifiable data points. Broadly, they can be seen as analogous to showing or hiding annotations. However, these are binary choices, while free scrolling in two dimensions and zooming in a third is essentially an infinite number of options. In such a scenario, the perspective — where the user scrolls to — is entirely a subjective position; that is, it is not based on any objectively-identifiable point within the data. Even if the view position cited is taken from a specific node about which it can be centred, this does not take into account zooming in or out on this position. Zooming in on a node presents a very different perspective to zooming out (it is the difference between focusing on a local cluster and viewing the whole graph; or levels in between).

Though it undoubtedly presents technical challenges — the digital environment is, at its root, predicated on the discrete — this can be seen rather as the consequence of a rhizomatic conception. Rather than being determined by discrete points, which are representations of molar extensities, the collection seen as an intensive assemblage moves the site of meaning outside of discrete entities into the connections between them. As such, it is essential to be able to refer, as it were, to the spaces between nodes, however necessarily imperfect the solution.

By separating the encoding of references *to the data* from the *perspective on it*, it is possible to indicate — albeit imprecisely — particular clusters or set of clusters even though they do not have determinable boundaries nor are formally identifiable. The *LettersGraph* system achieves this by taking the (constantly updating) camera position from the Sigma.js library and encoding the three dimensions (*x* and *y* axes and a zoom ‘dimension’) in a “camera” variable in the URL:

```
camera=102.47486615586426,48.96224595798349,0.0866496491381  
632  
[ x-axis ][ y-axis ][ zoom  
]
```

This allows, when the graph is reloaded with a URL containing this variable, to centre the view on the ‘cited’ perspective. This is not, of course, without problems. Firstly, reloading the graph on a different sized screen centres the graph in a slightly different position (for the moment, I can see no way around this). Secondly, it may not be clear exactly which ‘indeterminate’ position the user intends to refer to. Evidently, this latter is problematic; though, I would suggest, not necessarily more problematic than citing a particular portion of a very large webpage by providing just the URL.

6.9. Using the *LetterGraph* system

In this section, I will discuss various kinds of inferences that can be made using the *LetterGraph* system, as well as indicate many of its more obvious shortcomings, which can be taken as a basis for further development.²⁴³

Viewing the letters in the context of the whole collection is, I think, the principle benefit of this mode of exploration, especially compared to more internalising, hierarchically organised approaches. This is particularly the case in a more heterogeneous collection, where there is no natural ordering principle: a hierarchical organisation prevents there being a “View Next Letter” button (or anything of the sort), because there is no obvious ‘next letter’. The effect is rather like a book where one cannot flip pages, but must return to the contents page or index first. While I would argue this is a benefit overall, it must also be recognised that this kind of approach is potentially disorientating: there are no concrete points or determined structure to cling on to. As Dreyfus argued, in relation to hyperlinked, web-like information systems (most notably the Web itself):

[W]hen everything can be linked to everything else ... the vast size of the web and arbitrariness of the links make it difficult for people desiring specific information to find the information they seek. (Dreyfus 2009, 11; quoted in Robinson & Maguire 2009)²⁴⁴

While this illustrates very succinctly the problems inherent in a rhizomatic form of knowledge organisation, I think there are some important distinctions that can be made.

²⁴³ In this I am, no doubt, somewhat remiss in citing my perspectives on how I think this prototype system can be improved, without having engaged in any more than casual user testing.

²⁴⁴ Of course, search engines go some way to alleviating such problems, but (as discussed earlier) searching in the context of a humanities research site is not really about finding the ‘right result’.

Firstly, the *Letters* collection is not ‘vast’ in any objective comparisons to genuinely large collections (say, a digital library). Secondly, the links created in *LettersGraph* are not arbitrary; rather, I would suggest, they instantiate a few obvious conceptions, particularly physical letters being connected to physical senders and physical recipients (search terms are, I admit, more abstract). Most significant, however, is the holistic view provided. This differentiates it from more hypertext-bound systems (moving from page to page by clicking links), which function more like the Eco’s view of the library as a labyrinth: the user is *inside* the network and moves from room to room following heterogeneous paths without knowing, really, where they are going until they have got there. The *LettersGraph* view is different, in that it is an overview, an external perspective: connections are made plain beforehand, not as the result of a user’s essentially experimental movement.²⁴⁵

However, this external perspective is not without problems, most notably that it is essentially a fraudulent position to have on a rhizomatic system. A rhizome, as Deleuze and Guattari argue, has no ‘outside’ by definition. (Deleuze & Guattari 1987, 8); thus “we can only describe the rhizome from an internal perspective.” (Schuh 2016, 11) We can unpick this problem in one of two ways. The first is to suggest that the user is taken up themselves by the rhizome: they form (to somewhat oversimplify the technological mediation) connections to each and every node on the screen. But this really amounts to a positivist claim: there are no connections to what is absent. Thus the system is itself closed off. As such, a second conceptual approach (which amounts to the same thing, albeit in reverse) would be to ask *what is this external perspective?* Effectively, it is the imposition, however instantaneously realised, of a new, overarching hierarchy, capped by a kind of vantage point from which we

²⁴⁵ The distinction is that between being stuck in a maze and having a bird’s-eye-view of the maze; there is, I would suggest, a substantial difference between being able to trace paths with your finger from above and having to follow the paths *in situ*.

can see the delimitation of a closed system. The view afforded is the *Letters of 1916* itself, as it is realised at a given point in time.

As such, it is not ridiculous to suggest that the *Letters of 1916* itself be put as a node on the graph, connected to each letter to indicate its presence in the collection. But this would be the collection simply articulating itself. The problem is not so much that this statement is a tautology, but that it is redundant. A *Letters of 1916* node on the graph is completely analogous to the higher-vantage-point perspective. From a design perspective, therefore, it seemed inappropriate to add an extra set of edges, which will only contribute to the already lurking possibility of information overload, for no added informational value. At the same time, it means we must always be careful about how the graph is 'read', always guarding against positivist knowledge claims about the real world based on an perspective that, while useful, cannot be taken as a ground on which to construct meaning.

With this in mind, I would like to briefly illustrate some of the inferences that I think can be made from this presentation of the collection. They are principally concerned with the effects of repositories on the shaping of the collection. As noted earlier, the general effect of adding repositories into the graph is a large distortion of the person–letter–person network (the historical-social network, as it were). To take a rather naïve view, we might take from Figure 57 (below) the impression that Mabel Fitzgerald, Séamus Ryan and Robert Chalmers were central, 'hub' figures for local clusters, which are, in turn, connected by a series of intermediary figures.

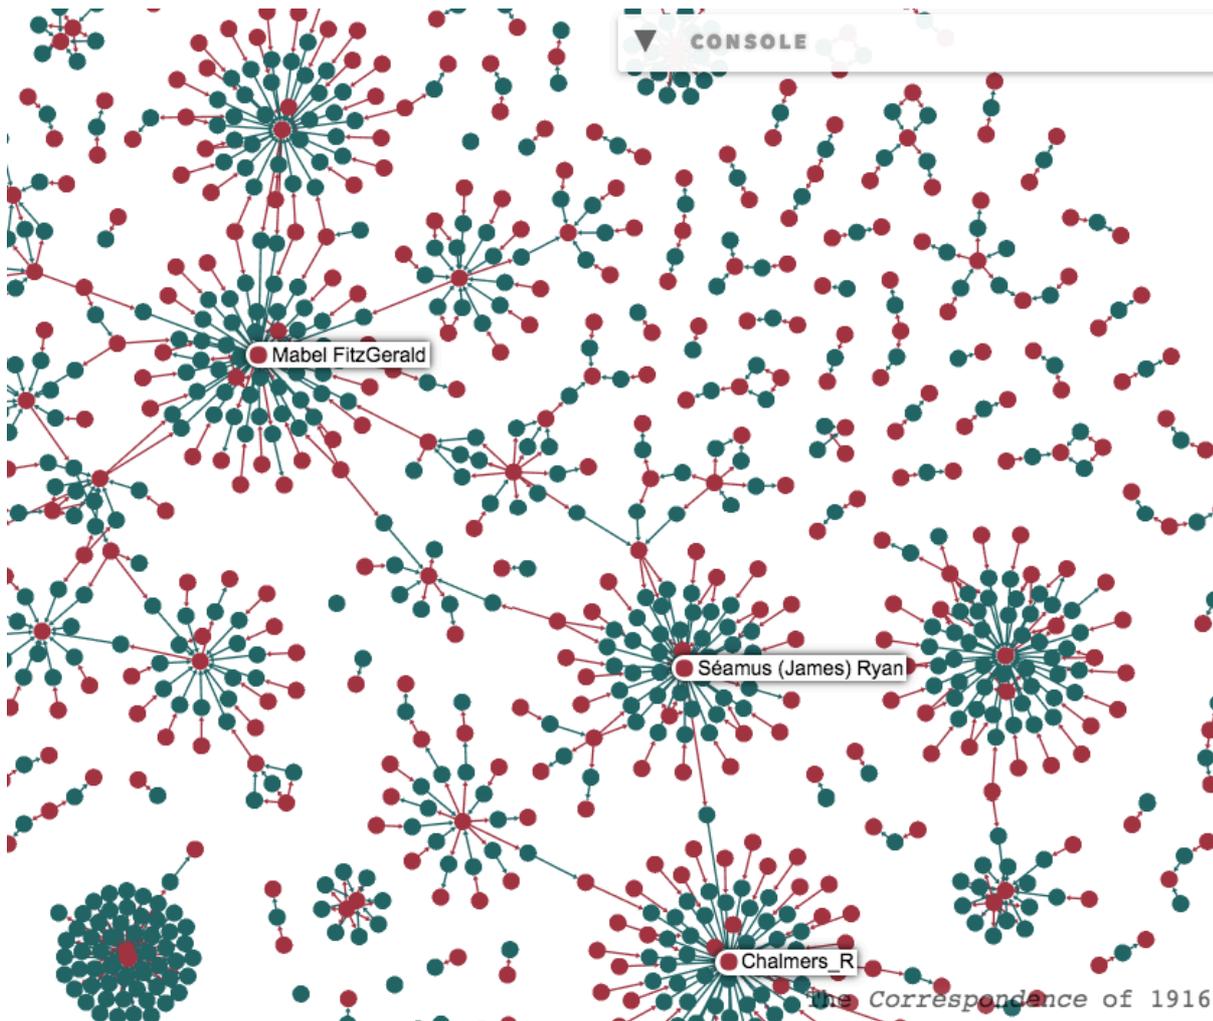


Figure 57: Mabel Fitzgerald, Séamus Ryan, and Robert Chalmers appear ostensibly to be ‘hub’ figures in a complex series of connections.

However (to take the example of Mabel Fitzgerald illustrated in Figure 58, below), we see two elements. Firstly, that Fitzgerald forms one of the key ‘component persons’ of the University College Dublin (UCD) Archive collection, along with Nancy O’Rahilly. This is made visually apparent by the ‘weight’ of letters connecting Fitzgerald to the UCD Archive that, under the effects of the layout algorithm draws her towards the centre of the UCD cluster. (All ‘significant’ figures with regards to a repository appear close to the repository node for this reason.) Moreover, it becomes apparent by looking at the colour of the letters–to–Fitzgerald edges (blue) that the letters involving Fitzgerald were *received* by her, rather than sent (indeed, the only two letter sent by her in the collection are to O’Rahilly and John Redmond, and are thereby also in the UCD archive). Here we see instantiated a suggestion

made in the previous chapter: that archival collections are founded on the collections of important figures, and specifically the letters that they received rather than sent. Indeed, looking more widely in this manner, we see that the key ‘component persons’ with regards to a particular repository typically receive letters from a wide number of persons (and often only one letter per person). Thus, receivership can be conceived — in the context of this collection — as fundamentally more important characteristic of the letter than authorship. By the same logic, the *disconnect* between senders and recipients in much smaller collections need not be indicative of the fact that they formed isolated communities; rather that the letters that survived are based on individuals’ personal connections (as the recipients), not that there were no connections.

