

Collaborative Cell Assemblies: Building Blocks of Cortical Computation ^{*}

Ronan G. Reilly

Department of Computer Science
National University of Ireland, Dublin
Belfield, Dublin 4
Ireland
Ronan.Reilly@ucd.ie

Abstract. The goal of this chapter is to propose a candidate for consideration as a building block of cortical computation. The candidate in question is the collaborative cell assembly (CCA), and I will make the case in the rest of this chapter that the concept provides a unifying framework for accounting for the mechanisms of perceptual "binding" in the sensory areas of the brain, as well as providing a mechanism for the incremental construction of cognitive and language systems from simpler components.

1 Computational Building Blocks

If one looks at how the brain processes sensory information or controls motor activity, one is struck by a surprising design feature. Rather than there being a single brain region involved in, say, object perception, a number of regions collaborate, contributing a specialization to the achievement of the overall perceptual goal. The cortex appears to function as a collection of collaborating specialists, none of which can solve the problem alone, but can only do so when each works together. Another feature of this pattern of specialization and modularity is a phenomenon of complementarity apparent in the organization of, for example, sensory processing streams. Grossberg (in press) has argued that pairs of complementary processing streams should be considered the basic cortical functional unit. A frequently cited example of complementarity in vision is the division of processing between the so-called "what" and "where" pathways. Ungerleider and Mishkin (1982) have proposed that the two major output fiber tracts from the occipital lobe, the inferior and superior longitudinal fasciculi, are responsible for carrying information about the object identity and object location, respectively. The inferior fasciculus travels into the temporal lobe, the superior fasciculus into the posterior parietal region. Notwithstanding their separate trajectories, these pathways communicate extensively. Consequently it is a good example of what Grossberg (in press) has identified as a complementary

^{*} This research was funded by an Enterprise Ireland Basic Research grant

pair of collaborative processing streams, with information from both streams required to provide a coherent percept.

At first glance, it may seem an inefficient way of constructing a perceiving-acting system by fragmenting its operations along lines that often seem counter-intuitive, and more importantly which apparently give rise to the computational cost of integration further downstream. Of course, part of the counterintuitive quality arises from the fact that we are conscious of the unity of our perceptual experience of the world. It therefore seems anomalous that different brain regions should be responsible for what appear to be intimately integrated aspects of a percept (e.g., its color and form, in the case of vision). There is then the further apparent challenge of integrating and coordinating the disparate outputs from each of the specialist modules in order for the system to mount a coherent response to a specific perceptual event. As I will argue later, the challenge may be more apparent than real.

1.1 Cortical and software engineering

Some informative analogies can be made between cortical "design" and approaches to the construction of large-scale and complex software systems. In the last decade, with the demand for increasingly complex operating system and application software, a monolithic approach to software development has no longer proved feasible. The dominant approach to software engineering is now object-oriented (OO) design and implementation. The advantage of object-oriented design is that it encourages the modularization and encapsulation of functionality, with communication between modules being limited to messages in a mutually agreed format. Flow of control in an object-oriented program is not centralized, but distributed throughout the objects or modules comprising the system. A key attraction of an object-oriented approach is that it facilitates, more readily than earlier approaches, the addition of new functionality, without the need to re-engineer the entire system. It also facilitates the re-use of already existing functionality for new purposes. It should be emphasized that re-use does not inevitably follow from an OO design, rather re-usability requires as a foundation the modularity that an OO approach provides. The price paid for the benefits of OO is the need for an agreed interface between objects or modules. Furthermore, if re-use is to be facilitated, this interface needs to be kept as simple as possible.

How does all of this relate to cortical design? In the domain of complex system design, be it cortical engineering or software engineering, several common constraints are at work: the need to manage the complexity of the problem, coordination of the elements of the solution, the need for incremental modification. In both the software and cortical cases, I will argue, complexity has been managed by decomposition of the task into smaller components. Coordination of these components has been achieved by means of a common interface, and incremental modification by keeping the interface as simple as possible. In the biological context, the issue of decomposition is uncontroversial. However, what I mean by a "common interface" will be the focus of a substantial part of this chapter.

