PERGAMON

Computers, Environment and Urban Systems
24 (2000) 283–303

Computers, Environment and Urban Systems

www.elsevier.com/locate/compenvurbsys

Techniques for defining school catchment areas for comparison with census data

J. Pearce*

North West Regional Research Laboratory, Department of Geography, Lancaster University, Bailrigg, Lancaster LA1 4YB, UK

Abstract

School performance tables have been used in the UK with the intention of giving parents an indication as to how well local schools are performing in order that they may make a more informed decision about which school to send their child to. It is likely that the performance measurements are not only a result of the success of the school but also dependent on the background of the children that attend it. School performance tables do not provide information upon the social context in which the school is set. It is, therefore, important to be able to provide information about the social characteristics of a school's catchment area. This is technically difficult because firstly school catchment areas are not strictly defined and secondly census units have no link to catchment geography of a school. This paper offers three techniques that may be suitable for establishing the link between the point information available on the performance of the school and the areal census data. Each technique is tested to see how well this link has been established in each case. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Voronoi polygons; Primary school league tables; Location-allocation

1. Introduction

School performance and school league tables have become a highly emotive issue in the UK. Following the Conservative government's decision to give parents a greater choice in the school to which they sent their children, it was decided to publish primary school performance indicators tables in 1996. This followed on from tables previously published on secondary school performance. League tables are intended to give an indication as to how well a school is performing and are used to compare the success of neighbouring schools. The tables are intended to allow

* Tel.: +44-1524-593762; fax: +44-1524-847099.
E-mail address: James.Pearce@lancaster.ac.uk (J. Pearce)

PII: S0198-9715(99)00045-9

parents to make a more informed decision on choosing the most suitable school for their child. As a result of this choice, there has been an increased interest amongst parents as to the success of the schools in their local area.

Critics (e.g. Gibson & Asthana, 1998) have argued that league tables are flawed in terms of the way in which they are measured and in that they give no indication of the context in which the school is placed. In order to tackle the second issue, it is important to develop techniques that can accurately allow a link to be established between the data set on school performance and contextual information such as census data. Once techniques are available, the school performance results can be compared to a range of socio-economic indicators. If this type of information were available, it is more likely that parents will be able to consider the success of the school, not just based upon the position of the school in the league table but also by putting the perceived success or failure of a school into the correct social context. It seems unreasonable to make direct comparisons between schools in prosperous middle class areas with nearby inner city primary schools without considering these differences in the socio-economic background. It is possible that a school fairly high up the school table may not in fact be doing as well as a school lower down the table that is set within a poorer neighbourhood. This paper puts forward three potential methods of differing degrees of sophistication and compares how successful they are.

2. Background

Primary school performance tables provide details of school performance as measured by the number of children aged 10 or 11 (key stage two) who satisfy a government-defined expected minimum standard. The tests were formulated primarily to provide indicators to Local Education Authorities who in turn were obliged to make the results available to the public in the form of the 'league tables' as they have commonly become known (as they are published from highest score at the top to lowest score at the bottom). This information is then interpreted by schools, parents and the public as a measure of the success of the school.

It was the intention that the results would provide an incentive, through the competition between schools that would arise, for schools firstly to attempt to raise standards but also to provide parents with information to help them choose which school to send their children to. The results show the percentage of students who achieved this government-defined minimum in three key subject areas: English, maths and science. A further variable called the total score is included which is a summation of the English, maths and science percentage scores giving a total out of 300. All children sat the test on the same day in May 1997. Children who were absent that day were assessed by the teacher over the school year. The data are available for all schools in England with the exception of special schools and those schools that had less than 11 pupils sitting the test. The results were compiled and verified by the Department for Education and Employment in November. The results were published digitally on *The Times* web site on 27 January 1998.

The format and presentation of the results in this league table format encourages government and parents to interpret the position in the league table as being largely indicative of the efforts and success of the school. However, the results of the tests clearly show that there are large variations in school performance across most scales of analysis. Fig. 1 shows the variation in the total score across Lancashire. Even across this region there are clear differences in levels of achievement. Low scores are common in the urban centres of places such as Preston, Blackburn and Burnley while scores are higher in rural and suburban areas around the urban centres of the county. The wide disparity in the total score between the schools would appear to have a strongly defined geography. There is clearly not a random distribution in the total scores and there are clear clusters of high and low scores in the region.

There is, therefore, a line of argument that the results are to a greater or lesser extent dependent upon the social make-up of the area in which they fall. The performance of a school is strongly contextualised in the background of its catchment

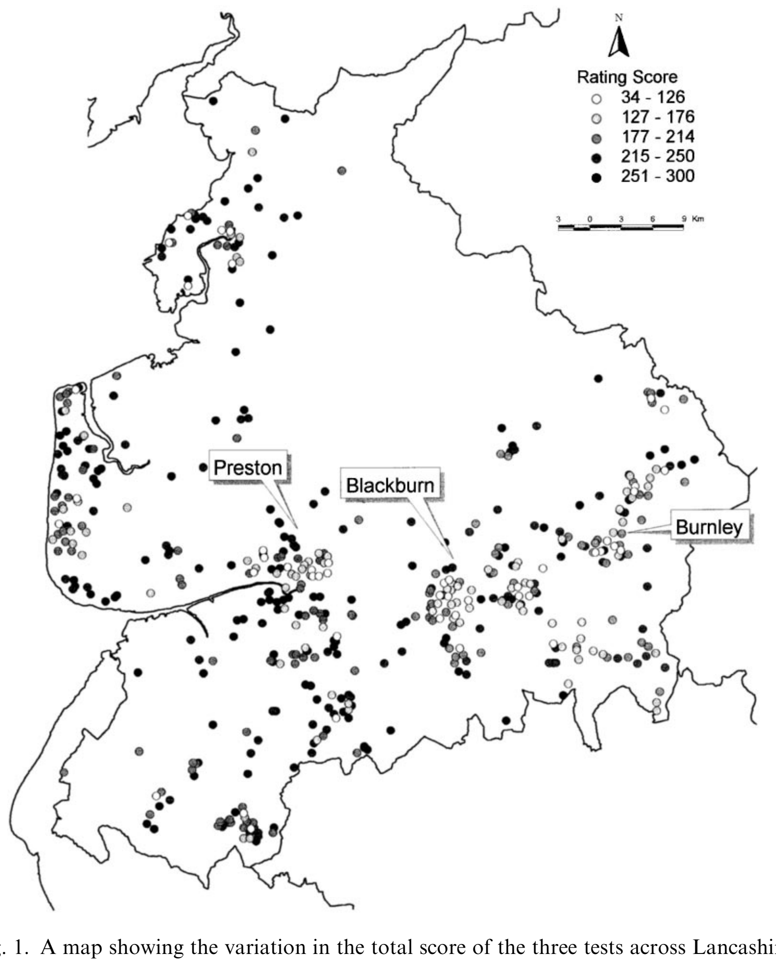

Fig. 1. A map showing the variation in the total score of the three tests across Lancashire.

area (Gibson & Asthana, 1998). Various studies such as Conduit, Brookes, Bramley and Fletcher (1996) and McCallum (1996) have shown that there is a clear relationship between school performance and a number of indicators of poverty and social deprivation. It is widely accepted that pupils from homes that have a higher income, with well-educated parents who encourage and are able to help their children to succeed at school, are more likely to do well and perform better in indicative tests.

