Taylor & Francis
Taylor & Francis Group

Mapping the changing residential geography of White British secondary school children in England using visually balanced cartograms and hexograms

Richard Harris, Martin Charlton & Chris Brunsdon

Social Science

Mapping the changing residential geography of White British secondary school children in England using visually balanced cartograms and hexograms

Richard Harris [a], Martin Charlton[b] and Chris Brunsdon[b]

[a]School of Geographical Sciences, University of Bristol, Bristol, UK; [b]National Centre for Geocomputation, Maynooth University, Maynooth, Ireland

ABSTRACT
In the context of debates about segregation within the UK, this paper maps the residential geography of two groups of White British school children, one of which was in secondary school in 2011 and the other in 2017. To present that geography, hexograms are introduced as a complement to visually balanced cartograms, both of which seek to address the problems of invisibility and distortion encountered with more conventional choropleth and cartogram maps. The nature of these problems is introduced, our solutions discussed, and the methods applied to the case study, which allow changes in the geography to be seen.

1. Introduction

The objective of this paper is to chart any changes to the geography of where White British children in two middle stages of secondary school education were living in England in 2011 and 2017. To achieve this, we present what we call hexograms as a way of mapping areal data. These are a visual method that allows areas to be represented as equally sized hexagons on an underlying cartogram map. As such, they are a cross between visually balanced cartograms, hexagonal binning and tile maps (see below), designed to tackle two representational issues common to area-based maps of population distributions: the problem of invisibility, found in conventional choropleth maps, and the problem of distortion, created by cartograms (Harris, Charlton, Brunsdon, & Manley, 2017a).

The study considers the limitations of conventional choropleth and cartogram approaches for mapping spatial distributions across areas of varied size. It does so in the context of a response to the UK Government-sponsored *Casey Review: a review into opportunity and integration* (Casey, 2016), which reignited debate on whether Britain is becoming more socially and ethnically divided. Newspaper headlines such as that published in the national *Metro* newspaper imply that it is: 'Diverse yet divided: UK is growing apart' (December 5, 2016). However, such impressions run largely counter to a range of empirical studies that show residential ethnic segregation to have fallen within the UK between 2001 and 2011 (Catney, 2015; Harris & Owen, 2017; Johnston, Poulsen, & Forrest, 2013). There is an exception. Considered as a whole, the White British in London, metropolitan areas and other large cities are the only group for which segregation from other ethnic groups increased from 2001 to 2011 in England and Wales (Catney, 2013). Cantle and Kaufmann (2016) have said that there is 'a growing isolation of the White majority from minorities in urban zones'.

A limitation of these existing studies is that they are based on census data collected only once a decade, the most recent measuring the Britain of seven years past. Little is known about recent trends (but see Lan, Kandt, & Longley, 2018 for an innovative study using consumer data). What we are interested in exploring is whether the apparent spatial contraction of the White British from urban locations has continued over recent years.

To that end, we follow Harris (2017) and turn to an alternative source of data, which is the National Pupil Data (NPD) for England for the years 2011 (the year of the last UK Census) and 2017 (the most recent data). The NPD has been described by the Department for Education as one of the richest education datasets in the world, holding a wide range of information about students. However, it generally covers only pupils in state schools in England (which is most: approximately 93% of all pupils are in state schools). The data used here are all pupils in two of the middle years of compulsory secondary education, who are aged 13–14 or 14–15 and attending a state school. We look at where the White British of those pupils are living and map any changes.

2. Choropleth maps, cartograms and the problems of representation

Our cartographic starting-point is Figure 1, which maps the percentage of the Census population that was counted as White British per English local authority (LA) in 2011 and, alongside, the corresponding percentage of the NPD data for the same year.

Measuring ethnicity is not unproblematic. Which of the pupils identifies as White British in each year is a matter of constrained choice: for the NPD, pupils or their parents select their ethnicity from a set of categories that are similar to those used in the UK Census. They may also refuse to provide details (1.0% of the pupils in the 2011 data has an unknown ethnicity, including refusals; 1.3% in 2017). Because it is a partial choice and a reflection of a person's own, sometimes changing self-identity, the ethnic classification fluctuates for some pupils over time (see https://nationalpupildatabase.wikispaces.com/Ethnicity). The growing instability of ethnic identity has been identified in the Census (Simpson, Jivraj, & Warren, 2016).

Nevertheless, there are interesting and potentially important differences between the two maps. Unfortunately, few of those are obvious because many occur in the major urban conurbations where the spatial patterns are hard to see. This is especially true of the places in and around London which form the conurbation towards the bottom right of the map.