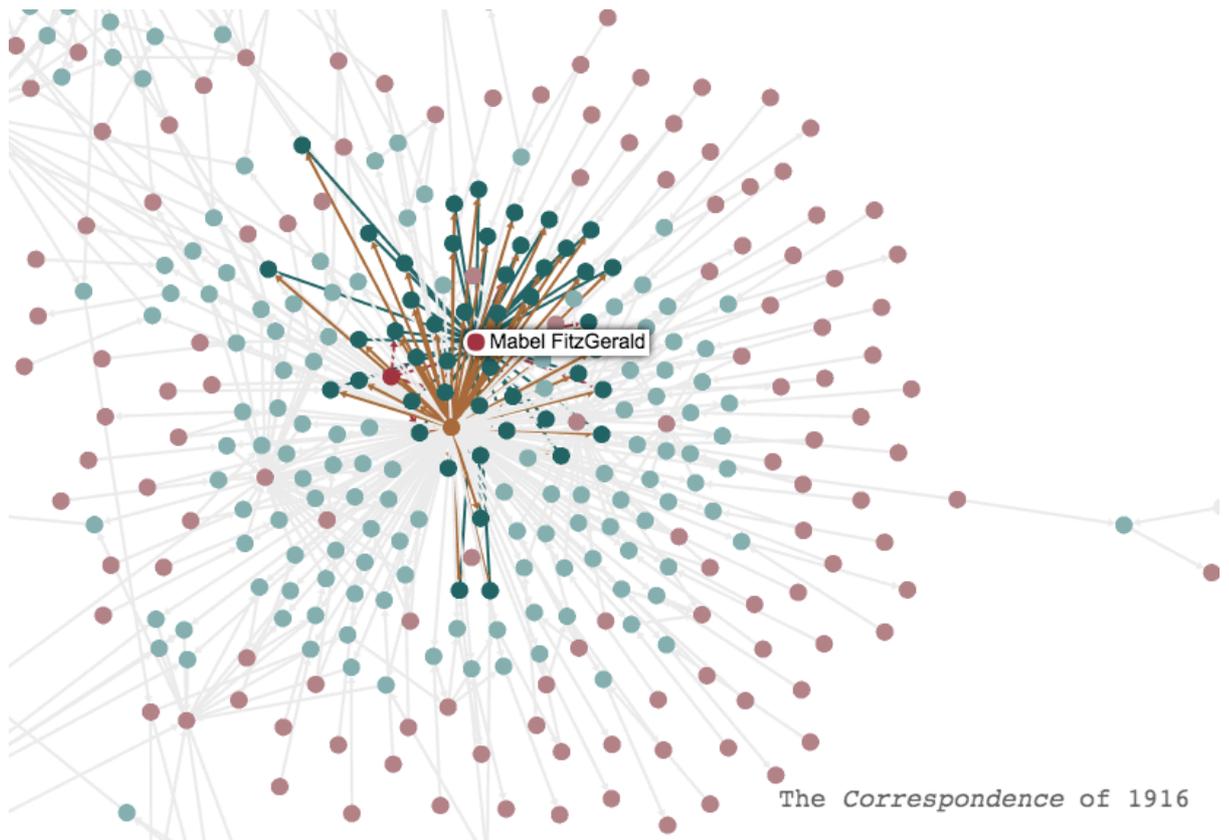


Figure 58: Mabel Fitzgerald shown as one of the key ‘constituent persons’ of the University College Dublin Archives (the brown node).

What this form of conceptualisation allows is the identification of points of doubt. It would be (I assume) naïve to suggest that Fitzgerald never sent any letters (or only two), any more than (a slightly greater assumption) that she received a lot of unsolicited letters from complete strangers. Instead, we may posit the existence of letters that are not in the collection. This is, in itself, a supposition, but then we can fall back to *reading* the letters received by Fitzgerald. (To ‘prove’ a point, as it were, the letter from Harry Stockman to Fitzgerald, 31st May 1916, begins: “I was very glad indeed to have your fine letter”.²⁴⁶) This shows, I think, the possibility of more fully considering the social structures that underpin correspondence that are, essentially, masked by the restrictions of a document-centric approach. The idea, to again reiterate, is the not the positivist assertion of knowledge, but the providing of sufficient information to the user to be able to experiment while appreciating the biases inherent in the collection. What placing the repositories on the graph does, therefore, is to indicate the most significant and distorting factor of the collection’s creation. (Figure 50, page 379, illustrates the extent to which the social-historical correspondent network presented is tied closely to the repositories from which the letters were collected.²⁴⁷)

The same kinds of exploration can be applied to the presence of search terms in particular sets of letters. Adding *Dublin* and *Belfast* as search-term nodes to the graph shows a very clear divide based on the repository, into three sets. We should not, I think, be surprised that a wide range of Dublin-based archives contain letters containing the term *Dublin* and not *Belfast*; nor, by the same token, that the Public Record Office of Northern Ireland (PRONI) contains letters containing *Belfast* and not *Dublin*. Looking at the PRONI letters, in

²⁴⁶ Harry Stockman. "Letter from Harry Stockman to Mabel FitzGerald, 31 May 1916." *Letters of 1916*. Schreibman, Susan, Ed. Maynooth University: 2016. Website. <http://letters1916.maynoothuniversity.ie/explore/letters/918>.

²⁴⁷ The *connections* between repositories are, generally, few and far between, typically arising from one or two letters from one ‘key component’ person in one repository to that of another.

particular, we see that large number of them are the Sir Wilfred Spender collection, comprising letters written from the First World War to his wife in Belfast; as such, it is hardly surprising that they mention Belfast (a ‘subject’ pertinent to the *recipient*, his wife) and not Dublin. Likewise, letters contained in the Dublin-based archives were predominately sent to Dublin (hence finding their way into those archives), and therefore feature the term *Dublin*.

The intersection — letters in which Belfast and Dublin co-occur — is more interesting. Firstly, they are almost all found in the National Archives of Ireland,²⁴⁸ suggesting they are written in an official capacity (a glance at the recipients, including Robert Chalmers and Sir Matthew Nathan, seems to confirm this). Many, indeed, are reports, written by local officials to those in higher authority. To a degree, then, we may make the positive claim that this is a real phenomenon; but at the same time, it allows this claim to be made with a degree of caution. That the *Dublin–Belfast* intersection is predominately found in official correspondence suggests that they represent those letters that are more likely to have survived. It is not, in other words, evidence of the fact that few ‘normal’ people wrote from Dublin to Belfast (of vice versa), only that they are (or might be) absent from the collection for some reason (having been lost, for instance). What I think this graph-based approach provides, therefore, is an ability to see a possible phenomenon and investigate it more fully, above all by diving into the letters themselves to determine the validity of any assertion.

A final potential problem is the inability to identify a sender or recipient, and in particular how to represent this absence of knowledge. (For the sake of this discussion, I will consider a missing *sender*; the logic, of course, applies equally in the case of a missing recipient.) When restricted to a document-centric perspective, this poses less of a problem. If the sender of a

²⁴⁸ The exception is a letter written to Piaras Béaslaí in Dublin, which refers to a mutual acquaintance having returned to Belfast.

particular letter cannot be identified, it does not change the fact of the letter's existence (it can be represented with the sender listed as 'Unknown'). When considered as a network, however, the absence of a sender has a serious effect on the shape of the network. For one thing, the letter node on the graph appears as if it came from nowhere (i.e. there is an outgoing edge from the letter to the recipient, but no inward edge from the sender). There are no genuinely satisfactory solutions to such a problem.²⁴⁹

6.9.1. Possible Criticisms

There are several criticisms that can be made of the *LettersGraph* system as it is implemented in this prototype. These stem in part from my own experience of using the system and from feedback received from informal user testing;²⁵⁰ and feedback from my presentation at the DiXiT Interfaces conference in Graz. The most general criticism is that the whole system is far from familiar for most users, requiring a steep learning curve. This is, I think, an entirely valid criticism. Certainly, the graph interface requires a particular mode of reading (especially drawing meaning from visual patterns), which is not particularly intuitive. From my own experience using the system, other problems are obvious. That there is a lot of information is undoubtedly true (and somewhat the point), but the bigger problem is one of making sense of the information, especially in tightly clustered nodes. For instance, it would clearly be helpful if the letter nodes and persons nodes were a different shape, or used some kind of icon rather than a circle. Overlapping edges, particularly in tight clusters, can make it difficult to follow the myriad connections (sender and recipient to letter, letter to repository or search-term). At present, the design is limited by the constraints of the graph visualisation library

²⁴⁹ Having *one* node labelled 'unknown' is obviously nonsensical: every letter in the collection with an unidentified author was not authored by the *same* unknown person.

²⁵⁰ Thanks especially to colleagues from my DiXiT secondment in the Centre for eHumanities in Cologne.

used, which is in turn pushing at the limits of new browser technologies. More mature technology and libraries may, in the future, allow a more expressive visual design.

However, I would not view ‘unfamiliarity’ as a reason not to do something; rather it could be seen as a requirement for a greater explanation of how it works. Indeed, while interface design is widely recognised as an important and ongoing debate in digital textual scholarship (the work of Ruecker *et al.*, 2011, represents perhaps the most comprehensive analysis of interface design for digital humanities projects), it is certainly true that there is no consensus on a standardised design for digital editions. In general, there seem to be two camps: those who favour the development of a standard interface for digital editions, and those who are more keen to experiment with the new paradigms afforded by the digital medium (normally those working on interface design). The DiXiT “Scholarly Editions as Interfaces” conference in 2016 provided a good overview of these positions. Robinson’s paper (“Why Interfaces Do Not and Should Not Matter for Scholarly Digital Editions”) reflects the former tendency (in line with his views expressed elsewhere; see Robinson 2005) Shillingsburg articulates a similar perspective. (Shillingsburg 2014) My view is clearly towards the opposite position, especially when a particular design affords new possibilities not available in an existing design. Indeed, even in print, the design of scholarly editions was not fixed, and though there was more consensus, scope for experimentation remained. The synoptic pages of Gabler’s *Ulysses* are a notable example. (As Eggert, 2009, notes, these were criticised in some quarters for being difficult to read — though, as Eggert argues, objectively no more difficult than a conventional critical apparatus.) In the case of *LettersGraph*, I feel the development of a new interface is justified on the grounds that no existing approach achieved the same aims. (One tester commented that they felt it was something they could get used to in time — which I hope to be the case).

A much more potent criticism is the general sense of disorientation caused by there being no stable points of reference in the graph layout. Another tester contrasted this to the letters being placed on a map, with geographical locations serving as a mode of orientation. This is a problem for this system, from a conceptual perspective as much as its instantiation. By taking as a paradigm the collection as indeterminate, and aiming to represent it as such, there can be no static points of reference that are not arbitrarily imposed (and thus distorting to some degree). Nevertheless, I think some mode of orientation — particularly between different versions of the graph — is important. I would therefore propose a mechanism by which a user could, when using one graph version, ‘pin’ or highlight nodes; upon loading a new version, these nodes (in their new positions) would still be highlighted. Some kind of animated transition — especially if the user could control it — would also provide visual clues. Moreover, there ought to be some kind of mechanism for locating nodes within a graph, without completely restructuring the graph by adding search terms as nodes. Again, this would require significant development beyond the prototype as it is presented.

An obvious flaw in the system is the synchronic nature of the graph. It is not that a user would draw the (obviously naïve) conclusion that all the letters were written at the same time, but rather that there is no way of knowing how the connections developed over the course of the represented time-frame. There are several possible solutions, but each is, in its own way, unsatisfactory. One is to follow the paradigm that is already established and place date ranges (say, months) as nodes on the graph, allowing them to be subject to the basic layout algorithm. This has two problems. Firstly, that the months nodes could be placed anywhere in the graph depending on their connections to letters (thus losing a sense of chronology). Secondly, any date range is an arbitrary construct, leading to a degree of over-determination, especially at the margins: it would matter a great deal whether the letter was written on the 31st of one month rather than the 1st of the next month, when in reality the distinction is (in

most cases) somewhat trivial.²⁵¹ A second approach is to provide a fixed position on a timeline for the letter nodes (running, say, earliest to latest horizontally across the screen). This is less misleading, but at the same time undermines the general utility of physical proximity as a measure of importance within a cluster; it will, rather, tend to pull clusters apart. A third option, which I think is preferable (though not ideal), is to add some form of date-range slider, which would cause those letters written within the selected date range to be highlighted. While not misleading, this is potentially disorientating: it is hard for a user to keep track of which nodes were highlighted, for instance. Moreover, it clashes with the click-to-highlight-connections mechanism. Undoubtedly, further work is required in this regard.

A similar criticism may be made of the lack of information shown about individual letters (without clicking on them). Arranged as it is, all letters are weighted equally, regardless of length or content. This could be rectified in two ways. One is to use the *Letters* project-assigned categories, placing each as a node on the graph. However, while these are not arbitrary, each letter may be assigned a different number of categories. As such, those letters with more categories will appear more centrally (they will connect multiple categories) than those with only one category. This distinction is rather artificial. Some form of topic model (such as LDA) may rectify this, by generating rather than imposing categories, and assigning category weights (which can be represented as weighted edges) to each letter, such that they would comprise in some measure all the generated topics. This would, I think, be particularly useful: it would allow the user to see the co-influence of topics and people and repositories on the collection. Moreover, as each letter would be assigned to *every* topic to a degree, there would be no need to visually represent the connections (mere proximity would represent

²⁵¹ This is not helped, moreover, by some letters not having a precise date. The *Letters* project deals with these using the TEI `@notBefore` and `@notAfter` attributes to specify a possible range. In the case that this range spans a month-boundary, which month should it be connected to? (If both, this will undoubtedly distort the graph in highly misleading ways.)

this). This is, I think, an approach worth investigating, though it would present considerable technical challenges (most notably, would the topic modelling be carried out in advance, in which case the number of topics was fixed; or would the user be able to request a number of topics, in which case a considerable computational workload would be placed on the server infrastructure).

As a final suggestion for representing the semantic content of letters, an entirely different layout could be conceived of. Each letter could be represented as semantically connected to *every other letter* as a weighted edge representing similarity (calculated by, for instance, Euclidean distance or cosine similarity metrics). It is difficult to see what the outcome of this would be, and particularly the degree to which it undermined the usefulness of the other connections. Again, this is something that would require a degree of experimentation. As the same time, it illustrates the general utility of a graph-based system, in that it can (at least in theory) represent any kind of data or metric that can be modelled as a connection. Provided the system is not used alone as a basis for empirical claims, almost any layout could in some ways be useful in fostering an experimental approach to knowledge creation.