A number of studies from the educational literature have explored the factors that influence children's educational development. Wedge and Prosser (1973) looked at the incidence of social disadvantage in a sample group and the attainment of those children considered to be disadvantaged. Three criteria were considered to be crucial to define disadvantage: family composition (five or more children or a single-parent family); low income; and poor housing. They found that one in six of the extremely disadvantaged sample was receiving special help compared to one in 16 for non-disadvantaged groups. On average, disadvantaged children were 3.5 years behind other children in their reading scores. Mortimore and Blackstone (1982) claim that material disadvantage may both limit access to education and make it more difficult for children to benefit from it. Douglas (1964) carried out principal component analysis on environmental and parental factors. Three factors emerged with heavy loading: housing (overcrowding); parental interest (support and attitude to the child's education); and the educational level of the parents and the occupation of the husband. Garner and Raudenbush (1991) tested for the existence of neighbourhood effects on educational attainment among some 2500 young people who left school between 1984 and 1986 in one educational authority in Scotland. Using 1981 census data, they developed hierarchical linear models to explore the relationship between home neighbourhood and educational attainment. The authors discovered a significant negative association between deprivation in the home neighbourhood and educational attainment.

There is a clear distinction to be made between schools that are successful and schools that are high up the league table. Likewise, a school that is low in the league table is not necessarily failing. A useful source of contextual information that may help to explain some of the variation in school performance schools, outlined above, can be provided by the 1991 UK census. A variety of variables that can be linked to educational achievement and social deprivation are available. Technical difficulties arise when trying to establish relationships between census data and the school performance data while attempting to link together the two rather different geographies.

3. Linking primary school results to census data

It is clear that there is a need to link areal census data and the point school performance data. The development of techniques to do so would enable school performance indicators to be contextualised in the social environment in which they operate. This would allow comparisons to be drawn between the schools in the data set over different scales of analysis. By comparing results between schools, more appropriate conclusions could be drawn on the level of achievement of each school.

There are clear technical obstacles to overcome in making the link between the two types of data. In effect, what is required is a set of key census values that are indicative of the social characteristics of the school catchment areas. Two key issues make the definition of school catchment areas problematical. Firstly, the concept of a school catchment area is a vague entity with no known or delimited boundaries. School catchment areas have a rather ambiguous geography. This is because the decision as to which children attend which school can in some cases be the choice of the parents and in other instances be determined by the schools themselves. The choice that parents make is often related strongly to the demand for places at the school. This means that children can travel long distances to a school and, therefore, makes it difficult to make assumptions about the characteristics of school catchment areas. Secondly, even if catchment areas were defined and known, census boundaries are unlikely to coincide with school catchment areas. In order that a link can be made between the school performance data and census data, it is necessary to make judgements about which census areas should be aggregated to provide information on which schools. This problem is represented in Fig. 2, where it can be seen that there is no common relationship between the census geography and the catchment area. The school catchment boundary can potentially incorporate both entire and partial census areas making the precise calculation of census data for the catchment area problematic.

The most accurate way to address this issue is to use the smallest unit available from UK census data: the enumeration district (ED). An ED represents a population of a few hundred people. For each ED a centroid location is available (easting and northing) that is an estimate of the centre of the population. Techniques can then be derived to allocate ED centroids to each school. This can be seen in Fig. 3

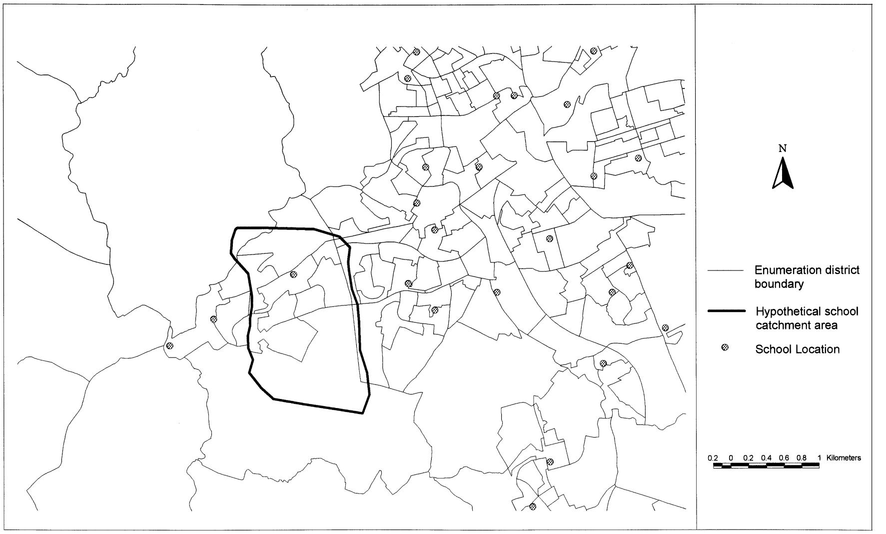

Fig. 2. A hypothetical school catchment area that shows the relationship between the catchment and the census units.