The problem, common to conventional choropleth maps of area-based data, is that the areas to be mapped – the English local LAs – vary from one another in both physical and population size. Furthermore, the places with most people living in them are amongst the smallest in area: the Spearman rank correlation between area and the 2011 Census count of residential population is –0.175. This means that many of the areas of most interest are amongst the ones rendered smallest on the map, some to the point of 'invisibility'.

Potential solutions to the problem include using a map insert to redraw London at a larger scale than the rest of the country. However, whilst this might work for the capital, it would be of no use to any other 'invisible' places in other parts of the map unless they too had their own map insert. An alternative, increasingly popular method, is the cartogram. Any map is a distortion of the real world, selectively choosing methods of projection and visualisation for the purpose of its design. Nevertheless, cartograms are distinct in breaking the more usual link between the size of an area displayed on the map and its true, physical size. Instead, areas are re-scaled in accordance to some attribute such as population size.

Cartograms are not new: their history is described by Dorling (1996); a more recent review is provided by Nusrat and Kobourov (2016). Over the 20-year gap between those publications, what was once an esoteric method of cartographic visualisation has become much more common as the range of software available to produce cartograms has expanded. There are several types of cartogram (see Table 1 of Nusrat & Kobourov); we specifically focus on contiguous cartograms that modify the boundaries, areas and shapes of areas but maintain adjacencies. The cartograms shown in Figure 2 employ the Dougenik, Chrisman, and Niemeyer (1985) algorithm but the arguments to be made apply equally to Gastner and Newman's (2004) diffusion method and others that have in common the broad principle of re-sizing areas by stretching and deforming them.

Elsewhere we have argued that ease of production has led to an uncritical adoption of cartograms in a number of academic papers, reports and media presentations with too little consideration to their effectiveness in communicating the data available (Harris, Charlton, Brunsdon, & Manley, 2017b). In the present case, the cartogram has some success. Noting that blue shading on the maps represents lower percentages ('cold spots') of White British populations, the cartograms enable it to be seen that those are concentrated within urban areas, including but not limited to London, and that they are usually colder for the NPD than for the Census population; that is, the percentages of the pupils that are White British in urban areas are often less than the percentages of the Census population.

It is possible that these differences are due to the NPD data being biased by the absence of pupils who attend fee-charging schools. Against this, we are reassured by a close correspondence between the percentage of the White British population that is White British per LA and the percentage of the NPD pupils: the Pearson correlation is 0.989. Fitting a regression line with a forced zero y-intercept to predict the percentage in the NPD data from the percentage in the Census produces a coefficient (slope) of 0.978. This is very close to a one-to-one relationship and has a model fit, measured by the R^2 value, of 0.997. Nevertheless, such a model tends to over-predict the percentage of the NPD pupils that are White British in local authorities where the Census percentage is less than 60, which is 31 of the 324 local authorities, amongst which 26 are in London. The over-prediction suggests that in areas where the White British are less prevalent in the Census population then, relatively speaking, they are even less so amongst the group of secondary pupils. This may imply that White British families and their children are moving out from urban places perhaps as a process of school choice (Butler & Hamnett, 2011). However, other explanations include broader demographic changes (the decreased number of White British pupils enrolled in secondary schools nationally), and also the possibility that in places (notably London) that are much more ethnically diverse than most of England, younger age groups are more
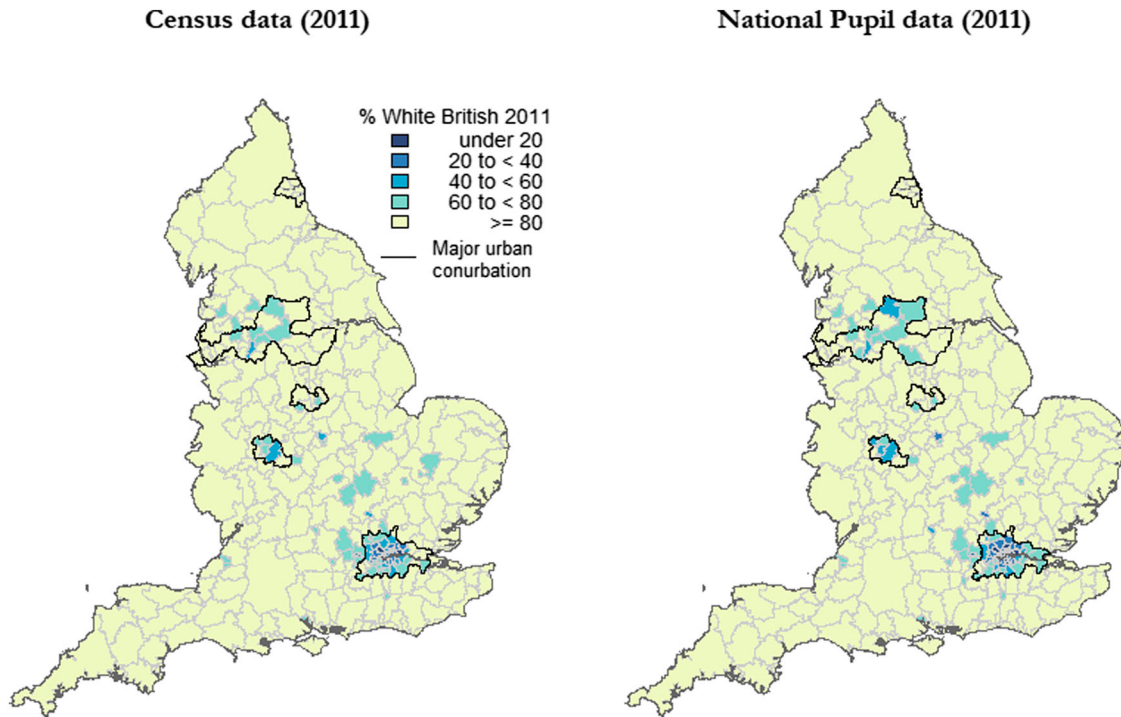
Census data (2011)   National Pupil data (2011)