6.10. Conclusion

In this chapter, I have described the creation of a graph-based interactive design for a digital edition. Its principle aim has been to overcome a document-centric approach, particularly in a case — correspondence — that derives much more of its significance from its social and historical context than other forms (most notably literature). Document centrality is, as argued elsewhere in this thesis, the result of a linear process of digitisation that necessarily operates predominately on single, discrete objects. At the same time, the collection as a whole comes to be seen as an empty vessel into which each digitised object is placed — a non-signifying whole. But it is necessary to recognise that the *closure* of the collection — by calling

a halt to the addition of more objects — creates a delimited, and thus signifying, whole, albeit one that cannot be attributed to anything beyond its own creation. By making the digital representation of the collection document-centric, the nature of this signifying closure is masked from users.

The function of a scholarly edition in this case cannot be achieved by encapsulating everything in a single argument — a regulative idea — for which the editor takes responsibility, because there is no single regulative idea that can account for a heterogeneous collection. Instead, we must be more explicit, making not just the remediated historical evidence, but the evidence-behind-the-evidence, available to the user: to not cut the strings connecting the edition to its creation, but to maintain them as an intrinsic part of the edition. The graph-based approach here allows the *mapping* of both the collection and the most important aspects determining its creation — thus representing the rhizomatic assemblage of the processes that gave rise to the collection: an *articulation* of its representation.

As such, while it cannot function as, in the words of Deleuze and Guattari, a Royal scientific tool — there is no rigid separation between ‘evidence’ and the contingencies underpinning that evidence — it may be seen as a vehicle for *nomad science*: an exploration of evidence in the context of its creation. This paradigm is, I would suggest, what characterises qualitative humanities research.

As a final point, it is worthwhile asking whether this approach (if not this interface) has any application beyond *Letters* (or, indeed, letters). This is difficult to determine. It is certainly not a general solution to all kinds of scholarly editions or digital collections. The strength of this approach is its application to a large set of texts that represent, or from which may be extracted, some kind of network; but whose representation of that network is substantially

determined by other latent factors. Relationships extracted from court records springs to mind as a possibility, especially where there is significant loss of textual material. I have been recently introduced to the (conveniently XML-marked-up) records of the Dáil Éireann (the Irish parliament), which would appear to hit a sweet spot between interest in a multitude of possible networks and interest in the text. Anything, in short, that demands close reading, and wider experimentation, to make sense of.

Conclusion

This thesis has not been about defining scholarly editions, or providing a theory of scholarly editing — at least not in the way that Robinson has in mind when he says that “[a] theory of scholarly editions should offer a set of principles to guide practice.” (Robinson 2013, 105) Scholarly editing is a pluralistic field, which, its history would suggest, resists being reduced to a set of principles. True, there are points at which a theory has become dominant in certain areas — Greg-Bowers in the Anglo-American sphere — but these were never absolute; counter-currents always existed and new ones emerged, as Eggert’s detailed analysis of various schools of editing shows. (Eggert 2009) Indeed, I think this thesis has shown that scholarly editing is really a multiplicity. A ‘school of editing’ might be a convenient label, but it is, at its heart, a population-of-editions. Each edition is its own school. An edition might be guided by the practices of a particular school, but it always has to consider the applicability of that school to the object being edited. That is to say, each edition re-invents a school, using pre-established principles where they stand up to the material at hand, but also making subtle modifications where required, or bringing about larger ruptures with established practice where the case can be made. This is because, as this thesis has argued, scholarly validity is something that must be created by the editor out of interaction with the source material, not simply imposed. If a literary work were really reducible to the Greg-Bowers approach to reproducing it (or, indeed, any other), most of an editor’s work would already be done in advance. Coming up with a new approach to editing is useful, but it must be seen as highly localised. It will never — a view I firmly stand by — become *the* scholarly editing.

Having said the above, I think this thesis in some way does produce a definition of a scholarly edition, even if it has required a circuitous journey to get there. It is not so much a definition

of what a scholarly edition *is*, rather what it *does*. A scholarly edition is a scholarly edition because it does what a scholarly edition does, specifically in the hands of the user. A scholarly edition is not a product but a functional capacity. A scholarly edition of a work *is not that work*, but a means of accessing it. Or, to put it better, ‘being a scholarly edition’ means a functional equivalence, between the source material and the output of a transmission. In Deleuzian terms, the source material and the edition have the same *affect*. As we have seen, affects are always in relation to something else — produced via interaction with something else. Text is an affect of documents: it is what is produced via an interaction that makes inscribed forms “resolve to language”, as Caton (2013) puts it (or, even better, text *is* the resolution-to-language). This makes the interaction of an expert palaeographer with a medieval manuscript very different to that of (to not cast aspersions on anyone else) *me* with the same manuscript. Thus producing the same expert–manuscript affect out of the user–edition interaction is not just about simple, if highly accurate, reproduction (e.g. a facsimile). It might be sufficient for another palaeographer, but it isn’t for someone else who might be interested in the text. Hence transcribing, describing, contextualising — the myriad paraphernalia that the editor surrounds the ‘text-proper’ with in order to mediate the same affect. In this thesis, I have neglected to talk overmuch of the role played by apparatus, critical commentaries and other paratexts. More fully considering their role in mediating affect is an interesting line for further study. At the same time, considered as ‘tools for mediating affect’, their role can equally be played in the digital medium by other forms than prose: a network diagram of letters contains much the same information as a written account of a group of people writing to each other.

At the same time, I think this thesis does suggest one fundamental ‘normative’ guideline — indeed, one that harks back to Bowers (1966): namely, that editors take *responsibility*. I do not mean this in quite the same way as Bowers, who criticised (among others) editors of

variorum editions for failing to decide which of the many texts they presented was the ‘right’ one (nor do I wish to suggest that scholarly editors are irresponsible!). Rather, this responsibility might best be called a ‘responsibility for hierarchies’. It is, perhaps, tempting to see a thesis underpinned by the theories of Deleuze and Guattari as implicitly opposed to hierarchies. But I think this is a mistake. Deleuze and Guattari do not deny hierarchies, or call for their wholesale tearing-down. “Never believe that a smooth space will suffice to save us.” (Deleuze and Guattari 1987, 551) Hierarchies — including hierarchies of representation — undoubtedly exist. What Deleuze and Guattari argue is that they are not universal but created, and we should understand the conditions of their created. An editor who carefully constructs a ‘hierarchy of representation’ of a literary work is not a tyrant overriding the potential of the user by presenting their (the editor’s) text of the work, as the so-called ‘un-editing’ movement would seem to imply; or they might be a dictator, but rather a benevolent one — and a necessary one. This is a role an editor must take, and must take consciously. An editor throwing remediated source material at the user is just as much a tyrant: they are still creating hierarchies of representation, whether they are conscious of it or not. But these hierarchies are worse: they exist, only in the shadows. It is no good ‘passing the buck’ to the user without telling them which bucks are passed.

Back to *functional equivalence*. Determining this functional equivalence is, as we have seen, problematic. This is because it is inherently dependent on the editor’s interaction with the source material — a unique event: the coming together of a unique editor and a unique document. There is nothing, especially on the side of the document, that pre-determines the result of this interaction. Another way of putting this, as Pierazzo does, is to say that no interpretation can be objective. (Pierazzo 2011) However, this places no fundamental restriction on interpretation, which is a problem. How can the user, to whom this interaction is being mediated, be certain of anything about the source material when the editor could

have done literally anything with the source material? This is not to suggest for a moment that scholarly editors do whatever they like with documents; rather, it is the restrictions that editors *do* place on themselves that, in a way, define what it means to be ‘scholarly’. Such restrictions, I argue, are born out of repeated, experimental interaction between the source-assemblage and the editor, in which the editor tests the limits of interpretation, and thereby the limits of the source-assemblage itself. This limit, borrowing Eggert’s term (Eggert 2009), I called the *regulative idea*.

Being determined in such a way has implications for the operation of a regulative idea. Most notably, it is not universally objective. However, it carves out a small space of ‘objectivity’: within the boundaries set — inside a little world created by the editor — subjective interpretations gain the status of objectivity. It is what Deleuze and Guattari call a *signifying regime*: the regime of, for instance, structural linguistics; of signifier and signified. Both the source material and the edition then *point back* to the regulative idea, as a “Supreme Signifier” in the words of Deleuze and Guattari (1987, 127). Interpretation, then, is the creation of a hierarchy, in which the regulative idea sits at the very top. This is, I think, the general logic by which scholarly editions, as traditionally conceived, operate; indeed, how they determine scholarship and rigour. It places the editor in a position of responsibility. A user of the edition ‘buys into’ the signifying regime created by the editor, accepting its premises, and in turn limiting their own interpretation of the edition in line with the regulative idea. However, if a user does not — if they use the edition as a doorstep — then all bets are off (again, a scholarly edition is not an *is*, but a *does*).

This kind of transmission I called a *reflexive transmission*. It is distinguished from an *articulated transmission*, which describes what it *does*, but does not (or cannot) make that ‘what’ the subject of a self-consistent regulative idea. A *naïve transmission*, by contrast, makes

no claims whatsoever beyond claiming to transmit something. The latter two, and naïve transmission especially, cannot, I think, be considered scholarly editions. They provide no basis, no foundation, from which to carry out further scholarship.

In later chapters, we saw the implications of this logic. Some source-assemblages are more easily formed into hierarchies — made subject to a regulative idea — than others. A work is easy, by this standard. It is singular and abstract. A work is the text that can be separated from its material instantiations. Even when a work is produced as a plurality — say, a variorum edition reproducing multiple versions of a work — the *idea* of the work (singular) still functions as a regulative idea. Reducing a document to such a hierarchy is more difficult. The text and its materiality are connected in myriad ways, creating little knots that the editor must somehow unpick. As we have seen, for instance, two documents can be united by a text that runs across them, but is then what is being considered the material surfaces or the text? Really, it is both; but an editor must elevate one or the other. Potential ambiguities abound. A rhizomatic system does not readily become hierarchical. An editor must break certain connections, elevating some components over others.

The very possibility of a regulative idea — reducing a system to a singular — is called into question by more heterogeneous collections of documents. This is, in part, because the documents themselves have a certain self-consistency, and may be made subject to their own regulative ideas, which may conflict with the wider regulative idea tying all the documents together. This system is what Deleuze and Guattari describe as a ‘radicle’; a plurality, but one that is subsumed under a greater whole. This may be seen most readily in documentary editions of, for instance, historical correspondence. Each letter has a life of its own, but its selection is assured by a higher regulative idea, such as authorship.

Where the possibility of a regulative idea fails completely is in even wider, more heterogeneous collections. These simply cannot be made the subject of a self-consistent regulative idea. They are determined only by what *is* done. A regulative idea to describe this becomes riddled with inconsistencies. It is the constant search for a higher level that would account for such a heterogeneity; and, when such a pinnacle is tentatively posited, it in turn demands the inclusion of other things that were not. In Chapter Five, I suggested that *Letters of 1916* is an example of such a collection. It is for good reason that it does not utilise the definite article, ‘*The Letters of 1916*’ — which would be to stake out a position for a regulative idea, a consistent signifying totality. Rather, I would suggest — and in no way disparagingly — it is better described by the partitive: it represents *some* letters of 1916. It can be described only by what *is* there, which cannot be reduced to a single idea. It is a rhizomatic system, “replacing the closed equation of representation, $x = x$ not y with an open equation: $\dots + y + z + a + \dots$ ” (Massumi 1987, xiii) Such a system cannot operate as a scholarly edition, at least via the logic of creating a signifying system.

A possible criticism that might be levelled at this ‘definition’ is that it does not draw very clear boundaries. Particularly with documentary editions, there is clearly a scuffle of sorts between the regulative idea of the whole and that of the component parts; things are far from clear cut. To treat this line like the crossing of the Rubicon is plainly overdoing it. I would argue, nevertheless, that a line being somewhat indistinct does not mean it is not there. Instead, I think, we might think of it more like crossing the Equator at sea. One might never know where the line is precisely, but it is there, and its having been crossed becomes increasingly apparent: different constellations come into view in the night sky; the Coriolis effect operates in the opposite direction. This is, I think, the case here: we might ignore minor aberrations in edge cases, but as collections increase in size and, more specifically,

heterogeneity, subsuming them into a signifying system, under a regulative idea, becomes impossible.

I think a reasonable argument can be made that such cases were ignored (or, better, *ignorable*) due to the nature of pre-digital media. It is not as if large collections of texts did not exist pre-digital, either as repositories for original documents or remediated collections. In such cases, however, the constraints of the medium, particularly having to deal with physical objects, meant that users necessarily engaged with individual components, one at a time. In Chapter Four, I raised Kline's warning of print-based documentary editors being potentially misleading in their selection and ordering of documents. However, for a user to be misled requires that user to engage with, and make arguments about, the collection at a wider level than its constituent components — which is less easy to do when the primary mode of interaction is a kind of close reading. With large digital collections, however, there is much more scope for considering the collection itself as an entity. As the final section of Chapter 4 argued, this comes about not simply by bringing computational analysis and 'distant reading' methodologies to bear. A hypertext system that enacts various lines of movement through the collection marks a shift from the individual-collected to the collection. Even apparently innocuous finding aids, such as full-text search, create 'meaningful' results by raising the collection-as-a-whole to the status of a meaningful entity.