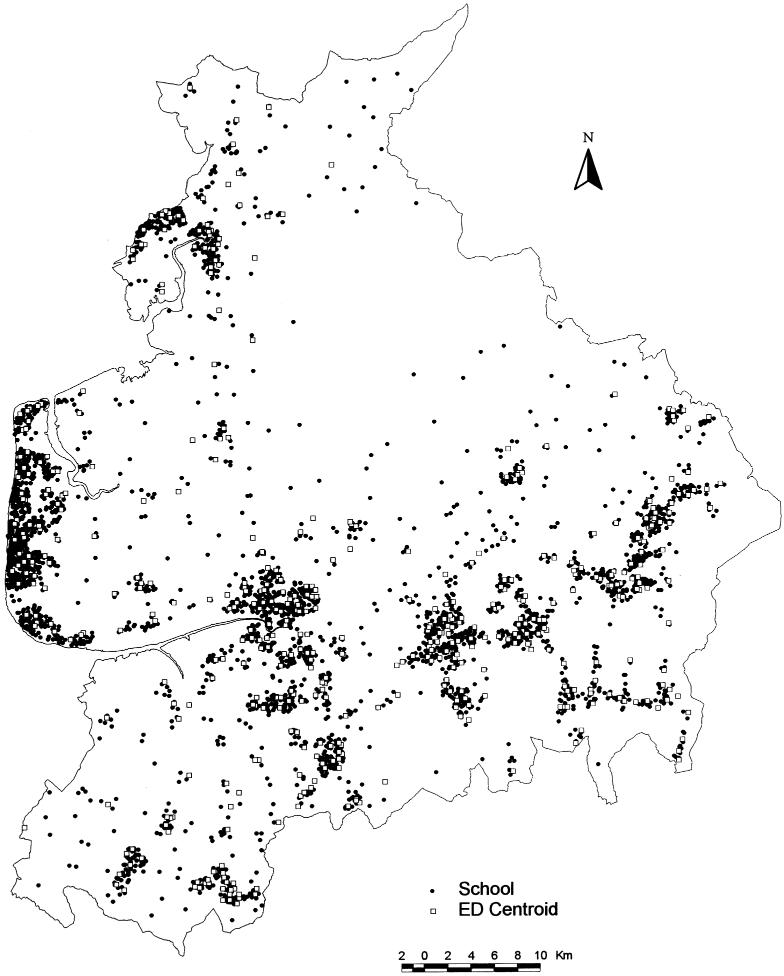

Fig. 3. School locations and enumeration district (ED) centroid locations in Lancashire.

which shows all the enumeration district centroids and school locations in Lancashire. In Lancashire there are 491 schools and 3047 ED centroids. The problem is, therefore, deciding how to allocate which ED centroids to which schools. Potentially the most suitable method of allocating ED centroids to schools, and hence defining school catchment areas, is to acquire information on the address of each pupil that attends each school in a study area. If the postcode of the child's address is known then an easting and northing can be obtained. From this it can be calculated which school the majority of the pupils in each ED attend and hence assign the ED to that school. This would provide an accurate depiction of the spread of children around schools and allow a relatively accurate description of the social characteristics of the school's catchment area.

This approach is very time consuming and also problematical in so far as the availability of such data is strictly controlled and access is largely limited by the schools and local education authorities. These data tend not to be held centrally and could only be obtained from each of the individual schools who are not obliged to provide these data. It is, therefore, desirable to use automated methods of allocating ED centroids to schools. This will allow schools that fall in similar or different types of catchment areas to be compared. Clearly, the method used can potentially have profound implications for any analysis of the relationship between school performance and the social environment in which the schools fall. This study considers three techniques that attempt to overcome this issue. Each one is discussed and then attempts are made to consider which one is the most successful by comparing the results of a regression model. The regression model calculates to what extent a set of census variables can be used as a predictor of the primary school performance results. The work focused on one county, Lancashire, as the computational time that it took to run some of the techniques was too great to consider using a larger data set.

4. Methodology

The method used to make a spatial association between the school performance data set and the census data will clearly have large implications for the final results. Different methods can associate different ED centroids and different numbers of ED centroids with a school. Depending upon the assumptions made, there are a number of methods that can potentially be used. The potential variation in the size of the catchment area and the ED centroids allocated to a school is great. Three different methods for allocating ED centroids to census data are discussed below. The first two methods treat the issue as a 'geometrical' issue and the space is dissected according to a set of assumptions. The third method uses linkages between the ED centroid and the school to treat the issue as one of location-allocation.

4.1. Voronoi polygon approach

Voronoi polygons have been used in a number of different applications in the social and environmental sciences to determine a mutually exclusive partition of space (Okabe, Boots & Sugihara, 1994). Voronoi polygons are a method for dividing an area into regions so that all locations closest to a particular sample point are enclosed within a single polygon (McDonnell & Kemp, 1995). Okabe, Boots and Sugihara (1992) highlight the wide range of applications in which Voronoi polygons have been utilised, including areas as diverse as archaeology and transportation. Applications have included their use to estimate unknown census boundaries (Boyle & Dunn, 1991a), unknown unit postcode boundaries in the UK (Boyle & Dunn, 1991b) and the estimation of unknown data such as rainfall values in Kansas and Nebraska (Haining, Griffith & Bennett, 1984).

Each boundary arc of the polygon is calculated by finding the perpendicular bisector of straight lines between pairs of neighbouring points. As a result, the area of the polygon is dependent upon the spacing of the points. A Voronoi polygon was constructed around each school to define a 'catchment area' for each one. This seems a sensible approach in the UK as traditionally children will tend to go to their nearest primary school. For instance, Lancashire schools that are over-subscribed prioritise their intake firstly through siblings currently in the school and then the geographically closest school. Fig. 4 shows the Voronoi polygons around the schools in Lancashire. The pattern seems broadly sensible as the area of the catchment increases with the sparser distribution of the schools, i.e. schools in rural areas have a larger catchment area that is more sparsely populated. One of the issues that this approach does not resolve is that the school catchment areas are not related to the size of the school as defined by the number of pupils. It seems intuitively sensible that larger schools will have larger catchment areas than neighbouring