Figure 1. Showing the percentage of (left) the Census population, and (right) the cohort of secondary school pupils that were White British in each English local authority in 2011.

comfortable expressing their ethnic identity as a category other than White British.

Unfortunately, the success of the cartogram comes at a price: the cost of re-scaling is to replace misrepresentation through invisibility with misrepresentation through distortion because the geography of the locations is warped. This exchange is conducted on unfavourable terms because the cartogram, which is dominated by places assigned a large scaling term, also results in areas that are too small to see. As such, it has failed to resolve the problem of invisibility (only substituting one set of invisible places with another) and compounded it with distortion.

The fundamental limitation that the cartogram encounters is that a scaling parameter with a strongly skewed distribution (the areas of the local authorities) cannot be exchanged for another that also is strongly skewed (the population counts) if the objective is to produce a map over which the entirety of the geographical patterning is legible. For the local authorities, the area of the largest is 10.3 times greater than the inter-quartile range. It would be about 2.86 if the variable was randomly sampled from a Normal distribution. For their population sizes, the ratio is very similar, at 10 so there is little visual gain. If we follow others and rescale not by population counts but by a more specific attribute, in this case, the numbers of people that are not White British, the ratio increases to 23.7. This draws attention to the places where fewer White British are living but the greater the amount of skew the more unrecognisable the geography of the country becomes (Figure 3).

None of this discussion means that a cartogram is unsuitable in every usage. In some cases, distorting the map to draw attention to a large presence (or absence) of a feature is the desired aim (see Hennig, 2014 and the examples at https://worldmapper.org/). In others, they can be used to show the clustering of, for example, disease cases amongst a variably sized population at risk. In all cases, cartogram producing software are agnostic to the choice of scaling term; it is up to the user to decide. With that in mind, Harris et al. (2017a) propose what is described as a visually balanced cartogram, which is a cartogram where a balance between visibility and distortion is sought by enlarging only those areas on the map that fall below a minimum, interpretable size and, therefore, need to be expanded in order to be seen. Although a cartogram can be defined as a map that is scaled in proportion to some thematic or population count of interest (thereby combining statistical and geographic information), our view is that it is any map that rescales the size of areas in accordance to an attribute other than their physical size. For us, the most important attribute is their visual interpretability. A similar logic is employed by Soetens, Hahne, and Walllings (2017), expanding areas sufficiently to represent each of a point outbreak of an infectious disease.