This is the crux of the problem. On one hand, we want to create large, heterogeneous text collections, and there are numerous advantages to doing so. Were *Letters of 1916* to create a 'signifying' collection subject to a regulative idea, it would have to focus on some discrete area: say, political letters, or letters very explicitly concerning the Rising; in other words, impose a kind of dominant narrative, to the exclusion of the everyday, lived experiences of people in Ireland. The latter are too disparate, too heterogeneous, to be moulded into a single,

coherent collection that *represents* something (necessarily *something* in the singular). The idea of providing a “candid ... glimpse into ordinary life, love and work” (*Letters of 1916* 2016) is anathema to the kind of system that is required to create a globally-signifying totality. On the other hand, having created such a heterogeneous collection, we do want to consider it as a collection, not just as an accumulated set of letters to be read one-by-one. We want to use digital tools, to ‘distant read’. (This is particularly important for *Letters*, as a lot resides beyond the physical confines of the single letter, at a higher level: at the level of the network of correspondence.) But in doing so we are on shaky ground: the data cannot be trusted. I do not mean that the accuracy of the data might be lacking (it may be, of course; though in the case of *Letters* it certainly isn’t); rather that we cannot know what the data represents. We want to use *Letters* to study Ireland in 1916, not to study the *Letters of 1916* collection. This is, as I argued above, the whole point of textual transmission: one is interested in what is being transmitted, not the transmission. (No one reads Gabler’s *Ulysses* because it is Gabler they are interested in.) But we cannot: the collection is asignifying; it represents only itself. We can, however, attempt to analyse just what it is as an object of study.

In Chapters Three and Four, I argued that understanding just what *is* the object of study in the case of digital collections of this kind requires an understanding of the process of its creation. Indeed, *reflexive transmissions* (\approx scholarly edition) also need an account of their creation. To this end, Chapter Three outlined a conception of a transmission process as a set of functions, that pass a textual assemblage between them and modify (interpret) it at each stage. What makes a reflexive transmission is a recursive structure of functions: trial, error, validation; a gradual development of a regulative idea. The creation of larger digital collections involves a more linear organisation of transmission functions — that is, there is no repetition that would allow for the development of a regulative idea, at least at the global level. However, there might be *localised* points of recursion, which could create local, typically

component-level, regulative ideas. Identifying such points allows the components at these points to be considered as reflexively transmitted.

Modelling the process in such a way also allows the identification of other points at which decisions made somewhat arbitrarily (at least, without reference to a regulative idea) become 'locked in'. In the case of *Letters of 1916*, the most significant of such points is at selection. The selection of a particular letter is determined by myriad factors, not least the choice of archives in which to look.

From this basis, I argued, it is possible to arrive at an alternative conception of a scholarly edition, though not one that creates a signifying system through the establishment of a regulative idea. In Chapter Six, I suggested that a digital collection that 'built in' the significant aspects of its creation could create an alternative kind of functional equivalence. The logic underpinning this is that what the collection *is* is inseparable from its creation; or, the other way round, that explaining the collection as an entity cannot be done without explaining its creation. A regulative idea works by reducing creation down to, ultimately, to its own single axiom: the decisions taken were taken because the thing being transmitted is what it is — at least within the signifying system created. Without this reducibility, any account is as long as the collection itself: it must be, because it needs to account for each and every component individually. This, though, is not quite true (or rather, it is the case for a *truly* heterogeneous collection). There are clusters of decisions that are reducible to a single decision. In the case of *Letters*, the most significant is the choice of repositories in which to look for letters. The approach therefore taken in Chapter Six is to model both letters and repositories as a rhizomatic assemblage (along with senders and recipients): to lay everything out on a single plane. This is important as repositories do not signify anything — they cannot

be elevated hierarchically. What they do, however, is form connections to letters, which can be interpreted, weighed up.

The *LetterGraph* prototype outlined in Chapter Six proposed modelling this kind of rhizomatic system as an interactive graph. This enables, I would argue, a different kind of functional equivalence: not of the whole, not universally and not automatically. Ultimately, it must be recognised that the collection as it exists, the network of correspondence generated from it, and the repositories from which the letters are drawn, are fundamentally inseparable. The system is a rhizome. These inseparable elements can be marginalised, hidden from the user, but their effects are still felt. Instead, the conception here does the opposite: it brings these connections out into the open. In doing so, it puts the onus onto the user to explore the collection and the correspondence networks therein, while also providing the tools to factor into their exploration the most significant decisions underpinning the collection itself. Accordingly, I think it can function as a scholarly edition.

The type of knowledge that can be drawn is, therefore, not universal — at least, not by default. But, it should be recognised, neither is a scholarly edition that is bound by a regulative idea. The only difference is that someone else — the editor — has already wrapped up the loose ends, marginalised out contingencies in the name of the regulative idea. Scholarly editing, traditionally conceived, might be regarded as an artisanal activity that produces a scientific instrument in the form of an edition. (Dahlström talks of a scholarly edition as a “scientific tool” in this sense: Dahlström 2004, 23) The kind of rhizomatic edition conceived of here is the opposite: a ‘production line’ type of process that produces something in need of ‘artisanal’ interpretation.

*

It is, finally, I think, important to ask several questions, most notably: How does this fit in with other theories? And where does it take us? In answer to the former, it should be first recognised that it in no way presents a ‘new’ theory of editing — that is, a new theory to supplant others. The arguments about functional equivalence and regulative ideas are really an articulation of an underlying basis for existing editorial schools. These arguments present *text* (and works and documents) from a very McGann-ian perspective, but not editing itself. Some editorial theory is always needed, whether it is borrowed, adapted or created anew for a given edition. Having a mode for considering these various approaches without simply juxtaposing them, and instead looking at the underlying logic of how they create transmissions, is, I think, a useful critical tool. It has the potential to aid our understanding of increasingly complex and hybrid workflows.

That said, the idea of a ‘rhizomatic edition’ can be seen as having some parallels with that of the ‘social edition’. While, however, the social edition as conceived by McGann brings the ‘social element’ of the creation of literary works to the fore, the approach turns the same kind of analysis towards the edition itself. In this respect, it is closer to Siemens’ ‘social edition’ of the *Devonshire Manuscript*, whose history of production is made available by use of a Wiki platform.

As for where this might take us in the future, the obvious answer is that it requires a lot more work. No digital system ‘embodies’ a theory, least of all a prototype system. While, I think, the theoretical basis for such an edition has been established here, whether it works in practice is another matter. The system presented here demonstrates potential. Exploring this potential further, however, requires a good deal more research, and particularly user testing. It is also evident that further development is required to the system itself. Many useful features could

— indeed, really *should* — be implemented (see the final section of Chapter Six). Most obviously, named entities within the letter-texts could be ‘externalised’, turned explicitly into references to additional nodes in the network. From a purely design standpoint, further avenues are open. (On an idle afternoon, I put the same collection into a three-dimensional graph library. Repositories and their letters really do appear like constellations. However, having to navigate such a space — using the same controls as 3D-shooter games — was more of a challenge.) It also remains to be seen whether other approaches to text analysis can be made compatible with such a system, or indeed theoretical conception. Finally, there is the question of exploring use-cases beyond a collection of letters that themselves ‘naturally’ form a network.

Appendix A: On Documents and Selection

The intention of this appendix is to provide a more full account of the selection of documents. In particular, it was spurred by a question put to me by one reader: ‘What of digital documents?’ This serves as an obvious riposte to the assertion made in Chapter 3 that documents are necessarily selected — in the very first instance — in their physical extensity (though not, in should be stressed, *by* their physical extensity). This, as I wrote the chapter in question, seemed a natural enough statement. One must, after all, pick up the document, and by extension everything else that happens to be stapled or glued to it — any Post-It notes — before one can ‘interpret’ any ‘textual matter’ that happens to be instantiated upon it. With regards to digital documents, this assertion appears less valid.

To begin to tackle this problem, we need to first return to a more obvious question: how do we know, having ‘grasped’ a document, that it is indeed a document? Its physical extensity is hardly a place to start. By no stretch of editorial reconceptualisation is the tea-cup I am currently holding a document. As Eggert suggests, a ‘document’ without text is not a document: it is a piece of paper (or a tea-cup with no writing on). This is all very well (and quite true), but the question remains: how do we establish that a document does indeed ‘contain’ a text, and is thereby a document and not a blank piece of paper, or a scribbled-on piece of paper?

In approaching this question, I would like to consider a ‘thought experiment’ expounded by Eggert in *Text as Algorithm and as Process* (2010). It relates to a prisoner (presumably about to be executed) who is invited to confess to his crime, and given a piece of paper and a pen with which to do it. Upon being taken away to be executed, the prison guards find the piece

of paper blank. Is there, Eggert asks, a document in this case? He cites a conversation with Dahlström in which they disagree on this point, Eggert contending that there is not a document. If I can guess at Dahlström's argument to the contrary,²⁵² I would suggest that he would point to the 'non-confession' actually being a text, and hence to the piece of paper being a document. For instance, we invoke the notion of authorial intent, the piece of paper surely contains the text of the prisoner's intention not to confess.

There are several arguments against this position. Eggert's suggestion that the piece of paper might then be made into a paper aeroplane and is therefore definitely not a document carries, in my view, little weight (for one thing, it could quite easily have been made into a paper aeroplane even if a confession had been written on it). What of Caton's argument that a text 'resolves to language'? In the most obvious sense, a blank piece of paper evidently does not resolve to language. However, in opposition to this, one could argue that (paralleling the view that it is the text of a 'non-confession') that it does in fact resolve to language — albeit a 'non-amount' of it. This may seem a dubious claim at first (after all, the same may be said of any blank piece of paper) but I think one that can be justified on the grounds that this is not any old blank piece of paper we are dealing with.

Let us side-step this question for a moment and instead consider what it would take to make a scholarly edition of the piece of paper in question. On one hand, it could be argued that it is impossible to extract and remediate the text of the piece of paper (there not being any there). However, the case could also be made that one could (for whatever reason) make a facsimile edition of the page. One could also 'transcribe' the text — surely the easiest editorial task going — as a big, empty space. Would one consider this a scholarly edition? No: but

²⁵² My apologies if this is not the argument!

not because it was empty, but because it is no more a scholarly edition than any random clump of text copied off of another piece of paper (one with writing on), were it accompanied by no other information.

This return us to the argument made in Chapter 2, that merely copying something does not produce an edition. What is lacking is *description* (an editorial rationale): specifically, of the determination of the boundaries of the document–interpretative assemblage created that gives rise to the editor’s interpretation.

Let us consider the assemblage in this case. It is more than just the piece of paper; it is whatever critical interpretative tools the editor brings to bear upon it: language and textuality, certainly; but also anything else contextual or historical that can be employed that, taken together, produce an overall coherence of interpretation. In this case, it also includes the prison guard who handed over the piece of paper to the scholarly editor with the words, “That’s his confession!”

Without *knowing* that it’s a (non)confession — without the prison guard part of the assemblage — it really is as Eggert says not a text and not a document. But ignoring this point is to undermine the premise of the thought experiment: Eggert has *defined* the piece of paper by its context (delimiting an assemblage) at the outset. The prisoner–paper–prison-guard–interpreter assemblage in this case constitutes the text (as a page-long ‘
, not no-text) and thus the status of ‘document’.²⁵³ By the same token, a brand-new blank piece

²⁵³ By the same token, a map of perfectly flat terrain with nothing in it is still a map of that terrain — given other interpretative mechanisms as part of that assemblage that creates the status of ‘terrain’ and not a problem with our eyesight.

of paper plucked fresh from the pack is not a document (the contention that it too was a non-confession would fail to stand up to the other facts of the case).

This example, obviously contrived as it is, serves to elucidate a specific point: that ‘grasping’ a document is not always a self-evident operation. The relative facility of grasping pieces of paper with writing on them — documents — creates an illusion of self-evidence. But this is nothing more than a particular aptitude (*affect*) of a piece of paper to be grasped by another seemingly ‘natural’ apparatus at hand: the human hand. As suggested in the introduction to this appendix, grasping-by-hand also pulls along things that happen to be attached to the page, such as Post-It notes (by the same token, grasping the Post-It note ‘pulls along’ the document with it). The peculiarity of physical assemblages is their response to physical action: anything that is part of a physical whole is grasped.

But, to take things to an extreme in one direction, we can imagine a piece of paper attached somehow to something far less ‘graspable’ by the human hand, say, a rhinoceros. One’s immediate reaction is to divide (mentally) the paper from the rhinoceros, and, if one dared, to attempt to pull the paper off of the rhinoceros. Only the most foolhardy person would consider the rhinoceros–paper assemblage an extensive unity to be grasped in its entirety. The ‘context’ of the rhinoceros is merely a physical hazard to be avoided

But the same is also true in the other direction, into the realm of the non-physical. The physical extensity of the non-confession-paper when grasped only physically produces only Eggert’s conclusion: that it is not a document. It is only when we grasp the document along with the explanation of the prison guard that it becomes a document (thus we grasp the document with our hands and the explanation with our ears). The question is then: is this also ‘contextual’ in the same way that a Post-It note (or, indeed, a rhinoceros) attached to

the document is? Evidently not in the *same* way (as if a Post-It note and a rhinoceros were contextual in the same way, despite both being physical!) But are there clear grounds to draw any dividing line? I would be hesitant to draw any. There are only differences in the *mechanics* of grasping — which is in reality a negotiation between different kinds of form, that is, interpretation. It is the grasping that creates the extensity.