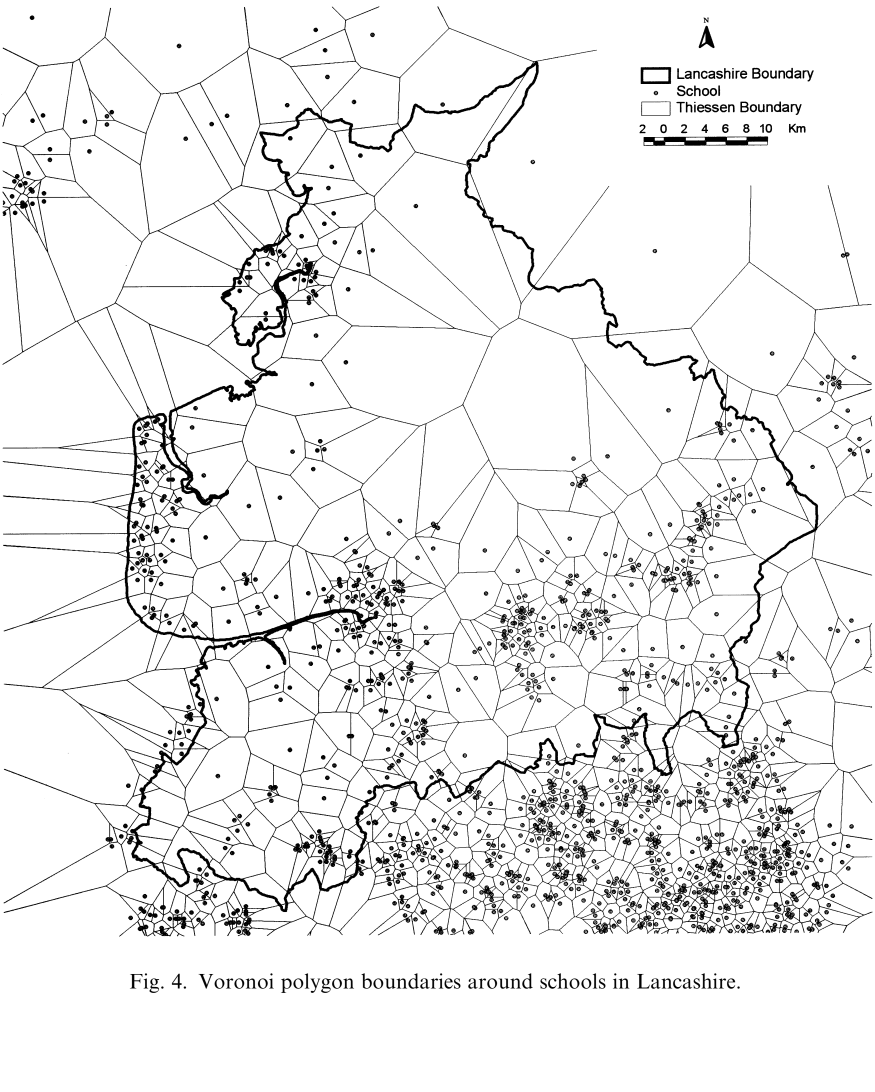

Fig. 4. Voronoi polygon boundaries around schools in Lancashire.

schools and that it is not enough just to rely upon the spatial distribution of the schools in the area.

4.2. Weighted Voronoi polygon approach

An alternative approach that attempts to address some of the issues outlined above is to construct weighted Voronoi polygons around each of the schools. Each polygon will be related to the size of the school. This seems appropriate, as other things being equal, larger schools will tend to have a larger catchment area. The boundary between two neighbouring schools will be weighted according to the relative size of the two schools.

The weighted Voronoi polygon approach has been used by Vincent and Daly (1990) to model the sphere of influence of different towns in the northwest of England for use in defining the pull of a town to shoppers. Boots and South (1997) use higher-order multiplicatively weighted Voronoi polygons to define store trade areas using assumptions on customer behaviour and characteristics of the store. This approach allowed the authors to incorporate more complicated overlapping trade areas. Cox and Agnew (1974) used weighted Voronoi polygons to divide the whole of Ireland into theoretical counties, which are then examined in terms of their relations to actual counties.

The construction of weighted Voronoi polygons is computationally more difficult than the Voronoi polygon method. The polygons cannot be easily constructed within any standard geographical information system (GIS) packages. Instead, specialist packages such as GAMBINI (Tiefelsdorf & Boots, 1997) are available to complete the task. GAMBINI will construct weighted Voronoi polygons that are proportional to a defined attribute, in this case the number of pupils. It will determine a 'sphere of influence around the generator points'. Multiplicative weighted Voronoi diagrams can be regarded as the equilibrium of a wave diffusion process around the generator points where the spread is proportional to the weight (Tiefelsdorf & Boots, 1997). Difficulties with the software arose when trying to create polygon coverages for use in standard GIS packages. Polygons were not properly built and it was difficult to maintain the topology as disjoint polygons were assigned different ID values. Other packages are available that will compute weighted Voronoi polygons and are likely to overcome some of these technical difficulties (Boots & South, 1997).

The resulting school catchment areas defined using this approach in Lancashire are shown in Fig. 5. It can be seen that larger polygons are constructed around the larger schools relative to their neighbouring school. This is to say that the position of the boundary between two neighbouring schools will be dependent upon the relative size of the two schools. This approach seems to overcome some of the problems of the standard Voronoi polygons as larger schools will tend to have a larger polygon and hence a larger catchment area. This is especially true in urban areas such as Preston and Blackburn where equal-sized catchment areas seem unrealistic. Larger schools are likely to have a greater pulling power and hence a greater catchment size. This phenomenon is reflected if weighted Voronoi polygons are used

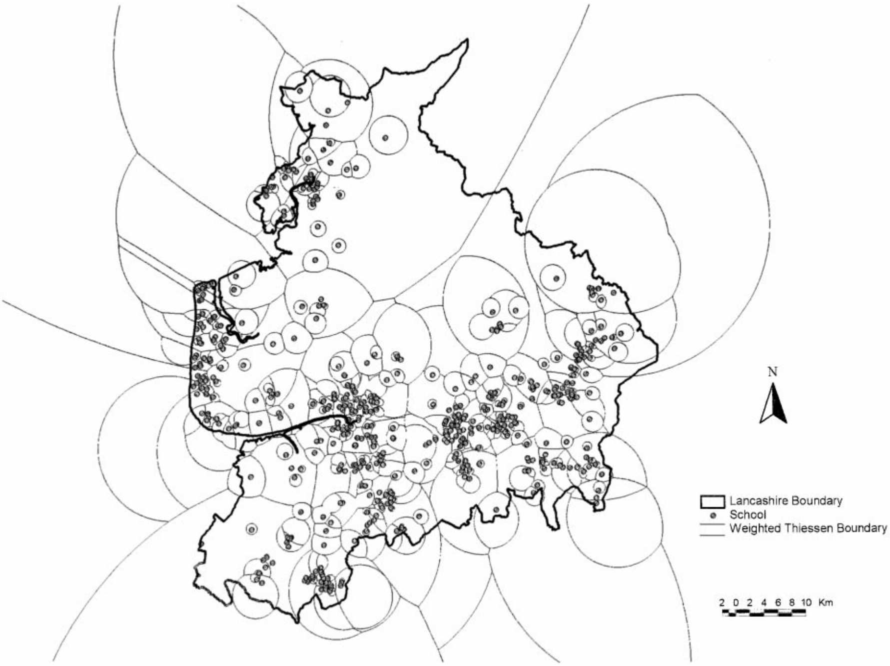

Fig. 5. Weighted Voronoi polygon boundaries around schools in Lancashire.

instead. Once again, as with the Voronoi polygon approach ED centroids that fell within the weighted Voronoi polygon were assigned to that school.

Fig. 5 raises some of the issues that need to be addressed when implementing this approach. It can be noted that some of the polygons around the schools are so small due to the small size of the associated school proportional to neighbouring schools. As a result, no ED centroids fall within the polygon. In fact some polygons are so small it is not apparent from the diagram that they exist. This is unsatisfactory as no EDs will be allocated to that school and hence no catchment information on the social characteristics is available for those schools. Equally, the ED centroids will instead be associated with another school. It is also noticeable that some of the polygons are so large that they tend to engulf some of the smaller polygons. In effect the smaller polygon sits within the larger polygon. It is debatable how realistic this situation is. It raises the question of whether a great number of students will in effect travel through another school's catchment area to reach their destination school.