Figure 4 is based on this idea with the smallest interpretable unit set at a target of 0.02 squared-inches. If an error is defined (loosely) as the percentage amount by which the original map and the cartogram do not overlap, then the gain from the visually balanced cartogram is clear: for the population cartogram the error is 15%,
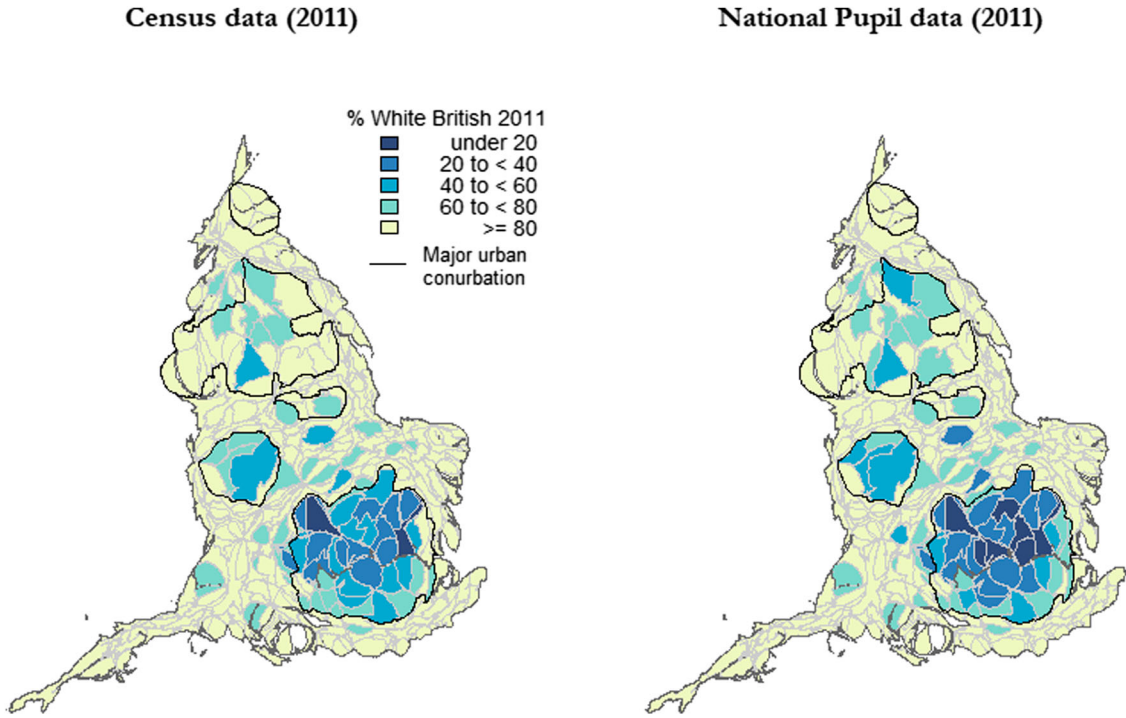
Census data (2011)

National Pupil data (2011)

Figure 2. The maps of Figure 1 redrawn with a population-based cartogram (the scaling parameter is the 2011 Census population count per local authority).

for the attribute based cartogram it is 24% and for the balanced cartogram it is only 5.9. More precisely, the error may be defined as:

$$\epsilon = 100 - 50A_{x \cap y}\left(\frac{1}{A_x} + \frac{1}{A_y}\right),$$

where A_x is the area of the original map, A_y is the area of the cartogram, and $A_{x \cap y}$ is the area of their geographical intersection.

Another way to consider the amount of distortion is to calculate the average displacement of the area centroids from their locations on the original map to their new locations on the cartogram. This is 28.49 km for the population cartogram, 43.97 km for the attribute cartogram and 13.78 km for the visually balanced cartogram. Both measures rely on the cartogram being placed within the same broad bounding box as the original map and oriented along the same
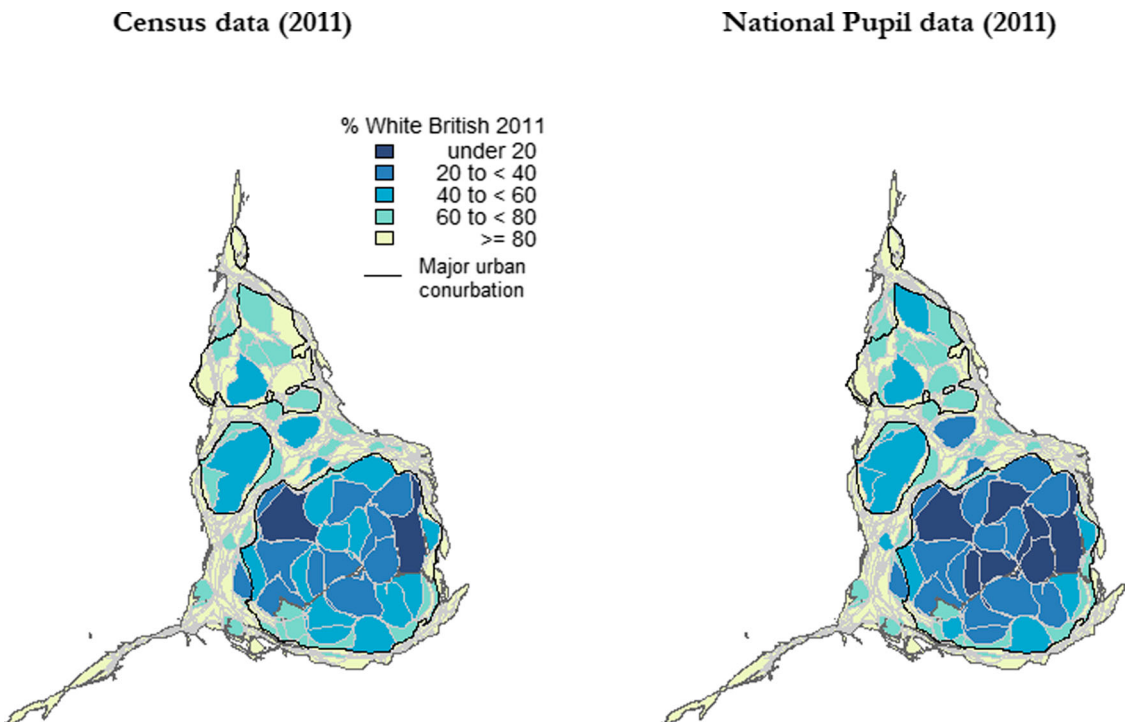
Census data (2011)