1.2 Evolutionary Constraints

A number of additional factors constrain the space of possible solutions to designing an organism with a large perceptuo-motor repertoire. The first is that it must be constructed by the incremental tinkering of an evolutionary process driven by natural selection. An incremental approach to system construction will inevitably tend to favor the addition of small specialists modules rather than a re-engineering of the entire system. As Simon (1969) has pointed out, complex wholes will tend to be developed incrementally over evolutionary time. This is because, within an evolutionary context, the progression from simple to complex must entail viable intermediate stages.

Another, and not unrelated factor is that the hardware upon which any additional functionality is built also acts as a source of constraint. Kaas (1987) has shown that as the neocortex enlarges, there is a natural tendency for regions to form, with a high-degree of local inter-connectivity within a region, and limited inter-regional connectivity via reciprocal cortico-cortical projections. This phenomenon of regionalization arises from structural constraints imposed by the two-dimensional nature of the cortical sheet, and the physical limits on cortical connectivity. Although there is still considerable debate about the role played by neocortical size in human cognitive capacity (Passingham, 1982), some neuroscientists now argue that human brain size is only marginally larger than that for a primate scaled-up to human proportions (Roth, 2000). The uniqueness of the human brain does not rely so much on quantities of neurons, but rather on the increased structural complexity based on the formation and acquisition of additional cortical modules, specialized for a variety of evolutionarily advantageous task.

There are, however, dangers in seeing the human brain as simply a scaled-up primate brain, which in turn is a scaled-up mammalian. The view of evolution as a ladder leading inexorably to a human pinnacle has long been abandoned by evolutionary biologists, but seems still to have an inescapable hold on neuroscientists (Preuss, 1993). Contemporary evolutionary biologists prefer to use a tree-metaphor when describing the progress of evolution. Branches of the tree represent divergences from a common branch, leading to species-specific specializations. Therefore, while acknowledging that increased cortical complexity is an important part of the story, a more interesting and more challenging question is the nature of the species-specific specializations that took advantage of neocortical enlargement and that led us down the hominid branch of the tree.

An additional aspect of evolution which should not be ignored, and to which Clark (1997) very eloquently alerts us, is that it will often blur the boundary between the individual and the environment. The solutions it arrives at are very often distributed, not only across individuals, but also seamlessly across the individual and the environment. While this may seem self-evident to evolutionary biologists, it is still seen as somewhat radical in the cognitive sciences. For example, recent research in the area referred to as "interactive" vision emphasizes the need to take account of the informational content intrinsic in the variation and stability of the environment, and the fact that vision is just one element

in the perception–action behavioral loop (Churchland, Sejnowski, & Ramachandran, 1997; O’Regan & Noë, in press; Ballard, Hayhoe, & Pook, 1998). Much of this work harks back to Gibson’s pioneering theory of ecological optics (Gibson, 1966), but has been given a fresh impetus by attempts to build artificial vision systems. What it tells us is that adaptation is pervasive and unpredictable in the configuration of environment and organism that gets exploited to gain leverage in the competition for survival.

1.3 Towards a Computational Building Block

So where does that leave us in our quest for a cortical computational building block? Well, one element of the story must be the easy ability to put together and interface different specialized modules – the common interface mentioned earlier.

When I first read about the architecture of the human visual system as an undergraduate, I was impressed by the computational complexity that I thought must be involved to bind together the different sets of feature detectors active in different regions of the visual area. Only latterly, did it occur to me that perhaps, the so-called “binding” problem, is not really a problem at all, but something that is supported by an intrinsic property of the way the cortex computes. It is something so intrinsic, that evolution has exploited it on many occasions. Take for example the evolution of color vision in primates. It turns out that primates are in a minority among mammals in having color vision. The rodent-like ancestor of mammals, which emerged during dinosaur period, was of necessity a nocturnal creature. It took advantage of the nocturnal niche left by the dinosaurs who had poorer control of their body temperatures and were consequently restricted in their nighttime activities. As a nocturnal creature, there was a greater requirement for rods than cones in their visual system, consequently monochromatic vision was favored by evolution in this context. When the dominance of the dinosaurs waned, the daytime niches became occupied by these early mammals. Nonetheless, it appeared that only in certain of these new environments was color vision an advantage; one of these being the arboreal environment of early primates.

