

# Investigating fixture propensity in a schools sports network using latent space models

Ian Hamilton, University of Warwick, Coventry, U.K.

February 2020

## 1 Abstract

Variants of a latent space model are fitted to a network of bilaterally determined school rugby fixtures with the aim of understanding the factors affecting the propensity for a fixture to exist. The primary importance of geographical proximity in the determination of these fixtures supplies a spatial ‘ground truth’ against which the performance of the models is assessed. Conclusions are drawn on the nature and size of the influence of known attributes, and on the suitability of applying latent space models for investigating communities in such contexts.

## 2 Introduction

Network latent space models position nodes in an unobserved latent space, with the probability of an edge existing between any two nodes being related to the proximity of the two nodes in the latent space. Modelling in this way can allow a number of features common to networks to be captured — transitivity; homophily by attributes; and clustering — as well as often allowing for informative graphical representations.

In the context of networks, transitivity is the phenomenon that two nodes, which each share an edge with the same third node, will have a higher probability of having a tie with each other than a pair that do not. Homophily by attributes describes the greater propensity for an edge to exist when two nodes share observed attributes. For example in a friendship network these could be attributes such as age, sex, geographical location, or recreational interests. Transitivity and homophily by attributes will both lead to clustering, but it is not uncommon to observe clustering beyond what may be explained by these features. This may be due to homophily by unobserved attributes, self-organisation of actors, the popularity of particular actors, or endogenous attributes such as position in the network (Handcock et al., 2007).

Early latent space network models used multidimensional scaling, and while these captured transitivity and homophily by attributes they were reliant on the arbitrary choice of a distance measure, leading to variable interpretations. Hoff et al. (2002) proposed a stochastic model in which the latent space positions may be estimated through standard statistical techniques. Handcock et al. (2007) extended that model to account for clustering by assuming that the latent positions are drawn from a finite mixture of multivariate normal distributions.

In this paper latent space models are applied to a network of schools rugby fixtures. This provides an interesting data set on which to apply these approaches, for a number of reasons. The fixtures are bilaterally determined rather than centrally scheduled, so without inspecting the network one may not discern the factors impacting the existence of a fixture. There would seem to be reason to expect that there may be clustering beyond that due to transitivity and homophily. First because there are likely to be factors not captured in available data that account for unobserved homophilies, such as sport orientation within school culture, and social networks between relevant staff. Second because the fixture information is publicly accessible, which might increase the potential for further clustering through self-organisation; if a school sees that a number of the schools they play have an opponent in common who they do not play they may consider proposing a fixture.

Importantly there are reasons to believe that the set of fixtures may be informative. Rather like the considerations of what will make a good Monte Carlo simulation, fixtures change from year to year but not drastically, and the process of fixtures being scheduled and changed has been taking place for many decades, with the first school rugby union fixture taking place over a hundred years ago. Together these considerations suggest that there has been time for relevant factors to exert an influence such that the observed situation represents a steady state with respect to the schools' current characteristics, allowing inferences to be informative. From the point of view of assessing the effectiveness of the latent space models, since one may reasonably expect that a primary consideration in agreeing a fixture would be geographical proximity, then there is a 'ground truth' homophily against which to assess them. It seems plausible that other homophilies, such as playing strength or school type expressed in various ways, may also have an influence on the probability of a fixture between two schools; some of these are considered, along with an assessment of the relative importance of these effects.

The paper proceeds in section 3 by discussing the data, with details of its collection, and a brief analysis highlighting some features and inter-relations of the school covariates considered. In section 4 the latent space models of Hoff et al. (2002) and Handcock et al. (2007) are fitted. In section 5 the influence of the different covariates is explored. Section 6 provides some concluding remarks.