National Pupil data (2011)

Figure 3. Showing an attribute-based cartogram (the scaling parameter is the Census count of population who are not White British in each local authority).

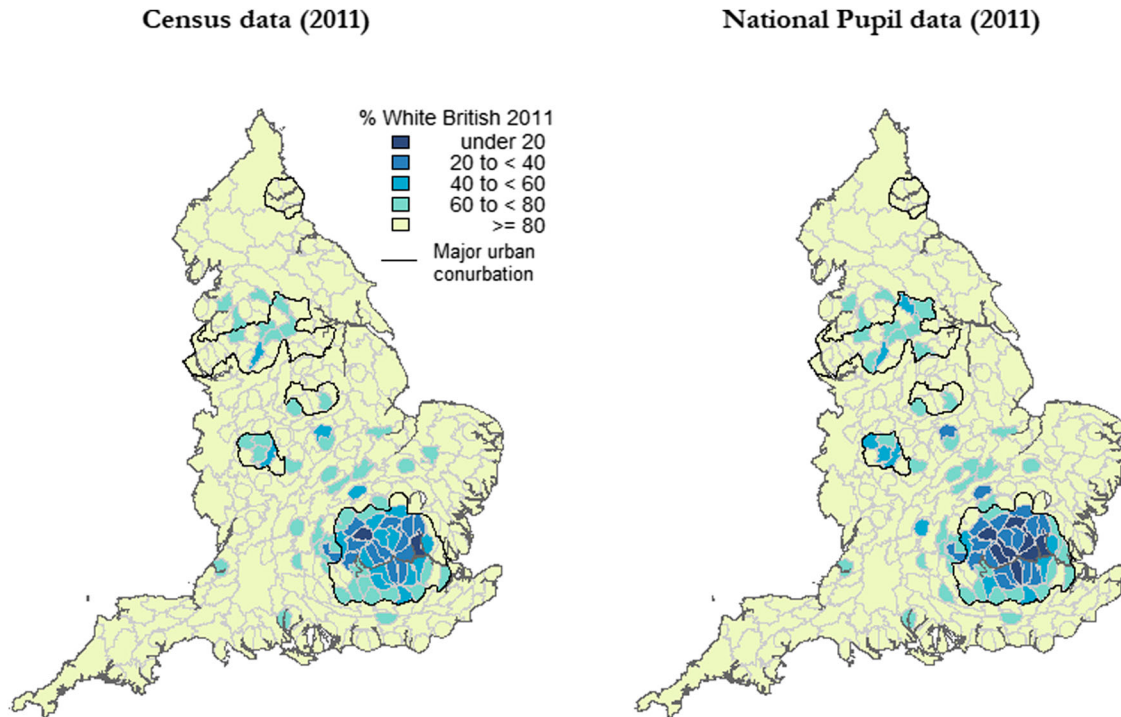Census data (2011)   National Pupil data (2011)

Figure 4. A visually balanced cartogram offering a better compromise between visibility and distortion.

vertical and horizontal axes (i.e. there is no rotation applied wholescale to the map nor a recentring of it). On these two measures, plus a third criterion of visibility, the visually balanced cartogram is the best of the maps.