What this story of the re-emergence of color vision in mammals tells us is that the facility for adding new modules to the sensory system, and integrating their outputs to complement the output from other modules, is one that is supported by some intrinsic features of cortical computation. In the next section I will discuss what this computational support might be.

2 Collaborative Cell Assemblies

There is general agreement that some form of population coding involving an ensemble of neurons, along the lines first proposed by Hebb (1948), is implicated in cortical representation of sensory and motor patterns.

2.1 Representation

Before proceeding, however, it is important to clarify what I mean by the vexed term "representation." According to Haugeland (1991), a system that uses internal (in our case, cortical) representations must: (a) be able to map unreliable environmental inputs onto relevant representations; (b) use these representations to "stand in" for entities in the environment; and (c) manipulate the representations within a larger, systematic, framework. The main source of disagreement among cognitive scientists centers on criterion (c), though there have been one or two well-articulated proposals arguing against the need for representations of the "standing-in" variety altogether (e.g., Brooks, 1991).

For those committed to the concept of some form of representation, it is possible to identify two extremes: at one end, there are those who argue for representations with properties of context independence and compositionality of the type associated with formal linguistic and logical representations (Chomsky, 1986; Fodor & Pylyshyn, 1988; Hadley, 1999). The newer and more radical view is that the brain is ultimately a dynamical system, and that "classical" representations of the type just described are an entirely inadequate framework within which to view its operation (Kelso, 1997; Port & vanGelder, 1995). According to these dynamicists, the differential rather than the predicate calculus is a more appropriate tool for describing brain function. One finds, however, that the two opposing camps often draw on different psychological and developmental phenomena to support their respective cases. The classicists tend to focus on language and language-related reasoning tasks, whereas the dynamicists tend to emphasize action-related activities in ecologically realistic contexts. The view taken here will follow that of Clark (1997) in arguing that there is a role for representation, not necessarily of the punctate compositional variety favored by the classicists. These representations can be function specific, and particularly come into play for what Clark has called "representation-hungry" problems. In other cases, more dynamical approaches to solving the problem can be used. To paraphrase Clark (1997), we should (a) beware of putting too much into the head – the brain, body and the environment are all candidates for providing problem-solving support, and (b) beware of rigid assumptions about the nature of internal representations – they can be, as the need arises, special purpose, action related, or even classical.

There are a number of potential pitfalls in discussing representation in the brain. Implicit in many theories of neural representation is the notion of a given representation, or type of representation, being located or contained in some specific brain region. This is partly motivated by evidence from lesions, which suggest that localized damage to areas of the cortex (e.g., Broca's) has clearly defined behavioral consequences. However, there is a danger in going from these lesion studies to the position that the relevant brain region "contains" the representations relevant to a particular behavior. A more plausible account, and one in keeping with a more detailed analysis of the lesion evidence, has been proposed by Damasio and Damasio (1994) in their theory of convergence zones. They argue that complementing the observed parcellation of function is an in-

terdependency of representation. Within their view, representations in the brain are dispositional; that is, regions of the brain implicated in a concept comprise its representation. Retrieval of information involves a process of reconstruction through the activation of these interconnected and distributed representations. Convergence zones serve to tie many of these distributed components together. As Damasio et al. put it:

The essence of this framework, then, comprises reconstruction of entities and scenes from component parts and integration of component parts by time correlations. The requisite reactivation is mediated by excitatory projections. (Damasio & Damasio, 1994; p.73)

Damasio et al. emphasize the dynamical nature of the retrieval process, and the fact that a lesion serves not just to remove one or two components from the distributed representation, but also to force a reconfiguration of the whole retrieval process. Consequently, a lesion to Broca's region can be viewed as not destroying syntactic knowledge stored at the locus of the lesion, but disrupting the coordination of widely distributed knowledge used in sentence construction.