4.3. Location-allocation

Methods that treat the issue of allocating census centroids to schools by defining suitable spatial areas where the geometrical dissection is based upon a set of assumptions and rules is clearly useful but it is possible that other important factors could be incorporated into the model. The third approach differs from the previous two as it does not consider the issue as one exclusively of dissecting the space in the most suitable way. It instead considers the issue as one of supply and demand. This is to say that there is a capacity (the number of pupils aged 10 or 11 sitting the test) and a known resource (the number of children of the correct age who live in the ED from the 1991 census). The 1991 census was the best available estimate of the number of children but it will not be totally accurate due to migration in the period between the collection of the two data sets; however, it is necessary to assume an approximate balance. It is possible to determine how they can be best allocated as limited by these attributes and also by distance between the school and the ED centroid. The location-allocation algorithm can, therefore, be used to match the ED and school totals.

The methodology used in this approach is represented in Fig. 6. A network was constructed between the schools and the EDs by assigning arcs from the schools to all ED centroids within 15 km of each school. This was considered to be the greatest distance that a child was likely to travel to a primary school. It was important to use a suitable maximum travel distance because if a lower distance was used then this would have implications upon the results as EDs may be wrongly assigned. Using the location-allocation functionality in ArcInfo, the children in the EDs were allocated to the schools. The allocation process was impeded by the distance between the school and the ED centroid. Therefore, the objective functions were to minimise the distance between the school and the centroid and to allocate the right number of children. The result of this was that closer ED centroids and hence children living closer to the school could be allocated more easily than ED centroids

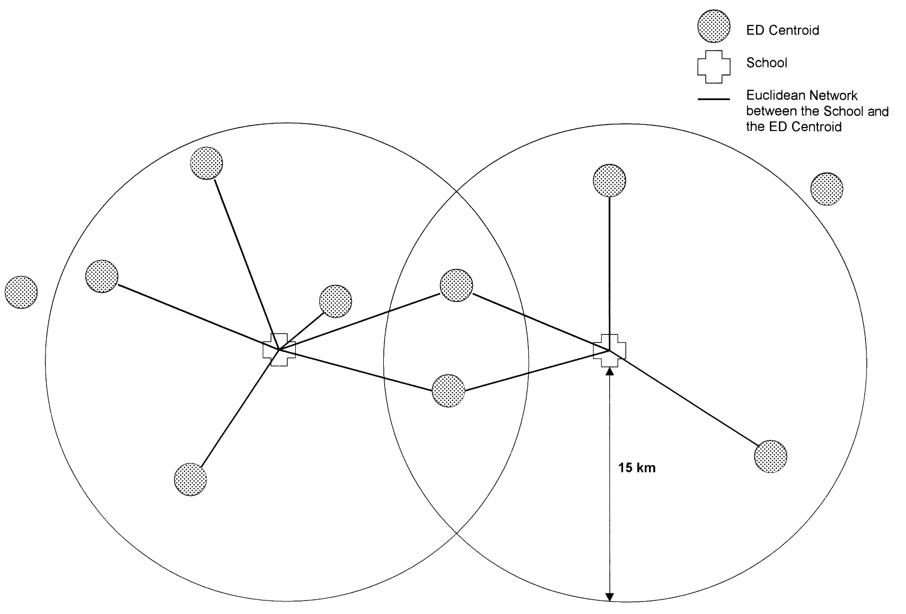

Fig. 6. Allocating enumeration district (ED) centroids to a school using a network based upon distance and allocation of the supply and demand.

that were further away from the school. The assumption is still that children are more likely to be sent to a closer primary school.

The approach seems to at least offer an alternative if not extend beyond the above spatial tessellation methods outlined above. The location-allocation approach takes into account how many children there are living in each ED, how far the children live from the school (Euclidean distance) and what the capacity of the school is (in terms of the number sitting the tests). The methodology is clearly limited firstly by the accuracy of the 1991 census data in 1998 as it possible that there have been substantial fluctuations in the population of children in this period. Secondly, the Euclidean distance is used instead of incorporating the road network travel distance between the school and each ED centroid. However, this approach is still potentially a better estimate than the geometric approaches outlined above.

5. Comparison of methods

A direct comparison of how well the three methods perform is difficult to undertake without data on the address of every child in each school. With this information, it would be possible to calculate which school each ED would best be associated with and compare this to the three modelled results. This would help to determine which method made the most accurate link between the school and the census information.

As discussed above, access to the data on individual pupils, especially within Lancashire, is difficult because it is not centrally held and schools are often unwilling to provide this information. Clearly, the collection of these data from every primary school would be both difficult and time consuming. As a consequence, methods were developed in an attempt to give a proxy measurement of the success with which the three methods linked together the data on school performance and census information.

One method that could be used to compare the success of the three models in defining school catchment areas is to firstly calculate the weighted values for a number of indicative census variables using each method in turn and then to compare how good the results are as a predictor of school performance. These results would give an indication of how successfully the link has been made between the two sets of data. The results of this can be used to determine how good they are as a predictor of school performance. From this, one can start to judge how well each model performs in making the link between the school and the census data in order to define the catchment area of the school. In the past, a wide range of census variables have been linked to direct indicators of school performance and also as indicators of social deprivation. The census variables used in this study are shown in Table 1. Each method allocated a different set of ED centroids to each school so for almost every school, three different catchment areas have been defined. For each of these methods, the social characteristics were computed for the catchment area of each school to provide a weighted percentage value for each census variable. These percentage values were weighted by the number of children sitting the test in each ED.

Table 1
Census variables used to define the social characteristics of the school catchment area

Abbreviation	Census variable
white	Percentage of white residents
black	Percentage of black residents
asian	Percentage of Asian residents
nocar	Percentage of households without a car
twocar	Percentage of households with two or more cars
locau	Percentage of residents living in local authority owned housing
priv	Percentage of residents living in privately rented accommodation
lonep	Percentage of lone parent households
ovcr	Percentage of overcrowded households (greater than one person per room)
noheat	Percentage of households without central heating
ftfem	Percentage of females with children aged 0–15 working full time
llti	Percentage of residents with a long term limiting illness
a511	Percentage of population of primary school age
unemp	Percentage of economically active who were unemployed
migr	Percentage of residents who were migrants from a different ward
qualifs	Percentage of residents with higher educational qualifications
profman	Percentage of residents who are professional and managerial workers
unsk	Percentage of residents who are unskilled workers

Using the census information on school catchment areas computed above, it is suitable to fit regression models that can test how well a set of the variables can be used to predict school performance scores. The success of the regression models can give an indication as to how well the models have established a link between the school and the catchment areas, if one assumes that the census variables used should have a strong causal relationship with school performance. It should be noted that the strength of the R^2 may not be a simple indicator of the relationship between school performance and the socio-economic variables. Other characteristics than the ones used in the models may have a significant influence on the test scores. If this is the case then it is possible that the results may in some way be distorted. However, the fact that all the methods show a relationship to many of the socio-economic variables is evidence that the relationships observed are real, and it is reasonable to prefer a method which makes the relationships stand out from the complicating factors.