3. Methods

The Main Map displays the pupil data using what we have described as hexograms. These are based on creating a cartogram where the minimum size of each area is that which allows each LA to be represented by its own, unique hexagon in a process of hexagonal binning.

Hexagonal binning has become a popular method of data visualization that is used to aggregate points of data on a map or chart (such as a scatter plot) into groups based on their location on the chart: a hexagonal grid is overlaid upon the points, the number of points in each hexagon is counted, and the hexagon is shaded to represent the number (and also, if the size of the hexagon is fixed, the density) of points within each hexagon. For geographical point data this is useful for representing spatial variation in, for example, the number of cases of a disease. Here, however, we use hexagonal binning for a different purpose where what we are seeking is a visual representation of area, not point data, and we want each area to be represented by its own, unique hexagon.

There are two ways to achieve this. The first is to raise the number of bins to the amount producing no conflicts (no shared hexagons). This entails the least geographic distortion of the map but the

resulting hexagons are too small to be visible. The second is to increase the size of the areas where the clashes are occurring but this will distort the map. We therefore encounter the now familiar trade-off between invisibility and distortion: fewer bins means greater visibility; more bins means less distortion. There is no one perfect solution but in making a judgement, it is helpful to gauge the expected number of conflicts for each number of bins as this will indicate how many places need enlarging. A graph of the two suggested that 23 bins are suitable for the present study.

The algorithm creating the hexograms works as follows. To begin, an initial cartogram is created that increases the size of the smallest areas in the map. Hexagonal binning is then applied to the centroids of the cartogram, identifying areas that conflict. To address these conflicts, an attempt is made to move the centroids to different parts of the conflicted areas to see if that will separate them. If not, the areas are enlarged some more. The process iterates through these procedures but, as a last resort, will attempt a third way to achieve a resolution, which is to increase the number of bins. As a final stage, some careful re-arrangement of the hexagons can return those that are less clustered towards the edge of the map back to their original positions, resulting in a non-tessellating tile map that, where possible, reflects the initial geography of the LA controids in the original map. The centroids of the hexagons in the tile map have a displacement of 26.34km from their original positions.

It might be argued that representing each area with equally sized hexograms suggests some equivalence

between places. That is correct: each represents one observed value, which is the percentage of the pupils in each LA who are White British. This is a little different from when, for example, tile maps are used in UK electoral studies, where each tile represents a constituency and – although these do also vary in their population size – each counts equally under the national election's 'first-past-the-post' system: their equivalence is determined by the electoral system. Since that is not the case here, should the hexograms be scaled by population size, which is more usual for cartograms? That is a possibility and an area for future research. For now, however, our interest is in the rates of White British population and how those are spatially distributed across the LAs. The spatial distribution of the population size of each LA is not our immediate concern and we certainly do not want it to dominate the maps by the geographical distortion that cartograms typically generate. What is of interest to us is the visual representation of the percentages and the geographical variation.

Much of the computational time to produce the map is spent on creating and re-creating the cartogram, which is slower than producing a tile map directly (see, for example, McNeill & Hale, 2017). However, there is advantage in having the hexogram produce the tile map as well as the cartogram. The tile map is visually appealing, with all areas clearly visible within it. It also ensures sufficient space (within each hexagon) to include additional annotation, marking on the map the places where the numbers of White British pupils have grown. However, only the cartogram retains the topological connections between local authorities. They therefore offer complementary representations of the same underlying data and are employed together in the Main Map.

4. Results

The Main Map shows how the residential geography of the White British school pupils changes from 2011 to 2017. The upper and middle rows map the percentage of the pupils that are White British, in 2011 and in 2017. As before, the LAs are shaded to highlight 'cold spots', where the percentages are relatively low. Evident is that the percentages have mostly fallen over the period: they decreased in 309 local authorities and rose in only 15. The biggest decrease is in Carlisle which may reflect it having one of the country's fastest rates of ethnic minority population growth.

Across the country, the number of White British recorded in the extracts from the NPD data decreased from 870613 in 2011 to 751067 in 2017, a reduction of 13.7%. At the same time, the number in other ethnic groups rose from 246453 to 304404, an increase of 23.5%, which includes an increase in the Mixed ethnicty group from 41671 to 51695 (24.1%), and in the White Other from 38154 to 54255 (42.2%). The latter, especially, reflects immigration from EU and other countries but a component of the reduced number of White British might not just be demographic changes and migration but also the instability of the White British category discussed earlier (i.e. an increased willingness to identify with and to choose categories other than White British for oneself).