It would appear from the foregoing that the brain employs highly distributed form of representation. Pulvermüller (1999) has persuasively argued this case for word representations. He has marshaled evidence to demonstrate that words relating to a particular sensory modality or motor activity depend for their representation on the brain regions responsible for that modality or activity. Related to this view, O'Regan and Noë (in press), argue that representations of the visual world are fundamentally sensorimotor. They reject the idea that seeing derives from the activation of an internal replica of the outside world. Instead, they propose that seeing is a way of acting, a particular way of exploring the environment.

There is no internal representation whose activity generates the experience of seeing. The outside world serves as its own, external, representation. The experience of seeing occurs when the brain has mastery of the currently occurring laws of sensorimotor contingency (O'Regan & Noë, 1999, p.1).

Their anti-representationalist argument has some echoes of Brooks' (1991) position. Nonetheless, I believe that their central message of action being an integral part of visual perception is very much in keeping with the argument being developed here.

To summarize, representation in the brain appears to involve the collaboration of many different brain regions, where the collaborating regions are involved both in the acquisition of the signal giving rise to the representation, and in actions resulting from the perception of the signal.

2.2 Collaboration

It is necessary to unpack what I mean by the term "collaborative cell assembly." As has already been discussed, there is some agreement among cognitive

neuroscientists that cortical representations are realized as distributed patterns of activity over an assembly of neurons. The concept of a cell assembly is due originally to Hebb (1948) who defined an assembly as a set of neurons, either adjacent to, or distant from, each other that have become associated through the strengthening of their connections as a result of simultaneous activation. Pulvermüller (1999), among others, have proposed an additional characteristic, namely that uncorrelated firing of connected neurons will lead to the weakening of the connections between them. Abeles (1991) has also argued that the spatio-temporal pattern of firing in an assembly is a significant aspect of its functionality.

What size can a cell assembly be? Calvin (1996) defines a minimal cell assembly as a two-dimensional hexagonal arrangement of minicolumns about 0.5mm across, where a minicolumn is about 0.03 mm in diameter comprising about 100 neurons. The cell assembly involves a spatio-temporal pattern of activation within a hexagon. The motivation for the hexagon centers on the need for a reliable representation of a word or concept in the cortex. Calvin invokes a cortical selectional mechanism, which he refers to as a Darwin machine, for cloning and selecting concept representations. Within his framework, multiple copies ensure accurate transmission over noisy cortico-cortical connections. There is an implicit commitment within Calvin's work to a classical mode of representation, with the concomitant requirement that concept representations be discrete and composable. Although, the individual concepts themselves have a spatio-temporally distributed character, they are bounded and punctate when incorporated into larger sentence-like representations.

The cell assemblies envisaged by Pulvermüller (1999) involve both local clusters of neurons, along the lines proposed by Calvin, as well as neuronal groups at much larger distances, even involving traversal of the corpus callosum. The thrust of Pulvermüller's proposal is that cortical representations of words involve neuronal groups in regions that relate to their meanings. So, for example, cell assemblies representing words dealing with vision, will preferentially involve neuronal groups in the visual area of the cortex.

The collaborative cell assemblies proposed here are envisaged to be along the lines suggested by Pulvermüller; that is, involving neuronal groups encompassing all areas of the cortex. What is new about the proposal to be outlined below is that it focuses on the interaction between these types of cell assembly, and the possible role that development may play in this interaction. Of specific interest, are interacting cell assemblies in which there is an asymmetry in development, and the implication that this might have for building a cognitive system. Furthermore, we will not confine ourselves to assemblies representing words, but will extend the concept to deal with more general aspects of cortical functioning.

An important distinction among cell assembly theorists is what type of activity constitutes an active cell assembly. The simplest approach is to say that an assembly is a set of reciprocally connected neuronal groups that are, within some time window, above an excitation baseline. This is what is referred to as a "rate" code representational scheme. Problems arise when there is more than

one object being attended to, and consequently, more than one cell assembly is active. There is a need to be able to distinguish among multiple objects. A number of researchers have proposed that a temporal code is required in this case (Gray, Knig, Engel, & Singer, 1989; Crick & Koch, 1992; Singer, 1994). Therefore a set of neurons can be considered to be participating in a given cell assembly if its neurons are firing at a rate above some baseline, and have their activities synchronized.