Thus, to conclude this digression, the assertion made in Chapter 3 — that documents are in the first instance selected in their physical extensity — is true only insofar as it is our typical experience of documents. They are things one can pick up. When we pick them up, things attached to them are also picked up. If they are bound in a book, the rest of the book comes along too. (Not so the rhinoceros!) In one respect, our normal experience of documents — and experience therewith — is not disproven by extraneous examples comprising horned, savannah-dwelling fauna or hypothetical prison experiments. The problem with the universality of this view, for textual scholars, is the existence of digital documents.

These undoubtedly pose a problem to the assertion that documents are, in the first instance, selected in their physical extensity. Such questions are particularly important in the realms of digital forensics, particularly (in our case) for the scholarly editing of ‘born digital’ documents (in particular, see Ries 2018, Kirschenbaum 2013, 2014 on this subject).

To begin with, what exactly do we mean by the ‘physical extensity’ of a digital document? This is a question that Kirschenbaum poses: “Today, the conceit of a ‘primary record’ can no longer be assumed to be coterminous with that of a “physical object.” (Kirschenbaum 2013) This, as suggested in Chapter 1, should not — at least in the terms of this analysis — be taken as a rejection of the idea of digital documents as ‘physical’. Digital documents are not ‘virtual’; they are inscribed on a physical medium, like all documents. The physicality in

this case might be a hard-drive, or USB stick, or a genuinely ‘floppy’ 5¼ inch floppy disk. And yet, as Kirschenbaum argues, the physical material is not the digital document. This, I think, may be interpreted in the following way: that the physical medium and the digital document are not co-extensive. The question, then, is how do we (as editors) come to *know* the that there is a document there at all — or, better, how do we establish the extensity of a digital document, given that it does not (unlike a physical document) coincide with the extensity of the physical medium? We return, essentially, the logic outlined above: by what mechanism do we grasp the document?

To begin with, let us consider how digital documents are stored on a disk. It is typical to consider digital documents as encoded as a series of binary ones and zeroes (Smith 1998), which are used in turn to encode the patterns that correspond to, for example, ASCII text. However, it should be recognised equally that binary ones and zeroes are themselves an abstraction from the physical medium, where they are in turn physically recorded (for instance on a hard disk) as blocks of magnetically charged or uncharged segments of material. Undoubtedly, these can be ‘read’, most obviously by the reading head of the disk, which converts (that is, de/records and de/reterritorialises them) as binary signals.

While this mechanism undoubtedly makes ‘reading’ disk a rather different experience to reading physical documents, the abstract mechanism is the same. There is no magic about it; and, while there is undoubted complexity, it is not beyond analysis. Indeed, it is a great deal easier to fathom than the firing of neurons in the human brain provoked by black marks on a white page.²⁵⁴ While the charged/uncharged segments of disk are not immediately tractable

²⁵⁴ Thus, when Smith argues of digital documents that ‘there is, essentially, no object that exists outside of the act of retrieval’ (Smith 1998) the same may in fact be said of all

as language, they are also not random but meaningful. They encode information just as ink on paper does — though not, evidently, in the *same* way (a difference of *coding*). Accordingly, they are not closed to interpretation — even, which may sound strange, without a computer.

Of course, given the size and imperceptibility to sight of magnetically charged segments of disk, some technology intervening — for instance, the reading head of a disk. That said (and as argued in Chapter 1) the human eye is as much an intervening technology between ink-on-paper and the brain, or the human finger between braille-on-the-page and the brain. The point is that, given a list of extracted ones and zeroes, there is evidently something that can be interpreted. And, moreover, this is an interpretation that is not beyond the capacity (though it is almost certainly beyond the patience) of a skilled cryptographer. If we imagine a small disk (a few hundred bytes) carrying the encoded ones and zeroes representing a short text using ASCII encoding, the difficulty does not seem unsurmountable. Even without knowing the ASCII codes for letters, an editor can experiment: that is, testing hypotheses to see whether a given interpretation produces coherence. This is the method behind all interpretation, as discussed in Chapter 2. Of course, as with all ‘expertise’, an editor who knows the ASCII codes can shortcut the process of trial and error, but it is only ever a shortcut: the proof is in the production of sense. Deleuze and Guattari’s description of the “[deduction] of metastable molecular or quasi-molecular units (substances) upon which it

documents. There is only a higher degree of apparent self-evident coincidence between the physical medium and the ‘document’ in the case of paper, and an apparent ‘naturalness’ to reading writing on pieces of paper. As the other examples earlier in this appendix suggest, such self-evidence is only ever relative: the ‘act of retrieval’ is always a requirement in determining the document.

imposes a statistical order of connections and successions” (Deleuze and Guattari, 1987, 46) — which was never really a metaphor — is in this case really quite literal.

In this case, where there is nothing on the disk but a short text, the process of extracting meaningful text serves to delimit the document itself. Its extensity is the assemblage of disk, editor, and ASCII codes. From this point, we may envisage a more complex situation: multiple ‘documents’ (plain text ASCII files) stored on a disk. In this case, analysis of the binary on the disk would reveal (by, for instance, operating system-produced metadata) the presence of these two documents. In this instance, it is the introduction of this metadata (which denotes a boundary) into the document-assemblage that produces two documents.²⁵⁵

What distinguishes this mechanism of document-delimitation from the ‘grasping’ of a paper document in its purely physical extensity is the *order* in which the mechanisms are applied. In the case of a paper document, it is broadly the case that ‘physical grasping’ precedes further analysis. (Even then, as the prisoner’s non-confession, and a piece of paper attached to a rhinoceros suggest, this is hardly universal.) In the case of digital documents, there is a more complex negotiation: a flitting back and forth between analysis, of the ‘possible-document’ as *intensity* and the *extensive* boundaries that might or might not be drawn. The only difference is that physically picking up a document *in general* shortcuts this negotiation and makes the process broadly a linear one (physical grasping of the extensive document preceding

²⁵⁵ By way of parallels, we might imagine a single piece of paper upon which two ‘documents’ were written. In one sense, there is a single document (the physical extensity of the page); in another, two documents, as distinguished by, for instance, different handwriting, different subject matter, various ‘closures’ at a textual level (as one frequently finds in the *Letters of 1916* collection, two letters on the same page, with a ‘Yours sincerely...’ marking the division between the two). The difference lies only in the mechanism of ‘grasping’ the document.

interaction with the document-as-intensity — that is, interpretation).²⁵⁶ But, as Footnote 3 suggests, this is only ever in general, and the initial physical extensity might be undone by later interpretations. Picking up a piece of paper is as much an interpretative *function* — interpretation of a document by its susceptibility to being physically held by the human hand — as any other kind of interpretation. What matters, as argued in Chapter 3, is the order in which interpretative functions are applied, and which one has the last word. For this is what produces a hierarchy of representation.

The above examples have focused on rather trivial examples intended to illustrate the general idea. Undoubtedly, digital documents are more complex, and are mediated by a variety of mechanisms (at a physical level — hardware —, and via the interaction of various levels of software: operating systems, metadata, encoding formats... all the way up to, as Ries notes, the operation of specific word processing software: see Ries 2018). The operation of these various mechanisms produces, in many cases, contrives to produce complex textual (and, indeed, documentary) situations.

For one thing, the ease with which a digital document may be copied ‘perfectly’ in its entirety suggests a greater degree of complexity than a ‘material’ document. If the material of the disk is different, to what degree may it said to be the ‘same’ document (assuming, as argued early, that a digital document has no material embodiment)? At the most base level, it is, of course, not the same document, in the same way that two printed copies of the same edition of the same book are not the same document. But, as in all cases, ‘sameness’ is really a matter of interpretation — in this case, how we view the mechanism of copying. Two ‘identical’ copies

²⁵⁶ Perhaps this is too strong a distinction. Picking up a piece of paper one sees on a table is an experiment, after a fashion — just one quite likely to succeed. But it might fail. The piece of paper one sees may be a hologram.

of a printed book *are* the same if we view the mechanism of copying to be sufficiently accurate for the purposes of our interpretation. Accordingly, we may suggest a spectrum: a scribe copying a manuscript copies less accurately than print; and a computer may copy a file with very high accuracy (but even then, not necessarily infallibly). What a computer affords us is a relatively easy way to verify the validity of a copy of a document (one might, for instance, produce a checksum of the file).²⁵⁷ This is, of course, just as much an interpretation as any other; albeit one with a very high likelihood of producing the same result time after time.

Another example of digital textual complexity (as discussed in Ries 2018) arises from ‘auto-save’ functionality in word processing packages such as Microsoft Word. This feature automatically saves — somewhere in the file — previous versions of a document, which may be recovered by looking at the code produced by the program, even if they are not there on the surface. By ‘surface’, in this case, I am thinking of the visual representation of the page in Word, when one opens the file using Word. Previous text, which has been auto-saved and subsequently deleted, is not there. But this is only when one ‘grasps’ the document using one particular mechanism (opening it in Microsoft Word). If one uses some other tool to view the underlying code of the document (at a level, incidentally, far above the binary code), the ‘deleted’ sections may be viewable. In the case of more recent, XML-based versions of Microsoft Word files (.docx rather than .doc), this is particularly easy. The question²⁵⁸ in this case is: are there one or many documents?

²⁵⁷ Checksums — a cryptographic hash of a file’s contents — are often used when downloading software to ensure that the file downloaded is identical to that which the distributor of the software intends (and has thus not been maliciously interfered with).

²⁵⁸ As posed by Ries at the Gottingen ‘Text Analysis’ summer school of 2014 — a couple of months into my PhD — and which quite flummoxed me at the time.

Following the traditional nomenclature of textual criticism, we might be tempted to conclude that there are multiple documents, each comprising a particular unique version of a text (or, indeed, work). It is not that there is, necessarily, anything particularly wrong with such a view. But such a view requires a particular sequence of interpretations — that variants of a text comprise, in and of themselves, *versions*. As genetic criticism would suggest, by contrast, the idea of a version is far more fluid. A single correction on the same page (the same document, given a certain interpretation) can be seen as producing a new version of the text. But these attempts to draw parallels are only useful insofar as they accurately represent the situation. In this case, we may draw a number of conclusions, which may well be contradictory. For one thing, while we may say that there is a single *file* on the disk (whether or not it is one or more document), this is dependent on an operating system that views it as one file and a program that opens it as one file.²⁵⁹ A single Microsoft Word file, when opened in Microsoft Word, is a single file (and, logically, a single document). Peer inside the file (in the case of a .docx file, one may simply change the file extension to that of a .zip file and ‘unzip’ it) and one finds a number of files. As to whether the saved variants comprise different documents again depends on how they are opened. If the file is opened in such a way that the current text and all the revisions are visible, it could be called one document and a genetic-criticism kind of revision pertains. In other circumstances, one might be have cause to interpret the presence of multiple documents.

Given the above, it seems legitimate to ask whether there is any utility at all in the question of whether there are one or multiple documents in such circumstances. The answer is, of course, that it depends, given the mechanisms by which the one or many documents are

²⁵⁹ The tendency of some programs to produce temporary files all the time illustrates the point that what a program might think of as a single file is seen by the operating system as a set of temporary (and often hidden) files.

‘grasped’ as such — in other words, it is entirely dependent on interpretation. As a result, it hardly seems gratuitous to wonder whether documents exist at all, or — slightly less absolutely — whether the category of ‘document’ serves a useful purpose. Renear (2010), approaching the question from a quite different direction,²⁶⁰ questions whether documents actually exist at all. In answer this, I would suggest that it depends. Certainly, to elevate the document to a kind of universal ontological category would seem fraught with the danger of self-contradiction or, at least, obvious counterexample. (Actually, I would go so far as to suggest that such a thing is impossible.) As to whether it is useful to talk of documents is a different matter. Calling something a document as a result of having delimited it — and the function that is thereby served by the document that is ‘created’ — is clearly useful. Following Foucault’s notion of an ‘author function’, it is tempting to assert that there is equally a ‘document function’, something that plays the role of a document in a given editorial situation. This is obviously a necessity, otherwise textual scholarship would have nothing at all to operate on. What plays the role of a ‘document’ is, however, unique in any given process of textual transmission, and it is always something that is created by the editor, out of interpretation. The document as an entity is a product of whatever ‘mechanism of grasping’ is brought to bear on it such that it produces a text-bearing entity. Perhaps, for the avoidance of argument, it is better to call it a ‘source assemblage’.

²⁶⁰ Over the course of two papers, Renear makes the case that documents cannot be modified (without creating a new document) and, as a suggested resolution to the impasse thus created, argues that the notion of a document (whether they exist or not) is essentially redundant. His argument is, in my view, somewhat reductive, beginning as it does with his definition of a document as nothing more than a specific sequence of characters. (It is also fair to say that his approach, rooted in analytical philosophy, is rather at odds with the kind of Deleuzian reading attempted in this thesis. We are, I think, inclined to disagreement from the outset!)

Appendix B: List of Abbreviations and Technical Terminology

ASCII – American Standard Code for Information Interchange. A basic character-encoding standard (comprises numbers, upper- and lower-case letters and limited punctuation). See *Unicode*.

BaseX — a ‘native XML’ database. This type of database stores its data as hierarchical XML nodes, rather than as tables. See *eXist-DB*.

C++ — an object-orientated derivative of the C programming language.