Given these national changes, the LAs have been classified on two dimensions. First, those where the percentage *loss* of White British is greater than the national trend (higher– …) and those where the percentage loss is less than nationally (lower– …). Second, those where the percentage *gain* of other ethnic groups is greater than nationally (… –higher), and those where it is less (… –lower). That gives four types of LA: higher–higher, higher–lower, lower–lower and lower–higher. Of particular interest are places where the loss of White British is greater and so too is the gain of other ethnic groups. These (higher–higher) authorities are the ones likely to experience increased residential separation of the White British from other groups. Also of interest are places where the loss of White British is relatively low and the gain of other ethnic groups is relatively high. These (lower–higher) local authorities are likely to have more residential mixing.

The Main Map shows that it is too simple to portray the loss of the White British pupils as a purely urban phenomenon. Rates of loss are higher in some rural and smaller town locations, especially along the east coast, where the rates of growth of other ethnic groups are higher. At the same time, many of the locations with a lower rate of decrease in the White British are also locations with a higher rate of increase of other groups (albeit sometimes from a low base). This is especially true of the suburban, semi-rural and rural areas around London but also around other cities and urban conurbations too. In just a few locations, the number of White British pupils increased from 2011 to 2017. Three of those are in London. The hexograms enable the visualisation of information that would be obscured or hidden in conventional choropleth and cartogram maps.

5. Conclusions

In this paper, we have shown how balanced cartograms and hexograms can be used to tackle the (mis-)representational issues limiting more conventional choropleth and cartogram mappings, and to map the changing residential geography of two cohorts of secondary school pupils, one from 2011, the other 2017. Our approach is not a panacea but it is informed by the knowledge that more judicious selection of a scaling variable will better balance the problems of invisibility and distortion that affect area-based maps of population distributions.

The methods have been implemented as a proof of concept built on existing software libraries and are not at all optimised for speed. With the 324 local authorities, the hexogram took 5 minutes to produce on a one-year-old, regular Windows-based laptop. A tutorial on how to apply the methods using the free and open source software, R, is available at https://rpubs.com/profrichharris/hexograms.

With reference to the geography we have mapped, although the spatial retrenchment of the White British into more semi-rural and rural locations appears to have continued beyond the 2011 Census, it does not follow that the country is becoming more segregated because the evidence is that those locations are becoming more ethnically mixed (Harris, 2017) and some urban locations have had an increase in their White British populations. The notion of 'white avoidance' raised by Cantle and Kaufmann (2016) seems, at a minimum, to be an over-generalisation of more complex demographic changes and their causes. For longer term understanding of segregation patterns, those demographic factors must be considered more fully: whereas the numbers of White British have decreased in the secondary school years over the period, within primary schools they have increased from 785,717 to 852,735 over the period (an increase of 8.5%)

Software

Each of the maps and methods was implemented and drawn in R, version 3.4.3. The main software libraries used were cartogram version 0.0.2 (Jeworutzki, 2016), fMultivar version 3042.80 for the hexagonal binning (Wuertz, Setz, & Chalabi, 2017), and the spatial libraries GISTools version 0.7-4 (Brunsdon & Chen, 2014), sp version 1.2-7 (Bivand, Pebesma, & Gomez-Rubio, 2013; Pebesma & Bivand, 2005), rgdal version 1.2-16 (Bivand, Keitt, & Rowlingson, 2017) and rgeos version 0.3-26 (Bivand & Rundell, 2017) to spatially manipulate and visualise the data. The Main Map was finished in Serif Affinity Designer.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

ORCID

Richard Harris http://orcid.org/0000-0001-7943-9005

References

Bivand, R., Keitt, T., & Rowlingson, B. (2017). *rgdal: Bindings for the 'Geospatial' Data Abstraction Library*. Retrieved from https://CRAN.R-project.org/package=rgdal

Bivand, R., & Rundell, C. (2017). *rgeos: Interface to Geometry Engine - Open Source ('GEOS')*. Retrieved from https://CRAN.R-project.org/package=rgeos

Bivand, R. S., Pebesma, E., & Gomez-Rubio, V. (2013). *Applied spatial data analysis with R* (2nd ed.). New York: Springer.

Brunsdon, C., & Chen, H. (2014). *GISTools: Some further GIS capabilities for R*. Retrieved from https://CRAN.R-project.org/package=GISTools

Butler, T., & Hamnett, C. (2011). *Ethnicity, class and aspiration: Understanding London's new East end*. Bristol: Policy Press.