Hebb's original proposal concentrated on the conditions for establishing and maintaining a single active cell assembly, which could in principle extend over a number of spatially separated cortical regions (Hebb, 1948). Things get a little complicated when we consider the collaboration of multiple cell assemblies. Clearly, given the scale of the human cortex, at any given time there are many millions of CAs active, some of which will be related to the same perceptual/cognitive event. How these CAs interact and possibly combine is an issue that needs to be addressed if we are to understand their functionality in more complex situations.

I propose that the best place to start looking for the simplest forms of collaborative cell assembly is in the sensory areas of the brain. Furthermore, by taking a "collaboration" perspective, I hope to show that some current problems in understanding the underlying neural mechanisms of perception may be cast in a more informative light.

2.3 Binding as Collaboration

Following stimulation of the relevant sensory cortex, activation spreads collaterally and "forward" to secondary and association areas of the cortex. In addition reciprocal activation from these forward cortical areas feeds backwards to the primary sensory areas. In the case of the primate visual system (Van Essen, et al., 1990), there are converging and diverging feedforward and feedback routes through a variety of specialized modules. This picture is also probably true of other sensory modalities.

Singer (1994) describes the hypothetical interplay of activity between two visual areas in this complex chain of processing, namely areas V1 and V5 of the visual cortex. If we assume that a subject is viewing a moving geometric shape, various features of the shape will cause neurons that are optimally responsive to these features to start responding. Tangential connections in V1 will cause this activity to spread to neighboring neurons responsive to the same features. This spreading will be enhanced by continuous or contiguously positioned features in the input. Meanwhile, in V5, where neurons are optimally responsive to direction of motion, groups of neurons that are responding to the motion of coherent contours will tend to synchronize their activity. By means of backprojections, the synchronized activity in V5 will support and selectively reinforce the feature-based activity in V1. As Singer (1994) puts it:

Responses to contour elements that are far apart and have different orientations have a low probability of being synchronized by local interaction within V1. However, if these contour elements move coherently,

their coherence would be detected by neurons in V5. Responses to these contours would synchronize in V5, and through the backprojections increase synchronization probability for the respective set of neurons in V1. (Singer, 1994; p235)

The scenario that Singer describes above is paradigmatic of interacting cell assemblies. I will consider it to be the most fundamental form of interacting CA, since it is evolutionarily ancient, having been around as long as there has been sensory systems responsive to more than one stimulus dimension.

Let us look a little more closely at the possible dynamical properties of this basic form of CA interaction. The first question is how the separate assemblies start off synchronized in the first place. The simplest explanation is that it arises from the fact that the object that is the source of stimulation impinges on the different parts of the sensory system at the same time. So, for example, V1 and V5 receive input from the retina at the same time, and the cycle of synchronous activity is simultaneously triggered in both regions. Once a disparate set of sensory regions have become active, then through a network of reciprocal projections these regions influence each other. This interactive influence can serve to sharpen a noisy signal or augment an incomplete one. So, for example, the combination of information relating to motion, contour, color, can help uniquely to specify an object in the visual field. Additional top-down influences, such as activation of prior memories of the object will further reinforce and stabilize its associated cell assembly.

When the process of perception is viewed as a collaborative one, both within low-level areas of the visual cortex, and between lower and higher levels, it becomes clear that to characterize the task of integration of disparate features of the input as a binding "problem" (e.g., Crick & Koch, 1990) is perhaps to view it from the wrong angle. What evolution has done in constructing our sensory systems is to provide additional sources of constraint on the possible identity of an object in the visual world by tapping into the information immanent in the sensory signal. Relying on motion information, or contour information alone, will get you only so far (if you are an insect, however, this may be far enough). Each additional source of information combines multiplicatively to constrain the identity of an object. The more sources of constraint in the sensory signal that you can exploit to guide your actions, the greater your survival advantage. One can think of evolution in this context as striving to maximize information transmission, in the information theoretic sense, by enlarging an organism's available bandwidth.