Canvas, SVG, WebGL — types of graphical representation for web browsers. Canvas and SVG are very high level, with substantial browser support; WebGL, which can be used to render complex three-dimensional graphics, is more powerful, and, accordingly, lower-level. JavaScript – a high-level programming language, the default language for browser-based programming.

CollateX — a tool for collating and identifying differences between versions of texts. Developed by Dekker *et al.*

DDE — Digital Documentary Edition (a rarely-used abbreviation; normally articles will spell it out before shortening)

DOT, GML, GraphML, GEFX – formats for representing graphs; the former two are plain-text based formats; the latter use XML.

EEBO — Early English Books Online

eXist-DB — a ‘native XML’ database. See *BaseX*.

Gephi – a popular tool for viewing and interacting with graph data.

GraphML — an XML-based mark-up language for encoding graphs

HTML — HyperText Markup Language. A mark-up language derived from SGML (Standard General Markup Language) developed by Tim Berners-Lee as the standard document type for the World Wide Web. Now a standard administered by the WC3 (World Wide Web Consortium), though developed by a number of working groups and browser vendors

IndexedDB API — a browser-based ‘lightweight’ database for storing and querying small amounts of data in the user’s web browser.

LDA — Latent Dirichlet allocation; a topic-modelling algorithm

LMNL — a mark-up language (broadly similar to XML) allowing overlapping hierarchies to be directly encoded in a text (without the workarounds required in XML). Its use is substantially less than XML, not least due to its additional complexity and comparative lack of tools.

NetworkX, GraphTool, iGraph — three graph layout and analysis libraries for the Python programming language. See *Python*.

Neo4J — a popular ‘graph-oriented’ database.

NoSQL — a generic term for databases that do not fall into the traditional paradigm of table-based ‘relational databases’ (these may take a number of forms: document stores, graph databases and XML databases are examples).

OHCO — Ordered Hierarchy of Content Objects. A view of (digital encoding of) text, as an ordered hierarchy of content objects (e.g. chapters contain ordered paragraphs, which contain ordered sentences, which contain ordered words...) argued for by DeRose et al. Now typically viewed as inadequate for representing textual complexity. (XML is, by default, an embodiment of this model, though many work-arounds for its limitations have been proposed.)

Omeka — A php-based application for building online exhibitions.

oXygen — An XML editor. See *XML*.

Python — A high-level programming language, widely used in academia, for Natural Language Processing, Machine Learning, etc., as well as elsewhere (as a web development language) and a general-purpose scripting language.

RDF — Resource Description Framework. A graph model for describing entities and relationships in a graph model; the foundation of the ‘semantic web’.

Ruby on Rails — A web framework written in the Ruby programming language.

Scripto — An Omeka plugin for transcription of texts. See *Omeka*.

SPARQL — Query language for querying RDF databases.

SQL — Structured Query Language. A family of (broadly similar) query languages for querying relational databases.

TAG, TAGQL — Text-As-Graph, Text-as-Graph-Query-Language. A hypergraph model for encoding text, developed by Dekker and Birnbaum; it is designed to overcome problems with XML (especially issues of hierarchy). A very recent development with (undoubted) potential, but little uptake at present. TAGQL is its related query language.

VivaGraph, Sigma.js — JavaScript-based graph rendering and interactivity libraries

Angular, React, Vue – three popular client-side (i.e. running in the user’s web browser) frameworks

TeX, LaTeX — a typesetting system developed originally by computer scientist Donald Knuth

SGML — Standard Generalized Markup Language. A declarative mark-up language, descended from IBM’s General Markup Language

Unicode – A character-encoding standard. Comprises a substantially wider range of characters (especially for non-English languages) than ASCII.

XML — eXtensible Markup Language. A declarative mark-up metalanguage, also derived from SGML (a restricted subset of syntax). The language may be used to produce other mark-up languages ('vocabularies'). The TEI (Text Encoding Initiative) has used XML since version P4. HTML has also been implemented using XML (as XHTML), as have many other languages for processing XML (e.g. XSLT).

XSLT — eXtensible Stylesheet Language Transformations. An XML-based language for transforming XML documents (typically into other XML documents, or HTML)

XPath — a language for accessing hierarchical nodes in an XML document tree (see OHCO, XML)

XQuery — a query language (and functional programming language) for querying XML. The main language used for developing applications in eXist-DB.

Works cited

- Aarseth, Espen J. *Cybertext: Perspectives on Ergodic Literature*. John Hopkins University Press, 1997.
- Adkins, Brent. *Deleuze and Guattari's A Thousand Plateaus: A Critical Introduction and Guide*. Edinburgh University Press, 2015.
- Apollon, Daniel, et al., editors. *Digital Critical Editions*. University of Illinois Press, 2014.
- Andrews, Tara L., and K. U. Leuven. 'The Third Way: Philology and Critical Edition in the Digital age'. *Variants: The Journal of the European Society for Textual Scholarship*, vol. 10, 2013, pp. 61–76.
- Barnard, David T., et al. 'Generalized Markup for Literary Texts'. *Literary and Linguistic Computing*, vol. 3, no. 1, Jan. 1988, pp. 26–31.
- Barone, James E. 'Comparing Apples and Oranges: A Randomised Prospective Study'. *BMJ*, vol. 321, no. 7276, Dec. 2000, pp. 1569–70.
- Barthes, Roland. 'The Death of the Author'. *Image-Music-Text*, vol. 145, 1977.
- Beall, Jeffrey. 'The Weaknesses of Full-Text Searching'. *The Journal of Academic Librarianship*, vol. 34, no. 5, 2008, pp. 438–444.
- Bleier, Roman, and Richard Hadden. 'Capturing the Crowd-Sourcing Process: revisionDesc and Storing the Stages of Crowd-Sourced Transcriptions in TEI'. *TEI Conference and Members Meeting 2016*, 2016.
- Bode, Katherine. 'The Equivalence of "close" and "distant" Reading; Or, toward a New Object for Data-Rich Literary History'. *Modern Language Quarterly*, vol. 78, no. 1, 2017, pp. 77–106.
- Bogue, Ronald. *Deleuze and Guattari*. Routledge, 2008.

- Boot, P. *Mesotext. Digitised Emblems, Modelled Annotations and Humanities Scholarship*. Amsterdam University Press, 2009.
- Bordalejo, Barbara. 'The Texts We See and the Works We Imagine: The Shift of Focus of Textual Scholarship in the Digital Age'. *Ecdotica*, vol. 10, Dec. 2013, pp. 64–76.
- Bowers, Fredson. *Textual and Literary Criticism*. CUP Archive, 1966.
- Bowman, Frank Paul. 'Genetic Criticism'. *Poetics Today*, vol. 11 no. 3, pp. 627–646.
- Bradley, J. 'Texts into Databases: The Evolving Field of New-Style Prosopography'. *Literary and Linguistic Computing*, vol. 20, no. 1, Jan. 2005, pp. 3–24.
- Bricmont, Jean, and Alan Sokal. *Intellectual Impostures*. Profile Books, 2011.
- Broughton, William Michael. 'Encoding Text as a Nexus of Perspective Hierarchies: Case Studies in XML, LMNL, and RDF'. *TEI Conference and Members Meeting 2016*, TEI Consortium, 2016.
- Bryant, John. *The Fluid Text: A Theory of Revision and Editing for Book and Screen*. University of Michigan Press, 2002.
- Burghart, Marjorie. *TEI Critical Edition Toolbox*, 2015, <http://ciham-digital.humanum.fr/teitoolbox/index.php>.
- Burnard, Lou. 'Tools and Techniques for Computer-Assisted Text Processing'. *Computers and Written Texts*, edited by Christopher S. Butler, Blackwell, 1992, pp. 1–28.
- Burnett, Kathleen. 'Toward a Theory of Hypertextual Design'. *Postmodern Culture*, vol. 3, no. 2, 1993.
- Buzzetti, Dino. 'Digital Editions and Text Processing'. *Text Editing, Print and the Digital World*, edited by Marilyn Deegan and Kathryn Sutherland, Ashgate, 2009.
- Buzzetti, Dino, and Jerome McGann. 'Electronic Textual Editing: Critical Editing in a Digital Horizon'. *Electronic Textual Editing*, edited by John Unsworth et al., TEI Consortium, 2006, http://www.tei-c.org/About/Archive_new/ETE/Preview/mcgann.xml#body.1_div.2.

- Cappon, Lester J. 'A Rationale for Historical Editing: Past and Present'. *The William and Mary Quarterly*, vol. 23, no. 1, 1966, pp. 56–75.
- Caton, Paul. 'On the Term "Text" in Digital Humanities'. *Literary and Linguistic Computing*, vol. 28, no. 2, June 2013, pp. 209–20.
- Causser, T., et al. 'Transcription Maximized; Expense Minimized? Crowdsourcing and Editing The Collected Works of Jeremy Bentham'. *Literary and Linguistic Computing*, vol. 27, no. 2, June 2012, pp. 119–37.
- . "'Making Such Bargain": Transcribe Bentham and the Quality and Cost-Effectiveness of Crowdsourced Transcription'. *Digital Scholarship in the Humanities*, 2018.
- Cerquiglini, Bernard, *In Praise of the Variant: a Critical History of Philology*, John Hopkins University Press, 1999.
- Ciula, A., et al. 'Expressing Complex Associations in Medieval Historical Documents: The Henry III Fine Rolls Project'. *Literary and Linguistic Computing*, vol. 23, no. 3, Sept. 2008, pp. 311–25.
- Clausner, Christian, et al. 'The ENP Image and Ground Truth Dataset of Historical Newspapers'. *Document Analysis and Recognition (ICDAR), 2015 13th International Conference on*, IEEE, 2015, pp. 931–935.
- Coyle, Karen. 'One Word: Digital'. *The Journal of Academic Librarianship*, vol. 32, no. 2, 2006, pp. 205–207.
- Cummings, James. *Customising the TEI for Your Project*.
<http://slides.com/jamescummings/customisingtei>. Accessed Feb. 2018.
- Dahlström, Mats. 'How Reproductive Is A Scholarly Edition?' *Literary and Linguistic Computing*, vol. 19, no. 1, 2004.
- . 'Critical Editing and Critical Digitization'. *Text Comparison and Digital Creativity*, Brill, 2010, pp. 79–97.

- Dahlström, Mats, et al. “As We May Digitize”—Institutions and Documents Reconfigured’. *Liber Quarterly*, vol. 21, no. 3/4, 2012, pp. 455–474.
- de Biasi, Pierre-Marc. ‘Edition horizontale, édition verticale. Pour une typologie des éditions génétiques (le domaine français 1980-1995).’ *Éditer des manuscrits. Archives, complétude, lisibilité*, edited by Georges Benrekassa, Jacques Neefs, and Béatrice Didier, 1996.
- Deegan, Marilyn, and Kathryn Sutherland. *Transferred Illusions: Digital Technology and the Forms of Print*. Ashgate Publishing, Ltd., 2009.
- Dekker, Ronald Haentjens, and David J. Birnbaum. ‘It’s More than Just Overlap: Text As Graph.’ *Balisage: The Markup Conference*, 2017. Web.
<https://www.balisage.net/Proceedings/vol19/print/Dekker01/BalisageVol19-Dekker01.html>.
- DeLanda, Manuel. *Deleuzian Ontology: A Sketch*. 2006, <http://www.situation.ru/>
- Deleuze, Gilles. *Difference and Repetition*. 1968. Translated by Paul Patton A&C Black, 2004.
- . *Logic of Sense*. 1969. Bloomsbury Publishing, 2004.
- . *Desert Islands and Other Texts, 1953-1974*. Semiotext(e), 2004.
- Deleuze, Gilles, and Félix Guattari. *A Thousand Plateaus: Capitalism and Schizophrenia*. Translated by Brian Massumi, University of Minnesota Press, 1987.
- Deppman, Jed, Daniel Ferrer, and Michael Groden. 2004. ‘Introduction: A Genesis of French Genetic Criticism.’ *Genetic Criticism: Texts and Avant-Textes*, edited by Jed Deppman, Daniel Ferrer, and Michael Groden, pp. 1-16. Philadelphia: University of Pennsylvania Press.
- Derrida, Jacques. ‘Archive Fever: A Freudian Impression’. *Diacritics*, vol. 25, no. 2, 1995, p. 9.

- DeRose, Steven J., et al. 'What Is Text, Really?' *ACM SIGDOC Asterisk Journal of Computer Documentation*, vol. 21, no. 3, 1990, pp. 1–24.
- Drosnin, Michael. *The Bible Code*. W&N, 1997.
- Drucker, Johanna. *Performative Materiality and Theoretical Approaches to Interface*. Vol. 7, no. 1, 2013. *Digital Humanities Quarterly*,
<http://www.digitalhumanities.org/dhq/vol/7/1/000143/000143.html>.
- . 'Distributed and conditional documents: conceptualizing bibliographical alterities'. *Matlit*, vol. 2, no. 1, 2014, pp. 11–29.
- Eco, Umberto. *From the Tree to the Labyrinth*. Harvard University Press, 2014.
- Eggert, Paul. 'Document and Text: The "Life" of the Literary Work and the Capacities of Editing'. *Text*, 1994, pp. 1–24.
- . 'Text-Encoding, Theories of the Text, and the "Work-Site"'. *Literary and Linguistic Computing*, vol. 20, no. 4, Nov. 2005, pp. 425–35.
- . *Securing The Past*. Cambridge University Press, 2009.
- . 'Text as Algorithm and Process'. *Text and Genre in Reconstruction: Effects of Digitalization on Ideas, Behaviours, Products and Institutions*, edited by Willard McCarthy, 2010, <https://books.openedition.org/obp/660?lang=en>
- Evans, Fred. 'Deleuze, Bakhtin, and the "Clamour of Voices"'. *Deleuze Studies*, vol. 2, no. 2, 2008, pp. 178–88. *Edinburgh University Press Journals*.
- Faulhaber, Charles. 'Textual Criticism in the 21st Century'. *Romance Philology*, vol. 45, no. 1, 1991, pp. 123–236.
- Finneran, Richard J. *The Literary Text in the Digital Age*. University of Michigan Press, 1996.
- Fischer, Franz. 'All Texts Are Equal, But... Textual Plurality and the Critical Text in Digital Scholarly Editions.' *Variants: The Journal of the European Society for Textual Scholarship*, vol. 10, 2012, pp. 77–92.