## 3 Data

### 3.1 Background

The network to be investigated is based on all matches in the Daily Mail Trophy in the 2015–16, 2016–17 and 2017–18 seasons. The Daily Mail Trophy is an annual tournament between some of the best school rugby teams in England, and a single school in Wales. In order to qualify for a ranking in the Daily Mail Trophy a school must register and compete against a minimum of five other teams in the tournament. Schools who registered for the tournament but were unable to complete five eligible matches in any given season are maintained in our analysis, despite being excluded from a ranking in the tournament. There are in total 118 schools included, playing between 3 and 37 matches over the course of the three seasons. The matches played as part of the tournament are typically only a subset of the matches played by these schools, with other matches taking place with non-tournament schools, or as part of a centrally scheduled knock-out competition. In almost all cases however, if they have played a match against one of the other tournament teams, outside of the centrally scheduled knock-out competition, it would be classed as a Daily Mail Trophy match and would appear in the data set. So the existence or absence of a match with another school is an accurate representation of the bilaterally arranged fixtures within the set of schools in the tournament. The data set omits matches played against opponents who do not appear in the tournament. The network being considered is undirected and weighted. Each node is a school, and each edge has a value equal to the number of matches between the schools over the three seasons. This network is considered in relation to school-level data representing variables that could contribute to some observable homophily.

### 3.2 Data collection

The fixtures and results for the Daily Mail Trophy were kindly shared by SOCS, the organiser of the tournament, but are also available at the tournament website ([www.schoolsrugby.co.uk/dailymailtrophy.aspx](http://www.schoolsrugby.co.uk/dailymailtrophy.aspx)).

For each school, data were also collected on:

1. annual fees (Fees)
2. year of foundation (Founded)
3. the number of boys in sixth form (6th Form boys)
4. the proportion of pupils in the school that are boys (Percent boys)
5. the proportion of pupils that are boarders (Percent boarders)
6. whether the school is privately or state funded (School type)

7. whether the school played one or two terms of rugby (Term type)
8. performance rating of the rugby-playing strength of the school (Rating)

While it would be plausible to consider other variables, these were either not readily available, for example the size of sports bursaries or proportion of pupils from overseas, or were substantially accounted for by this set of covariates, for example if a school is co-educational throughout or just in sixth form, or not at all, is well captured in the proportion of pupils that are boys and the number of sixth form boys. All covariate data were collected during the first week of June 2019. This means they are not contemporaneous with the fixtures occurring in the period 2015–2018. This is not expected to have materially impacted any conclusions because the covariates are not subject to large year on year changes, .

Fees information was sourced from the schools’ own websites. Annual fees were taken to be the minimum fees for full-time education of a pupil in Upper Sixth form, not accounting for bursaries or scholarships. Thus it is zero for state-funded schools, standard day fees for schools that admit day pupils, and minimum boarding fees for schools that admit only boarders.

Year of foundation was generally more readily available on the Wikipedia page for the school than the school’s website and so the Wikipedia date was used. This was originally included as it was suspected that it might define a meaningful similarity between schools. But in collecting the data it became clear that the trajectories of schools were very diverse and it was often even difficult to be sure of a definitive foundation date with occurrences of, for example, schools moving geographically, amalgamating, moving from private to state or vice versa, and renaming not uncommon. All of these contributed to subjectivity in the definition of year of foundation. It is included in the analysis however as it provides a useful sanity check to some of the later methods in the degree to which those methods identify year of foundation as a non-informative covariate.

The data on the number of pupils, including the proportions of boys and boarders, and the absolute numbers of sixth form boys, were generally not available from schools’ websites. For privately funded schools, these data were collected from the Independent Schools Council (ISC) website ([www.isc.co.uk](http://www.isc.co.uk)). For boys and for girls, this reports the number of boarders, the number of day pupils, and the number of sixth form pupils. From these the proportion of day pupils, and the proportion of boys are calculated. However it should be noted that since schools admit pupils from different ages, the relative proportions as they pertain to a consistent age group, say 13–18, are unlikely to be the same. Given the purpose of including these proportions — identifying homophilies — and that no ready alternative was available, then this seems acceptable, as they still provide evidence on the nature of the school. Pupil numbers for state schools were taken based on the most recent relevant government report (Department for Education, 2019). These did not include a delineation for numbers of sixth form students or of boarders. For this, the most recent available Of-