2.4 From Symmetric to Asymmetric Collaboration

We can readily extend the mechanism described by Singer above from the domain of perceptual information integration to more system-wide collaboration. For example, Iverson and Thelen (1999) argue for the interdependence of cognition and motor control systems in the brain. They view brain computation as a

dynamical as opposed to logical process. In their view, language, cognition, perception, and action are cut from the same computational cloth. They summarize their position as follows:

In sum, our argument for embodiment rests on the necessity for compatible dynamics so that perception, action, and cognition can be mutually and flexibly coupled. Such dynamic mutuality means that activity in any component of the system can potentially entrain activity in any other... (Iverson & Thelen, 1999, p. 37)

In my view, what they propose is functionally equivalent to the collaborative cell assembly described by Singer, but this time potentially involving cortical regions at a remove from the sensory-motor "surface", and traversing sensory modality boundaries.

In this broader view of collaboration, it is useful to make a distinction between collaborations involving cell assemblies that are equally well developed, and those in which one partner in the collaboration is more developed than the other. I will refer to the former as symmetric collaboration and the latter as asymmetric. In the case of asymmetric collaboration, there is the possibility for the less well-developed cell assembly to exploit the functionality of the more developed one. Elsewhere (Reilly, 1995), I have referred to this as cortical software re-use. The term is inspired by the software engineering concept of re-use (Bigerstaff & Perlis, 1989), and motivates the proposal that the developing partner in a collaboration is able to re-use the repertoire of cell assemblies already established in the more developed cortical region. Within this view, the sensory-motor systems, can be thought of as providing the basic computational building blocks for higher-level language and cognitive functions, which are then adapted and modified to fulfill their new role. The cortical connectivity provided by reciprocal cortico-cortical projections between sensory, motor, and association areas, and more frontal cortical areas is a key mechanism for providing access to a repertoire of re-usable functions. The style of computation that this connectivity supports is best viewed as a process of dynamical entrainment along the lines proposed by Iverson and Thelen (1999) involving the synchronization of firing patterns in reciprocally connected cortical areas. The selection of re-usable modules from one domain for re-use in another is based on a structural isomorphism, possibly supported by resonance, between the dynamical firing patterns of the source and target domains. This might mean, for example, that a motor plan for complex object-manipulation task gets re-used in the service of syntax in language production (e.g., Greenfield, 1991; Reilly, in press). The selection of relevant re-usable functions is domain or content independent, and relies on relatedness at an abstract structural level. In the case of the previous example, it might rely on the inherent hierarchical and recursive nature of both motor plans and utterance plans. The key point here is that there may be, but not necessarily, a "semantic" connection between the re-used component and its new application.

There is a wide variety of circumstantial evidence to support the re-use hypothesis. For example, Greenfield (1991) observed parallels in the developmental

complexity of speech and object manipulation. In studying the object manipulation of children aged 11–36 months, she noted that the increase in complexity of their object combination abilities mirrored the phonological and syllabic complexity of their speech production. There are two possible explanations for this phenomenon: (1) It represents analogous and parallel development mediated by separate neurological bases; or (2) the two processes are founded on a common neurological substrate. Greenfield (1991) used evidence from neurology, neuropsychology, and animal studies to support her view that the two processes are indeed built upon an initially common neurological foundation, which then divides into separate specialized areas as development progresses. If Greenfield is correct in her analysis, this is good evidence of developmental re-use by the language domain of functionality initially established in the motor domain.

Another source of circumstantial evidence is Gerstmann’s Syndrome¹ (Res-tum & Sobota, 1983). This syndrome comprises the following cluster of co-occurring symptoms arising from left parietal lesions: (a) finger agnosia, (b) acalculia, (c) left–right disorientation, and (d) agraphia. What is particularly interesting from the re-use point of view is the association between finger agnosia and acalculia. The re-use interpretation of this association is that numerical abilities exploit the functionality already established for dealing with spatial arrangements of fingers, particularly the ability to individuate or recognize as separate and distinct each finger. One of the features of finger agnosia is not being able to tell without looking whether one or more fingers is being touched. Patients with this condition have sense of their fingers being an undifferentiated mass.