For this analysis, it was suitable to use a logit model as it has the effect of automatically weighting each observation, according to the number of pupils in each school (Wrigley, 1976). The dependent variable in each model was the total score achieved by each school. Where no centroids fell within the polygons created in methods one and two, the school was removed from the regression model. The results of fitting the logit model to the three approaches is shown for the Voronoi polygon method in Table 2 and with the weighted Voronoi polygon method in Table 3 and for the location-allocation approach in Table 4.

From Table 2, it can be seen that the Voronoi polygon model has significant values for 15 of the 18 census variables. This model has a pseudo R^2 of 39.6%. The

Table 2
Logit model, using Voronoi polygon linkage

	Estimate	SE	Parameter
1	1.008	0.05103	1
2	0.02912	0.004197	TWOCAR
3	0.007844	0.002540	QUALIFS
4	0.003102	0.0005797	WHITE
5	−0.1110	0.02253	BLACK
6	0.006513	0.001671	NOCAR
7	−0.004529	0.001068	LOCAU
8	−0.007881	0.001840	PRIV
9	−0.01619	0.004524	LONEP
10	−0.007354	0.001027	NOHEAT
11	−0.009612	0.001328	FTFEM
12	−0.01411	0.004505	A511
13	−0.01307	0.003271	UNEMP
14	−0.09789	0.01724	MIGR
15	0.004999	0.001376	PROFMAN
16	−0.005407	0.001781	UNSK

Table 3
Logit model, using weighted Voronoi polygons linkage

	Estimate	SE	Parameter
1	2.278	0.1526	1
2	−0.1040	0.02301	BLACK
3	−0.005677	0.002471	NOCAR
4	0.002645	0.001294	LOCAU
5	−0.006684	0.001969	PRIV
6	0.02830	0.005241	OVER
7	−0.007929	0.001133	NOHEAT
8	−0.003397	0.001524	FTFEM
9	−0.01637	0.005635	LLTI
10	−0.09128	0.008025	A511
11	−0.01149	0.004417	UNEMP
12	0.006912	0.001957	PROFMAN
13	−0.007186	0.001919	UNSK
14	0.007611	0.002907	QUALIFS

pseudo R^2 provides a statistic that is analogous to R^2 and can be used to measure the goodness of fit. The logit model for the weighted Voronoi approach, shown in Table 3, is a slightly more economical model with a pseudo R^2 of 40.3%. The final approach, the location-allocation method, provides a model with 12 significant values (Table 4) but this time with a pseudo R^2 of 45.8%. The variation in the pseudo R^2 for the three different methods of associating ED centroids with schools gives a strong indication of how well the link between the two data sets has been established. Although none of the three methods provides a very high R^2 it is clear

Table 4
Logit model, using allocation linkage

	Estimate	SE	Parameter
1	−1.219	0.07107	1
2	0.01504	0.001943	QUALIFS
3	0.03443	0.001674	WHITE
4	−0.09308	0.02306	BLACK
5	0.04103	0.002595	ASIAN
6	−0.003304	0.001253	LOCAU
7	−0.006278	0.001786	PRIV
8	−0.009867	0.0009748	NOHEAT
9	−0.006274	0.001574	FTFEM
10	−0.03485	0.005180	LLTI
11	−0.06573	0.007498	A511
12	−0.007967	0.003226	UNEMP
13	−0.01111	0.003196	PROFMAN

that there has been a highly significant link established between the census information and the school performance tables. It is possible that a very high pseudo R^2 is unlikely because there are no school characteristics included. Of the three methods, the location-allocation approach provides a far higher R^2 value than the other two approaches which suggests that the link between the two sets of data is stronger.

The five highest and five lowest residual values are shown for each of the regression models in Tables 5–7. It is not possible to directly compare the values between the models because a different number of schools were allocated using each method. However, the high positive scores indicate that the school is performing better than predicted by the model and the highly negative scores indicate that the school is doing worse than predicted. The residual values can give an insight into the factors

Table 5
Highest and lowest residuals, using Voronoi polygon linkage

School name	Residual value
Highest	
1. Holy Trinity Church of England, Burnley	9.093
2. Duke Street, Chorley	7.883
3. Cobbs Brow Cty, Skelmersdale	7.468
4. St. Peter's Church of England, Morecambe	5.419
5. St. Edwards's Roman Catholic, Darwen	5.365
Lowest	
1, Grosvenor Park Cty, Morecambe	−11.864
2. Hawthorns Cty Jr., Blackburn	−10.243
3. Hargher Clough Jr., Burnley	−9.350
4. Chaucer Cty, Fleetwood	−8.640
5. Spring Hill Cty, Accrington	−8.466

Table 6
Highest and lowest residuals, using weighted Voronoi polygon linkage

School name	Residual value
Highest	
1. Holy Trinity Church of England, Burnley	8.552
2. Lever House, Preston	7.037
3. Lancaster Road Cty, Morecambe	5.954
4. St. Mary's Roman Catholic, Fleetwood	5.919
5. St. Edwards's Roman Catholic, Darwen	5.887
Lowest	
1. Grosvenor Park Cty, Morecambe	−10.946
2. Chaucer Cty, Fleetwood	−10.197
3. Grange Park Jr., Blackpool	−8.899
4. Hargher Clough Jr., Burnley	−8.620
5. West End Cty, Morecambe	−7.964

Table 7
Highest and lowest residuals, using allocation linkage

School name	Residual value
Highest	
1. Ansdell Cty, Lytham St. Annes	8.685
2. Cobbs Brow Cty, Skelmersdale	7.288
3. Anchorsholme, T-Cleveleys	6.077
4. Lytham Church of England, Lytham St. Annes	5.660
5. Bolton-le-Sands, Church of England, Carnforth	5.252
Lowest	
Brockholes Wood Cty, Preston	−9.657
St. Joseph's Roman Catholic, Lancaster	−9.257
Stoneyholme Cty, Burnley	−7.678
Ightenhill Cty, Burnley	−7.231
St. Joseph's Roman Catholic, Preston	−7.176

that are influencing the success of the models. The location and nature of the schools that have the extreme residual values may indicate the types of school that are not complying with the assumption that the characteristics of the surrounding neighbourhood influence school performance.