Cantle, T., & Kaufmann, E. (2016). *Is segregation on the increase in the UK?* Retrieved from https://www.opendemocracy.net/wfd/ted-cantle-and-eric-kaufmann/is-segregation-on-increase-in-uk

Casey, L. (2016). *The Casey review: A review into opportunity and integration*. London: Department for Communities and Local Government.

Catney, G. (2013). *Has neighbourhood ethnic segregation decreased?* Manchester: ESRC Centre of Dynamics of Ethnicity. Retrieved from http://hummedia.manchester.ac.uk/institutes/code/briefingsupdated/has-neighbourhood-ethnic-segregation-decreased.pdf

Catney, G. (2015). Exploring a decade of small area ethnic (de-)segregation in England and Wales. *Urban Studies*, *53*(8), 1691–1709. doi:10.1177/0042098015576855

Dorling, D. (1996). *Area cartograms: Their use and creation* (Concepts and Techniques in Modern Geography 59). Norwich: Environmental Publications.

Dougenik, J. A., Chrisman, N. R., & Niemeyer, D. R. (1985). An algorithm to construct continuous area carograms. *The Professional Geographer*, *37*(1), 75–81. doi:10.1111/j.0033-0124.1985.00075.x

Gastner, M. T., & Newman, M. E. J. (2004). Diffusion-based method for producing density-equalizing maps. *Proceedings of the National Academy of Sciences*, *101* (20), 7499–7504. doi:10.1073/pnas.0400280101

Harris, R. (2017). Measuring the scales of segregation: Looking at the residential separation of White British and other schoolchildren in England using a multilevel index of dissimilarity. *Transactions of the Institute of British Geographers*, *42*(2), 432–444. doi:10.1111/tran.12181

Harris, R., Charlton, M., Brunsdon, C., & Manley, D. (2017a). Balancing visibility and distortion: Remapping the results of the 2015 UK general election. *Environment and Planning A*, *49*(9), 1945–1947. doi:10.1177/0308518X17708439

Harris, R., Charlton, M., Brunsdon, C., & Manley, D. (2017b). Tackling the curse of cartograms: Addressing misrepresentation due to invisibility and to distortion. Proceedings of the 25th GIS Research UK Conference, Manchester. Retrieved from http://huckg.is/gisruk2017/GISRUK_2017_paper_29.pdf

Harris, R., & Owen, D. (2017). Implementing a multilevel index of dissimilarity in R with a case study of the changing scales of residential ethnic segregation in England and Wales. *Environment and Planning B*. doi:10.1177/2399808317748328

Hennig, B. D. (2014). Gridded cartograms as a method for visualising earthquake risk at the global scale. *Journal of Maps*, *10*(2), 186–194. doi:10.1080/17445647.2013.806229

Jeworutzki, S. (2016). *cartogram: Create Cartograms with R*. Retrieved from https://CRAN.R-project.org/package=cartogram

Johnston, R., Poulsen, M., & Forrest, J. (2013). Multiethnic residential areas in a multiethnic country? A decade of major change in England and Wales. *Environment and Planning A*, 45(4), 753–759. doi:10.1068/a4688

Lan, T., Kandt, J., & Longley, P. (2018). Ethnicity and residential segregation. In P. Longley, J. Cheshire, & A. Singleton (Eds.), *Consumer data research*. London: UCL Press. Retrieved from http://discovery.ucl.ac.uk/10046615/1/Consumer-Data-Research.pdf

McNeill, G. M., & Hale, S. A. (2017). Generating tile maps. *Computer Graphics Forum*, 36(3), 435–445. doi:10.1111/cgf.13200

Nusrat, N., & Kobourov, S. (2016). The state of the art in cartograms. *Computer Graphics Forum*, 35(3), 619–642. doi:10.1111/cgf.12932

Pebesma, E. J., & Bivand, R. S. (2005). Classes and methods for spatial data in R. *R News*, 5(2). https://cran.r-project.org/doc/Rnews/

Simpson, L., Jivraj, S., & Warren, J. (2016). The stability of ethnic identity in England and Wales 2001–2011. *Journal of the Royal Statistical Society, Series A*, 179(4), 1025–1049. doi:10.1111/rssa.12175

Soetens, L., Hahne, S., & Walllings, J. (2017). Dot map cartograms for detection of infectious disease outbreaks: An application to Q fever, the Netherlands and pertussis. *Germany. Eurosurveillance*, 22(26). doi:10.2807/1560-7917.ES.2017.22.26.30562

Wuertz, D., Setz, T., & Chalabi, Y. (2017). *fMultivar: Rmetrics - Analysing and Modeling Multivariate Financial Return Distributions*. Retrieved from https://CRAN.R-project.org/package=fMultivar