While the evidence just outlined is suggestive, more compelling evidence will come from direct testing of hypotheses derived from the CCA proposal, in conjunction with computer simulation studies. Some preliminary hypotheses are discussed in the next section.

3 Implications of the CCA Perspective

I will divide the implications into two broad categories: (a) those relating to CCA as a cognitive science theory, and (b) those relating to the interaction between artificial neural network research and neuroscience. From a cognitive science perspective there are a number of testable hypotheses that can be derived from the CCA proposal, particularly relating to its re-use aspect. A number of these center on the issue of task interference. If language is constructed on a sensory–motor foundation, and continues to draw upon that underlying functionality, one should expect to find competing non–linguistic demands on this resource having an impact on language performance. So, for example, we should find subjects’ language performance impaired when asked to produce syntactically complex utterances while at the same time performing a motor task with, say, some form of recursive structure. Contrariwise, we would not expect to find motor

¹ Thanks to Dana Mackey for drawing my attention to this

performance impaired by the complexity of a language task. In a similar vein, subjects' performance on a mental arithmetic task should show some degradation if they are asked simultaneously to carry out a motor task involving their fingers. Obviously, considerable effort would have to be invested in designing appropriate control conditions for these type of experiments.

The CCA perspective also has a number of implications for the construction of computational artifacts to tackle difficult AI problems. One such implication is that an approach involving the incremental construction of solutions starting with simpler ones and using these as a foundation for more complex ones might be a fruitful one to explore. To some degree the work on neural network ensembles can be considered to be taking this approach, though motivated by reasons of reliability and redundancy rather than the learning of complex tasks (see Sharkey, 1999, for a comprehensive survey of this field). A marrying of CCA and ensemble approaches would, I believe, provide a powerful tool for the construction of large-scale neural network applications.

4 Conclusions

In this chapter I have proposed a candidate building block for cortical computation, which I refer to as a collaborative cell assembly (CCA). I have argued that CCAs arose originally from an evolutionary requirement for incremental increases in the functionality and bandwidth of sensory systems. However, in conjunction with the extended developmental period of the human cortex, they have provided a framework for re-using perceptuomotor "software" to construct the cognitive and linguistic systems. The CCA proposal can be used to generate a number of testable hypotheses relating primarily to resource competition in the execution of cognitive tasks that utilize underlying sensory-motor functions. Furthermore, implications are drawn regarding how best to construct artificial neural networks to tackle large scale problems.

References

- Abeles, M. (1991) *Corticotonics*. Cambridge, UK: Cambridge University Press.
- Ballard, D.H., Hayhoe, M.M., Pook, P.K. & Rao, R.P.N. (1997) Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, **20**, 723-767.
- Biggerstaff, T.J. & Perlis, A.J. (1989). *Software Reusability: Concepts & Models*. New York: ACM (Addison-Wesley).
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, **47**, 139-59
- Calvin, W.H. (1996). *The cerebral code*. Cambridge, MA: MIT Press.
- Chomsky, N. (1986). *Knowledge of language*. New York: Praeger.
- Churchland, P., Sejnowski, T., & Ramachandran, V.S. (1995). A critique of pure vision. In C. Koch & J.L. Davis (eds.), *Large-scale neuronal theories of the brain*. Cambridge, MA: MIT Press.
- Clark, A. (1997). *Being there: Putting brain, body, and the world together again*. Cambridge, MA: MIT Press.