6. Discussion

The results of the regression analysis obtained in the previous section indicate that the method used for making the association between school performance and census information has clear implications upon the final results. Each model used the same data but made different associations between the two sets of data and, therefore, the results provided different R^2 values. If the results are to be used to provide an indication of the success of each model it would seem that the third approach, that of location-allocation, provides the clearest results. Although the R^2 is less than 50%, it at least proves that there is a clear link between school performance and the social characteristics of the school's catchment area. This in turn suggests that the environment in which the school has to work has implications upon the position in the league table that the school will be placed.

The results of the comparison above show that the link between the school performance results and the census data is not as strong as one might expect using any of the three methods. This could be due to a variety of reasons. Firstly, it may be due to the weakness of the techniques in that they do not depict the geography of the catchment with enough precision. As a result ED centroids are not allocated to the most suitable school. If this were true then it would be no surprise that the R^2 values are so low. Secondly, the results may be a consequence of not choosing the most indicative census variables. The results may indicate that it is other variables that influence the results and hence other social reasons that determine school performance. Furthermore, the data use 1991 census variables which was 6 years before the tests were sat. It is possible that in that time there were changes in the social make-up of the catchment areas, particularly the number of 4-year-olds in the 1991 census may no longer be a suitable measurement of the number of 10-year-olds in the local area.

Thirdly, it is likely that the results of the analysis are influenced by more unusual circumstances such as the number and the position of church schools in the data set. This is supported by the tables of residuals (Tables 5–7). For each of the methods of linkage, it is common to find a church-supported school falling amongst the highest residual factors. In the three tables, eight of the 15 high residual values are church-supported schools. These schools tend to fall in disadvantaged areas and, therefore, the model predicts that they would not perform very well. The reason for the difference between the observed and the predicted can be attributed to the different size and extent of the catchment areas. It is clear that church schools have more widely dispersed catchment areas than other primary schools. This is highlighted in Fig. 7 where data were acquired from the Cathedral school, Lancaster, on the postcode of each child in the school. There were 171 children in the school of which 16 had missing postcodes. It is clear that if a catchment area were to be defined from this information, neither the Voronoi or weighted Voronoi polygon approaches are likely to be sufficient because neither will encompass the majority of these pupils. This can be seen in Fig. 7 where the associated Voronoi polygon and weighted Voronoi polygon of the school surround the point that represents the Cathedral school. It is clear that neither of these polygons is likely to be even a near approximation of the catchment area of this school. The map shows that the catchment area is highly dispersed, probably due to its status as a Catholic Church school. Catholic parents are likely to be more willing to send their child a greater distance to a Catholic school than they would if this were not a Catholic school. The three techniques that have been developed each use the distance from the school as a limitation to the association between schools and ED centroids. As a result, the

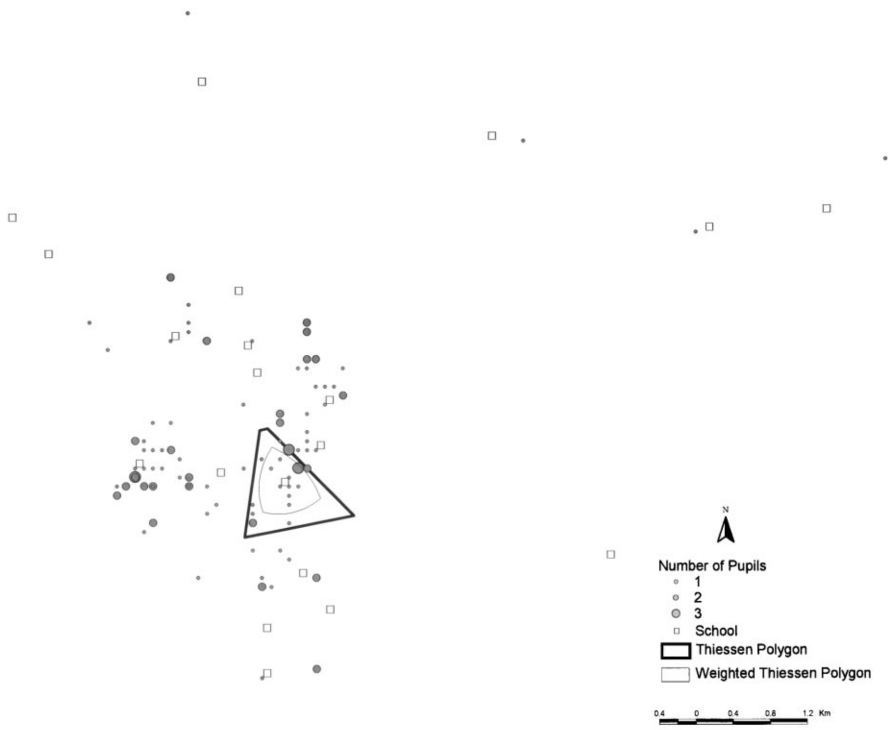

Fig. 7. Distribution of children attending a Roman Catholic school in Lancaster.

widely dispersed nature of church school catchment areas cannot be easily incorporated into the models. It may often be the case that children attending church schools come from different backgrounds than the majority of children attending nearby schools. If this is the case then none of the models can account for this phenomenon as the social characteristics of the school catchment are computed from only the local census values. This helps to explain the high residual values for many of the church schools. In line with this, it would be useful to investigate the influence that church schools have upon the results produced in this paper and also to attempt to compute the more complicated catchment areas of the church schools.

A further issue to consider is the influence that the schools that are not included in the data set could exert upon the results. Schools that had less than 10 children sitting the test are not included within the school performance tables. It is clear that if they were included within the models then school catchment areas and shapes would be altered. There is clearly some interesting work to be undertaken here. In line with this, there are further inaccuracies either within the data set or within the conversion of the postcodes into an easting and northing value. There are examples of schools in Lancaster, for example, where a school is clearly located in the wrong place on the map. This could be a result of the initial data set containing the wrong postcode or of errors associated with the postcode look-up table. Furthermore, it is commonplace that schools can have either the same postcode or different postcodes but the same easting and northing. When there is an identical match in the geocoding of two or more schools, the first two models can only draw one polygon and hence the catchment areas are identical. An example of this takes place within the data set in Carnforth, Lancashire, where three schools are assigned the same location. The method using the location-allocation techniques seems better suited to solving these issues as it does not rely solely upon the dissection of space and can model two closely located schools, which may help to explain its better performance in the regression analysis.