- Folsom, Ed, and Kenneth M. Price, editors. *The Walt Whitman Archive*.
<https://whitmanarchive.org/>. Accessed 21 Mar. 2017.
- Foucault, Michel. 'What Is an Author?', 1969. *Foucault Reader*, edited by Paul Rainbow, Pantheon, 1984, pp. 101–20.
- . *Archaeology of Knowledge*. Routledge, 1972.
- Gabler, Hans Walter. 'Textual Criticism and Theory In Modern German Editing'. *Contemporary German Editorial Theory*, edited by Hans Walter Gabler et al., University of Michigan Press, 1995.
- . 'The Primacy of the Document in Editing'. *Ecdotica*, vol. 4, 2007, pp. 197–207.
- . 'Theorizing the Digital Scholarly Edition'. *Literature Compass*, vol. 7, no. 2, Feb. 2010, pp. 43–56.
- . 'Beyond Author-Centricity in Scholarly Editing'. *Journal of Early Modern Studies*, vol. 1, no. 1, 2012.
- . 'Networking a Scholarly Edition—Networking Ulysses as a Digital Text and Research Site'. *Breac: A Digital Journal of Irish Studies*, 2015,
<http://breac.nd.edu/articles/networking-a-scholarly-editionnetworking-ulysses-as-a-digital-text-and-research-site/>
- Gochenour, Phil. 'Nodalism'. *Digital Humanities Quarterly*, vol. 5, no. 3, 2011.
- Gooding, P., et al. 'The Myth of the New: Mass Digitization, Distant Reading, and the Future of the Book'. *Literary and Linguistic Computing*, vol. 28, no. 4, 2013, pp. 629–39.
- Greetham, D. C. *Textual Scholarship: An Introduction*. Garland Publishing, 1992.
- . *Scholarly Editing: A Guide to Research*. Modern Language Association of America, 1995.
- Greg, W. W. 'The Rationale of Copy-Text'. *Studies in Bibliography*, vol. 3, 1950, pp. 19–36.

- . 'Bibliography - An Apologia'. *Collected Papers of W. W. Greg*, edited by J. C. Maxwell, Clarendon Press, 1966, p. 247.
- Grésillon, Almuth. *Éléments de Critique Génétique. Lire Les Manuscrits Modernes*. Presses universitaires de France, 1994.
- Hadden, Richard. 'Review: The Diary of William Godwin'. *RIDE: A Review Journal for Digital Editions and Resources*, no. 3, 2015, <https://ride.i-d-e.de/issues/issue-3/godwin/>. Accessed Feb 2018.
- Hahn, Trudi Bellardo. 'Mass Digitization'. *Library Resources & Technical Services*, vol. 52, no. 1, 2008, pp. 18–26.
- Harvey, P.D.A. *Editing Historical Records*. The British Library, 2001.
- Harvey, Anthony, et al., editors. *St. Patrick's Confessio: Hyperstack Edition*. Web. <https://www.confessio.ie/#>. Accessed 12 Oct. 2016.
- Hockey, Susan M. *Electronic Texts in the Humanities: Principles and Practice*. Oxford University Press, 2000.
- . 'The History of Humanities Computing'. *Companion to Digital Humanities*, edited by Susan Schreibman et al., Blackwell Publishing Professional, 2004.
- Holley, Rose. 'How Good Can It Get?: Analysing and Improving OCR Accuracy in Large Scale Historic Newspaper Digitisation Programs'. *D-Lib Magazine*, vol. 15, no. 3/4, Mar. 2009.
- Huitfeldt, Claus. "Multi-Dimensional Texts in a One-Dimensional Medium." *Computers and the Humanities*, vol. 28, no. 4/5, 1994, pp. 235–241.
- . 'Philosophy Case Study'. *Electronic Textual Editing*, edited by John Unsworth et al., TEI Consortium, 2006, http://www.tei-c.org/About/Archive_new/ETE/Preview/kiernan.xml.
- Jackson, Allyn. 'Book Review: The Bible Code'. *Notices of the American Mathematical Society*, vol. 44, no. 8, 1997, pp. 935–39.

- Jane Austen's Fiction Manuscripts: A Digital Edition*, edited by Kathryn Sutherland, 2010.
 Web: <http://www.janeausten.ac.uk>.
- Jockers, Matthew L. *Macroanalysis: Digital Methods & Literary History*. University of Illinois Press, 2013.
- Juola, Patrick, and Stephen Ramsay. 'Six Septembers: Mathematics for the Humanist'. *Zea E-Books*, Apr. 2017, <https://digitalcommons.unl.edu/zeabook/55>.
- Kaiser, Max. 'Putting 600,000 Books Online: The Large-Scale Digitisation Partnership between the Austrian National Library and Google'. *LIBER Quarterly*, vol. 21, no. 2, Jan. 2012.
- Kaufman, Micki, et al. 'Visualizing Futures of Networks in Digital Humanities Research'. *Digital Humanities 2017*, 2017, <https://dh2017.adho.org/abstracts/428/428.pdf>.
- Kelly, Aodhán. 'Tablet Computers for the Dissemination of Digital Scholarly Editions.' *Manuscrita*, vol. 28, 2015.
- Kirschenbaum, Matthew. 'The .txtual Condition: Digital Humanities, Born-Digital Archives, and the Future Literary'. *Digital Humanities Quarterly*, vo. 7, no. 1, 2013, <http://www.digitalhumanities.org/dhq/vol/7/1/000151/000151.html>.
- Kiernan, Kevin. 'Digital Facsimiles in Editing'. *Electronic Textual Editing*, edited by John Unsworth et al., TEI Consortium, 2006, http://www.tei-c.org/About/Archive_new/ETE/Preview/kiernan.xml.
- Kline, Mary-Jo. *A Guide to Documentary Editing*. Johns Hopkins University Press, 1998.
- Kretzschmar, William J. 'Large-Scale Humanities Computing Projects: Snakes Eating Tails, or Every End Is a New Beginning?' *Digital Humanities Quarterly*, vol. 3, no. 2, 2009, <http://www.digitalhumanities.org/dhq/vol/3/2/000038/000038.html>.
- Landow, George P. *Hypertext: The Convergence of Contemporary Critical Theory and Technology*. Johns Hopkins University Press, 1992.

- . *Hypertext 3.0: Critical Theory and New Media in an Era of Globalization*. Johns Hopkins University Press, 2006.
- Latour, Bruno. 'On Recalling ANT'. *The Sociological Review*, vol. 47, 1999, pp. 15–25.
- . *Reassembling the Social an Introduction to Actor-Network-Theory*. Oxford University Press, 2005.
- Lebrave, Jean-Louis. 'Rough Drafts: A Challenge to Uniformity in Editing'. *TEXT: Transactions of the Society for Textual Criticism*, vol. 3, 1987, pp. 135–42.
- LMNL Markup Dot Org*. 2008, <http://lmln-markup.org/>. Accessed Feb 2018.
- Logan, George, et al. 'Computer Based Publication of Critical Editions: Some General Considerations and a Prototype'. *Computers and the Humanities. Today's Research Tomorrow's Teaching. Conference Pre-Prints.*, University of Toronto Press, 1986, pp. 318–326.
- Mak, Bonnie. 'Archaeology of a Digitization'. *Journal of the American Society for Information Science and Technology*, (Preprint), 2013.
- Mandell, Laura. 'How to Read a Literary Visualisation: Network Effects in the Lake School of Romantic Poetry'. *Digital Studies/Le Champ Numérique*, vol. 3, no. 2, Aug. 2013.
- . *Visualizing Gender Complexity - A Word Says More than 1000 Pictures: Dynamic Data Visualization and Exploration for the Humanities*. Universität Hamburg, <https://lecture2go.uni-hamburg.de/l2go/-/get/v/19498>. Accessed Feb 2018.
- Mann, Thomas. 'Will Google's Keyword Searching Eliminate the Need for LC Cataloging and Classification?' *AFSCME*, 2005, <http://www.guild2910.org/searching.htm>.
- Martens, Gunter. 'What Is a Text? Attempts at Defining a Central Concept in Editorial Theory.' *Contemporary German Editorial Theory*, edited by Hans Walter Gabler et al., University of Michigan Press, 1995, pp. 209–232.
- Massumi, Brian. *A User's Guide to Capitalism and Schizophrenia: Deviations from Deleuze and Guattari*. MIT Press, 1992.

- McCarty, Willard. 'Modeling: A Study in Words and Meanings'. *A Companion To Digital Literary Studies*, edited by Susan Schreibman and Ray Siemens, Blackwell, 2008, <http://www.digitalhumanities.org/companion/view?docId=blackwell/9781405103213/9781405103213.xml&chunk.id=ss1-3-7>.
- McGann, Jerome J. *A Critique of Modern Textual Criticism*. University Press of Virginia, 1983.
- . *The Textual Condition*. Princeton University Press, 1991.
- . 'The Rationale of Hyper Text'. *Text*, vol. 9, 1996, pp. 11–32.
- . *Radiant Textuality: Literature after the World Wide Web*. Palgrave, 2001.
- . 'Database, Interface, and Archival Fever'. *PMLA*, 2007, pp. 1588–1592.
- . 'Electronic Archives and Critical Editing'. *Literature Compass*, vol. 7, no. 2, Feb. 2010, pp. 37–42.
- . *A New Republic of Letters*. Harvard University Press, 2014.
- McGarry, Fearghal. *The Rising: Ireland, Easter 1916*. Oxford University Press, 2010.
- McKenzie, D. F. *Bibliography and the Sociology of Texts*. Cambridge University Press, 1999.
- Mishra, Sundesh. 'Heterographies and Other Assemblages'. *BORDERLANDS e-Journal*, vol. 6, no. 2, 2007, http://www.borderlands.net.au/vol6no2_2007/mishra_heterographies.htm.
- Moretti, Franco. *Graphs, Maps, Trees: Abstract Models for a Literary History*. Verso Books, 2005.
- . *Distant Reading*. Verso Books, 2013.
- Myers, Victoria, et al., editors. *The Diary of William Godwin*. Oxford Digital Library, 2010, <http://godwindiary.bodleian.ox.ac.uk/index2.html>.
- Nail, Thomas. 'What Is an Assemblage?' *SubStance*, vol. 46, no. 1, Mar. 2017, pp. 21–37.
- Neuber, Frederike. 'Stefan George Digital: Exploring Typography In A Digital Scholarly Edition'. *Digital Humanities*, 2016, <http://dh2016.adho.org/abstracts/321>.

- Novick, Ben. 'Postal Censorship in Ireland, 1914–16'. *Irish Historical Studies*, vol. 31, no. 123, May 1999, pp. 343–56.
- Nunberg, Geoff. 'Google Books: A Metadata Train Wreck'. *Language Log*, 2009, <http://languagelog.ldc.upenn.edu/nll/?p=1701>. Accessed Feb 2018.
- Palmer, Carole L. 'Thematic Research Collections'. *Companion to Digital*, edited by Susan Schreibman et al., Blackwell Publishing Professional, 2004, http://nora.lis.uiuc.edu:3030/companion/view?docId=blackwell/9781405103213/9781405103213.xml&chunk.id=ss1-4-5&toc.depth=1&toc.id=ss1-4-5&brand=9781405103213_brand.
- Panagia, Davide. 'Intermedial Relations: A Critique of Reification Theory'. *SubStance*, vol. 46, no. 1, Mar. 2017, pp. 90–110.
- Pierazzo, Elena. 'A Rationale of Digital Documentary Editions'. *Literary and Linguistic Computing*, vol. 26, no. 4, Dec. 2011, pp. 463–77.
- . 'Digital Documentary Editions and the Others'. *Scholarly Editing*, no. 4, Mar. 2013.
- . *Digital Scholarly Editing: Theories, Models and Methods*. Ashgate Publishing, Ltd., 2015. [Advanced copy, HAL Archives-Ouverts.fr, 2014]
- . 'TEI: XML and Beyond'. *Text Encoding Initiative Conference and Members Meeting 2015*, TEI Consortium, 2015a.
- Pierazzo, Elena, and Peter Stokes. 'Putting the Text Back into Context: A Codicological Approach to Manuscript Transcription'. *Kodikologie Und Paläographie Im Digitalen Zeitalter 2 - Codicology and Palaeography in the Digital Age 2. Schriften Des Instituts Für Dokumentologie Und Editorik*, 3, BoD–Books on Demand, 2011, pp. 397–429.
- Price, Kenneth M. 'Electronic Scholarly Editions'. *A Companion To Digital Literary Studies*, edited by Susan Schreibman and Ray Siemens, Blackwell, 2008, <http://www.digitalhumanities.org/companion/view?docId=blackwell/9781405148641/>

9781405148641.xml&chunk.id=ss1-6-5&toc.depth=1&toc.id=ss1-6-5&brand=9781405148641_brand.