- Crick, F., & Koch, C. (1990). Toward a neurobiological theory of consciousness. *Sem. Neurosci.*, **2**, 263–275.
- Damasio, A.R., & Damasio, H. (1994). Cortical systems for retrieval of contents knowledge: The convergence zones framework. In C. Koch & J.L. Davis (eds.), *Large-scale neuronal theories of the brain*. Cambridge, MA: MIT Press, pp 201–238.
- Fodor, J. A., & Pylyshyn, Z.W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, **28**, 371.
- Gibson, J.J. (1966). *The senses considered as perceptual systems*. Boston, MA: Houghton Mifflin.
- Gray, C. M., Knig, A., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, **338**, 334–337.
- Greenfield, P. (1991). Language, tool and brain: The ontogeny and phylogeny of hierarchically organized sequential behavior. *Behavioral and Brain Sciences*, **14**, 531–595.
- Grossberg, S. (in press). The complementary brain: A unifying view of brain specialization and modularity. *Trends in Cognitive Sciences*.
- Hadley, R.F. (1994). Systematicity in connectionist language learning. *Mind and Language*, **9**, 247–271.
- Haugeland, J. (1991). Representational genera. In W. Ramsey et al. (Eds.), *Philosophy and connectionist theory*. Hillsdale, NJ: Erlbaum.
- Hebb, D.O. (1948) *The organization of behavior*. New York: Basic Books.
- Iverson, J.M., & Thelen, E. (1999). Hand, mouth, and brain: The dynamic emergence of speech and gesture. *Journal of Consciousness Studies*, **6**, 19–40.
- Kaas, J.H. (1987). The organization of neocortex in mammals: Implications for a theory of brain function. *Annual Review of Psychology*, **38**, 124–151.
- Kelso, J.A. S. (1997). *Dynamic patterns*. Cambridge, MA: MIT Press.
- O'Regan, J.K., Noë, A. (in press). A sensorimotor account of vision and vision consciousness. *Behavioral and Brain Sciences*.
- Passingham, R.E. (1982). *The human primate*. San Francisco: W.H. Freeman.
- Port, R. & vanGelder, T. (Eds.) (1995). *Mind as motion*. Cambridge, MA: MIT Press.
- Preuss, T. M. (1993). The role of the neurosciences in primate evolutionary biology. Historical commentary and prospectus. In R.D.E. McPhee (Ed.), *Primates and their relatives in phylogenetic perspective*. New York: Plenum Press, pp. 333–362.
- Pulvermüller, F. (1999). Words in the brains language. *Behavioral and Brain Sciences*, **22**, 253–336.
- Reilly, R.G. (in press). On the relationship between object assembly and language production: A connectionist simulation of Greenfield's hypothesis. *Behavioral and Brain Sciences*.
- Reilly, R.G. (1995). Sandy ideas and coloured days: The computational implications of embodiment. *Artificial Intelligence Review*, **9**, 1–18.
- Restum, W. & Sobota, W.L. (1983). Gerstmann's syndrome: A new perspective to some controversial issues. *Archives of Physical Medicine and Rehabilitation*, **64**, 499.
- Roth, G. (2000). Is the human brain unique. Paper presented to Mirror Neurons and the Evolution of Brain and Language, Hanse Institute for Advanced Study, Delmenhorst, Germany.
- Sharkey, A. (1999). *Combining artificial neural networks: Ensemble and modular multi-net systems*. London: Springer-Verlag.
- Simon, H. (1969). The architecture of complexity. In H. Simon (Ed.), *The sciences of the artificial*. Cambridge, UK: Cambridge University Press.

- Singer, W. (1994). Synchronization of cortical activity and its putative role in information processing and learning. In C. Koch & J.L. Davis (eds.), *Large-scale neuronal theories of the brain*. Cambridge, MA: MIT Press, pp 201–238.
- Thatcher, R.W. (1992). Cyclic cortical reorganization during early childhood. *Brain and Cognition*, **20**, 2450.
- Ungerleider, L.G., & Mishkin, M. (1982). Two cortical visual system. In J. Ingle, M.A. Goodale, & Mansfield, R.J.W. (Eds.), *Analysis of visual behavior*. Cambridge, MA: MIT Press, pp. 549–586.
- Van Essen, D.C., & Anderson, C.H. (1990). Information processing strategies and pathways in the primate retina and visual cortex. In S.F. Zornetzer, J.L. Davis, & C. Lau (Eds.), *An introduction to neural and electronic networks*. New York: Academic Press, pp. 43–72.