Other areas that need to be researched include a consideration of how school performances have changed through time. The data set that has been used has been available digitally for all years from 1996 up to and including 1998 and it is likely that the 1999 results will follow a similar format. Clearly, a steady improvement in school performance can be indicative of success within the school, changes within the catchment area of the school or alternatively just random fluctuations. Either way, it is important to know and understand any such changes. Likewise reductions in school performance may have a defined geography that needs to be investigated. Patterns may be indicative of changes locally in the areas surrounding the school, changes in the resources of the school and changes within the schools themselves.

The geographical area in which the models are applied is likely to influence the degree to which they are successful in making the association between two sets of data. The results suggest that the techniques that have been developed work better in different areas. It seems intuitive that in rural areas the distances between primary schools will be greater than in urban areas and so a child is more likely to go to their closest school. For the situation to be different the pull of another school must be large. Conversely, in more urban areas, there is a greater density of schools and hence the feasibility of attending the non-nearest school is higher. Therefore, it seems likely that the geography of primary school catchment areas is likely to be more complicated in urban centres and hence more difficult to model. The evidence for this is supported by the results shown above. They show that a more obvious dissection of space as determined by the location and density of the schools does not really address all the issues involved in defining school catchment areas. Instead, the method of location-allocation helps address some of these issues by modelling where the students are and where they have to go.

The results produced from each of the models indicate that there is a clear link between school performance and the social characteristics of the area. They also show that the design of the model has important implications upon the results. All of the techniques seem to have been successful to a certain degree with the location-allocation approach perhaps the most suitable, especially in urban areas. Further research could focus upon refining the methods used and upon special cases such as church schools.

7. Conclusions

The results of this study have demonstrated the need to contextualise a school within its local area. School performance tables can provide a false impression on how well a school is performing. This paper has suggested that the performance indicators cannot be considered in isolation from the social characteristics of its local area because the background of the pupils has a clear influence upon how well a school performs. It is, therefore, important to develop methods that link together the different geographies of the school performance tables and the census data.

This paper has suggested and tested three approaches that may be used for achieving this. Logit regression analysis has been used to try and establish how well the link has been made between the two data sets. The results of the analysis have shown that there is a clear relationship between school performance and the measures of social deprivation. All three methods show clear imperfections but they have still managed to establish links between the two sets of data. The location-allocation approach seems to offer the greatest potential for establishing a strong link. This approach does not treat the issue as one of geometry but rather considers it as one of supply and demand. The improved results are likely to be because the catchment area of the school is more complicated than dissecting the space around the schools. Instead, a consideration of where pupils are and where they need to go seems to provide a better model.

The relationship between test scores and socio-economic variables cannot be known perfectly because of the uncertainties in the linkage between schools and their catchment areas. However, the fact that three reasonable methods of allocating census zones to schools all give roughly similar results is evidence that there is such a relationship.

Further research could usefully focus upon other factors such as the influence of church schools upon the modelled results as it is apparent that church school catchment areas tend to have a rather different geography. Furthermore, it would be useful to consider the influence of schools that are missing from the dataset, as these often tend to occur in rural areas and also to consider any schools that are wrongly geocoded. By taking these issues into consideration it may be that the results of the model can be improved.

Acknowledgements

Thanks to Robin Flowerdew for his help and advice with writing this paper. This paper followed on from work for the Workshop on Investigating Locational Data held at Lancaster University, 9–10 July 1998, sponsored by the Quantitative Methods Research Group of the Royal Geographical Society/Institute of British Geographers. RRL.net, and the Regional Science Association International (British and Irish Section). The paper was presented at the European Research Conference: Socio-Economic and Geographic Information Systems in Espinho, Portugal, 22–27 May 1999, sponsored by the European Science Foundation and the Euroconferences Activity of the European Union. I acknowledge the use of Crown Copyright data from the 1991 census, including digital boundary data, purchased for the academic community by ESRC and JISC. Thanks also to Michael Gill and Clive Sabel for help in data preparation.

References

Boots, B. N., & South, R. (1997). Modeling retail trade areas using higher-order, multiplicatively weighted Voronoi diagrams. *Journal of Retailing, 73*, 519–536.

Boyle, P. J., & Dunn, C. E. (1991a). *Redefining enumeration district centroids and boundaries* (North West Regional Research Laboratory Research Report 7). Lancaster University.

Boyle, P. J., & Dunn, C. E. (1991b). Redefinition of enumeration district centroids: a test of their accuracy using Voronoi polygons. *Environment and Planning A, 23*, 1111–1119.

Conduit, E., Brookes, R., Bramley, G., & Fletcher, C. L. (1996). The value of school locations. *British Educational Research Journal, 22*, 199–206.

Cox, K. R., & Agnew, J. A. (1974). Optimal and non-optimal territorial partitions: a possible approach towards conflict. *Papers of the Peace Science Society (International), 23*, 123–138.

Douglas, J. W. B. (1964). *The home and the school: the study of ability and attainment in the primary school*. London: MacGibbon and Kee.

Garner, C. L., & Raudenbush, S. W. (1991). Neighbourhood effects on educational attainment: a multilevel analysis. *Sociology of Education, 64*, 251–262.

Gibson, A., & Asthana, S. (1998). Schools and examination results: contextualising school 'performance'. *British Educational Research Journal, 24*(3), 269–282.

Haining, R. P., Griffith, D. A., & Bennett, R. (1984). A statistical approach to the problem of missing spatial data using a first order Markov model. *The Professional Geographer, 36*, 338–345.

McCallum, I. (1996). The chosen ones? *Education* 187 (3), 9 January 1996, 12–13.

McDonnell, R., & Kemp, K. (1995). *International GIS dictionary*. New York: Wiley.

Mortimore, J., & Blackstone, T. (1982). *Disadvantage and Education*. London: Heinemann.

Okabe, A., Boots, B., & Sugihara, K. (1992). *Spatial tessellations: concepts and applications of Voronoi diagrams*. Wiley.

Okabe, A., Boots, B., & Sugihara, K. (1994). Nearest neighbour operations with generalized Voronoi diagrams: a review. *International Journal of Geographical Information Systems, 8*(1), 43–71.

Tiefelsdorf, M., & Boots, B. (1997). Gambini Multiplicative Weighted Voronoi Diagrams [On-line]. Available at: http://www.wlu.ca/~wwwgeog/gambini.htm

Vincent, P., & Daly, R. (1990). Voronoi polygon types and their use in GIS. *Mapping Awareness, 4*(5), 40–42.

Wedge, P, & Prosser, N. (1973). *Born to fail?* Arrow Books.

Wrigley, N. (1976). *Introduction to the use of logit models in geography* (Concepts and Techniques in Modern Geography 10, GeoAbstracts, Norwich).