---. *Edition, Project, Database, Archive, Thematic Research Collection: What's in a Name?* Vol. 3, no. 3, 2009. *Digital Humanities Quarterly*, <http://digitalhumanities.org/dhq/vol/3/3/000053/000053.html>.

Radford, Gary P. 'Positivism, Foucault, and the Fantasia of the Library: Conceptions of Knowledge and the Modern Library Experience'. *The Library Quarterly*, vol. 62, no. 4, 1992, pp. 408–24.

Rehbein, Malte, and Christiane Fritz. 'Hands-On Teaching Digital Humanities: A Didactic Analysis of a Summer School Course on Digital Editing'. *Digital Humanities Pedagogy: Practices, Principles and Politics*, edited by Brett D. Hirsch, Open Book Publishers, 2012, pp. 47–78.

Renear, Allen H., et al. *Refining Our Notion of What Text Really Is: The Problem of Overlapping Hierarchies*. 1993, <http://cds.library.brown.edu/resources/stg/monographs/ohco.html>.

Renear, Allen H., Karen Michelle Wickett. 'There Are No Documents'. *Proceedings of Balisage*, vol. 5, 2010, <https://experts.illinois.edu/en/publications/there-are-no-documents>. Accessed May 2019.

Ridge, Mia. 'From Tagging to Theorizing: Deepening Engagement with Cultural Heritage through Crowdsourcing'. *Curator: The Museum Journal*, vol. 56, no. 4, p. 2013.

Ries, Thorsten. 'The rationale of the born-digital *dossier génétique*: Digital forensics and the writing process: With examples from the Thomas Kling Archive'. *Digital Scholarship in the Humanities*, vol. 33, no. 2, 2018, pp. 391–424, <https://doi.org/10.1093/lc/fqx049>

Robinson, Lyn, and Mike Maguire. 'The Rhizome and the Tree: Changing Metaphors for Information Organisation'. *Journal of Documentation*, vol. 66, no. 4, 2010, pp. 604–13.

- Robinson, Peter. 'Creating & Documenting Electronic Texts'. *Creating and Documenting Electronic Texts: A Guide to Good Practice*, AHDS, 1993,
<https://ota.ox.ac.uk/documents/creating/cdet/chap3.html>.
- . 'Where We Are with Electronic Scholarly Editions, and Where We Want to Be'.
Jahrbuch Für Computerphilologie, vol. 5, 2004, pp. 123–43.
- . 'What Text Really Is Not, and Why Editors Have to Learn to Swim'. *Literary and Linguistic Computing*, vol. 24, no. 1, Apr. 2009, pp. 41–52.
- . 'Towards a Theory of Digital Editions'. *The Journal of the European Society for Textual Scholarship*, 2013, p. 105.
- . 'Why Digital Humanists Should Get out of Textual Scholarship'. *SDSE*, 2013,
[https://www.academia.edu/4124828/
SDSE_2013_why_digital_humanists_should_get_out_of_textual_scholarship](https://www.academia.edu/4124828/SDSE_2013_why_digital_humanists_should_get_out_of_textual_scholarship).
- Rockwell, Geoffrey. 'Crowdsourcing in the Humanities: Social Research and Collaboration'. *Collaborative Research in the Digital Humanities*, edited by Marilyn Deegan and Willard McCarty, Ashgate Publishing, Ltd., 2012, pp. 135–54.
- Romanello, Matteo. 'Exploring Citation Networks to Study Intertextuality in Classics'.
Digital Humanities Quarterly, vol. 10, no. 2, 2016.
- Ross, Donald. 'Aids for Editing "Walden"'. *Computers and the Humanities*, vol. 15, no. 3, 1981, pp. 155–62.
- Ruohonen, Keijo. *Graph Theory*. 2013
- Sahle, Patrick. 'Modeling Transcription'. *Knowledge Organisation and Data Modeling in the Humanities*, 2013, <https://datasymposium.wordpress.com/sahle/>. Accessed Feb 2018.
- . *Criteria for Reviewing Digital Editions and Resources*. RIDE, 2014, <http://ride.i-d-e.de/reviewers/catalogue-criteria-for-reviewing-digital-editions-and-resources/>. Accessed Feb 2018.

- . 'What Is a Scholarly Digital Edition?' *Digital Scholarly Editing: Theories and Practice*, edited by Elena Pierazzo and Matthew James Driscoll, 2014, <http://www.openbookpublishers.com/reader/483/#page/46/mode/2up>. Accessed 17 Jan. 2017.
- Sandford, Scott A. 'Apples and Oranges-A Comparison'. *Annals of Improbable Research*, vol. 1, no. 3, 1995, pp. 2–3.
- Savat, David, and Tael Harper. *Media After Deleuze*. Bloomsbury Publishing, 2016.
- Schloen, David, and Sandra Schloen. 'Beyond Gutenberg: Transcending the Document Paradigm in Digital Humanities'. *Digital Humanities Quarterly*, vol. 8, no. 4, 2014.
- Schreibman, Susan. *Versioning Machine 5.0*, 2010, 2016 2002, <http://v-machine.org/>.
- . 'Computer-Mediated Texts and Textuality: Theory and Practice'. *Computers and the Humanities*, vol. 36, no. 3, 2002, pp. 283–293.
- . 'Digital Scholarly Editing'. *Literary Studies in the Digital Age*, edited by Kenneth M. Price and Ray Siemens, Modern Language Association of America, 2013, <http://dlsanthology.commons.mla.org/digital-scholarly-editing/>. Accessed Feb 2018.
- , editor. *Contested Memories: The Battle of Mount Street Bridge*, 2015. <https://mountstreet1916.ie/mt-st-bridge-background/>. Accessed 22 Feb. 2018.
- Schreibman, Susan, et al., editors. *The Woodman Diary: Digital Edition*. 2015, <http://dhprojects.maynoothuniversity.ie/woodman/about-the-project/>.
- Schreibman, Susan, et al. 'Notes from the Transcription Desk: Visualising Public Engagement'. *English Studies*, vol. 98, no. 5, July 2017, pp. 506–25.
- Sculley, D., and B. M. Pasanek. 'Meaning and Mining: The Impact of Implicit Assumptions in Data Mining for the Humanities'. *Literary and Linguistic Computing*, vol. 23, no. 4, Sept. 2008, pp. 409–24.
- Shillingsburg, Peter. 'The Computer as Research Assistant in Scholarly Editing'. *Literary Research Newsletter*, vol. 5, no. 1, 1980, pp. 31–45.

- . *Scholarly Editing in the Computer Age*. University of Michigan Press, 1996.
- . 'How Literary Works Exist: Convenient Scholarly Editions'. *Digital Humanities Quarterly*, vol. 3, no. 3, 2009,
- . 'The Semiotics of Bibliography'. *Textual Cultures: Texts, Contexts, Interpretation*, vol. 6, no. 1, 2011, pp. 11–25.
- . 'Literary Documents, Texts, and Works Represented Digitally'. *Ecdotica: Rivista Di Studi Testuali*, vol. 3, 2013.
- Siemens, R., et al. 'Toward Modeling the Social Edition: An Approach to Understanding the Electronic Scholarly Edition in the Context of New and Emerging Social Media'. *Literary and Linguistic Computing*, vol. 27, no. 4, 2012, pp. 445–61.
- Simon, John. 'Editors and Critics'. *Newsletter of the Association for Documentary Editing*, vol. 4, 1981, pp. 1–4.
- Sinclair, Stéfán, and Geoffrey Rockwell. *Voyant Tools*. 2018, <https://voyant-tools.org/>. Accessed Feb 2018.
- Smith, Abby, 'Preservation in the Future Tense', *CLIR Issues*, no. 3, 1998, <https://www.clir.org/1998/05/clir-issues-number-3/>
- Smith, Daniel W. 'The Doctrine of Univocity: Deleuze's Ontology of Immanence'. *Filozofski Vestnik*, vol. 22, no. 1, 2001, pp. 163–179.
- Smith, Daniel, and John Protevi. 'Gilles Deleuze'. *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, 2015, <http://plato.stanford.edu/archives/win2015/entries/deleuze/>. Accessed Feb 2018.
- Sperberg-McQueen, C. M. 'How to Teach Your Edition How to Swim'. *Literary and Linguistic Computing*, vol. 24, no. 1, Apr. 2009, pp. 27–39.
- Stagoll, Cliff. 'Arborescent Schema'. *Deleuze Dictionary Revised Edition*, edited by Adrian Parr, Edinburgh University Press, 2010, pp. 14–15.

- . 'Events'. *Deleuze Dictionary Revised Edition*, edited by Adrian Parr, Edinburgh University Press, 2010, pp. 89–91.
- Sukovic, Suzana. 'Information Discovery in Ambiguous Zones of Research'. *Library Trends*, vol. 57, no. 1, 2008, pp. 82–87.
- Sutherland, Juliet. 'A Mass Digitization Primer'. *Library Trends*, vol. 57, no. 1, 2008, pp. 17–23.
- Sutherland, Kathryn. 'Anglo-American Editorial Theory'. *The Cambridge Companion to Textual Scholarship*, Cambridge University Press, 2013.
- Sutherland, Kathryn, and Elena Pierazzo. 'The Author's Hand: From Page to Screen'. *Collaborative Research in the Digital Humanities*, edited by Marilyn Deegan and Willard McCarty, Ashgate, 2012, pp. 191–212.
- Tanner, Simon, et al. 'Measuring Mass Text Digitization Quality and Usefulness: Lessons Learned from Assessing the OCR Accuracy of the British Library's 19th Century Online Newspaper Archive'. *D-Lib Magazine*, vol. 15, no. 7/8, July 2009.
- Tanselle, G. Thomas. *Rationale of Textual Criticism*. University of Pennsylvania Press, 1989.
- . 'The Varieties of Scholarly Editing'. *Scholarly Editing: A Guide to Research* edited by D.C Greetham, Modern Language Association of America, 1995.
- Thaller, Manfred. 'Text as Data Type'. *ALLC-ACH '96 Abstracts University of Bergen June 25-29 1996*, University of Bergen, pp. 252–54.
- Townshend, Charles. *Easter 1916: The Irish Rebellion*. Penguin UK, 2006.
- Vanhoutte, Edward. 'Defining Electronic Editions: A Historical and Functional Perspective'. *Text and Genre in Reconstruction: Effects of Digitalization on Ideas, Behaviours, Products and Institutions*, edited by Willard McCarty, Open Book Publishers, 2013, pp. 119–44.

- Vanhoutte, E., and R. V. den Branden. 'Describing, Transcribing, Encoding, and Editing Modern Correspondence Material: A Textbase Approach'. *Literary and Linguistic Computing*, vol. 24, no. 1, Apr. 2009, pp. 77–98.
- Van Hulle, Dirk. *The Making of Samuel Beckett's Stirrings Still / Soubresauts and Comment Dire / What Is the Word*. University Press Antwerp, 2011.
- Van Hulle, Dirk, and Peter Shillingsburg, 'Orientations to Text, Revisited'. *Studies in Bibliography*, vol. 59, pp. 27–44, 2015, <https://muse.jhu.edu/article/603626/summary>. Accessed May 2019.
- Van Peursen, W. 'Text Comparison and Digital Creativity: An Introduction'. *Text Comparison and Digital Creativity: The Production of Presence and Meaning in Digital Text Scholarship*, edited by Ernst Thoutenhoofd et al., BRILL, 2010.
- Weingart, Scott. *Networks Demystified*, 2015,
<http://www.scottbot.net/HIAL/index.html@tag=networks-demystified.html>. Accessed Feb 2018.
- Williams, James. *Gilles Deleuze's Logic of Sense: A Critical Introduction and Guide: A Critical Introduction and Guide*. Edinburgh University Press, 2008.
- Wimsatt, W. K., and M. C. Beardsley. "The Intentional Fallacy." *The Sewanee Review*, vol. 54, no. 3, 1946, pp. 468–488.
- Winterer, Caroline. 'Visualizing Benjamin Franklin's Correspondence Network'. *Mapping the Republic of Letters*, 2012,
<http://republicofletters.stanford.edu/casestudies/franklin.html>. Accessed Feb 2018.
- Whitmore, Michael. 'Text: A Massively Addressable Object'. *Wine Dark Sea*, 2010,
<http://winedarksea.org/?p=926>. Accessed Feb 2018.
- Wittgenstein, Ludwig. *Philosophical Investigations*. 1953. John Wiley & Sons, 2009.

Wolpert, David H., and William G. Macready. 'No Free Lunch Theorems for Optimization'. *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, 1997, pp. 67–82.

Zeller, Hans. 'Record and Interpretation: Analysis and Documentation as Goal and Method of Editing'. *Contemporary German Editorial Theory*, edited by Hans Walter Gabler et al., University of Michigan Press, 1995, pp. 17–58.

---. 'Structure and Genesis in Editing: On German and Anglo-American Textual Criticism'. *Contemporary German Editorial Theory*, edited by Hans Walter Gabler et al., University of Michigan Press, 1995, pp. 95–124.