sted report with such data was used ([www.gov.uk/government/organisations/ofsted](http://www.gov.uk/government/organisations/ofsted)). Since these were from previous years and so total numbers differed, the proportions of sixth form students or boarders were assumed to be constant, and the absolute number of sixth form boys was adjusted for the current total number of pupils. The total number of pupils was never materially different from current numbers, and so one may reasonably be confident that these numbers are accurate. However it should be noted that these Ofsted reports were quite commonly from as much as a decade ago. In one case, Ampleforth College, the population information was not available on the ISC website and so the method used was the same as for the state schools, with the most recent Independent Schools Inspectorate report ([www.isi.net](http://www.isi.net)) used, in place of an Ofsted report, to derive the relevant proportions. Whether a school was state or privately funded was determined by its classification in the most recent relevant government report (Department for Education, 2019).

In order to determine if the school played one or two terms of rugby, the month of the school's first and last fixtures of the 2018 season were considered, including those fixtures outside the Daily Mail Trophy competition. If these both lay in a single term, then it was interpreted as being a single term rugby-playing school. If there were two or more matches in the second term then it was deemed to be a two term school. If there was a single match outside of first term then previous seasons' fixtures as well as any information from the school's website was used to make the categorisation.

The performance rating was based on the method discussed by Hamilton & Firth 2020 [in preparation], applied to an aggregation of the three seasons' results. One may use this method to calculate a projected league points per match were each team to play every other team home and away in a round robin format. The method coherently accounts for the varying schedule strengths of participating schools, and provides a positive-valued measure of a comparable scale to the other covariates as well as an intuitive interpretation.

For all of these factors, when considering pairwise covariates the absolute difference is used. The binary variables of school type and term type take value 0 if identical and 1 if different for each pair. The two percentage variables of proportion of boys and proportion of boarders are taken as the absolute difference in these percentages.

Postcodes were taken from the Daily Mail Trophy website and then used via a Google Maps API in order to calculate travel times and distances between schools using the R package `googleway` (Cooley, 2017) and to project locations onto a relevant map. Distances were calculated as at Saturday 5th October 2019 12pm using the "best guess" methodology, assuming a journey by road. In order to plot schools geographically, latitude and longitude for each postcode was sourced from [www.freemaptools.com](http://www.freemaptools.com).

The five continuous variables are scaled to allow for better comparability and interpretability, specifically the following units are used:

Travel Time	hours
Fees	£10,000
Founded	centuries
6th Form Boys	100 boys
Rating	projected league points per match

### 3.3 Exploratory Data Analysis

The two binary covariates of whether a school is privately or state funded, and whether it plays one or two terms of rugby are compared first. Table 1 shows that there is a clear relation, with the private schools substantially more likely to play one term and state funded schools more likely to play two.

	One term	Two terms
Private	79	17
State	8	14

Table 1: Number of schools in the tournament playing one or two terms of rugby and private or state funded

The other covariates are presented in Figure 1. A number of things may be noted. First the individual covariate distributions are largely in line with expectations. A number of the covariates are bimodal. This is unsurprising in the case of fees, with one of the modes at zero, or in the case of the percentage of boys, with modes close to 50% and 100%, reflecting a predominance of all-through co-educational and single sex schools, but the bimodal nature to the year of foundation is less intuitive without more historical context. The number of sixth form boys is right-skewed with a mode around 150. The proportion of boarders has a clear mode at just above zero, reflecting a high proportion of day schools and some schools with just a handful of boarders. The remaining proportions are, perhaps not so intuitively, distributed quite evenly up to full boarding status, but without another clearly identifiable mode. Rating is left-skewed with its mode at around 2.7 league points per match, which suggests teams share a little over one bonus point per match. Importantly for our analysis, the continuous covariates seem largely to represent somewhat independent features of the schools. Interestingly there does seem to be a correlation between the proportion of pupils that are boarders and the fees. This is not due, as one might first suspect, to boarding fees being higher as they must also account for living expenses, since the minimum upper sixth form fees have been used in all cases, and there are only four pure boarding schools where those minimum fees are boarding fees. In all other cases they represent a day fee. A correlation of 68% between the proportion of pupils that are boys and the number of sixth form boys is less surprising. Here colours have been used in the scatterplots to

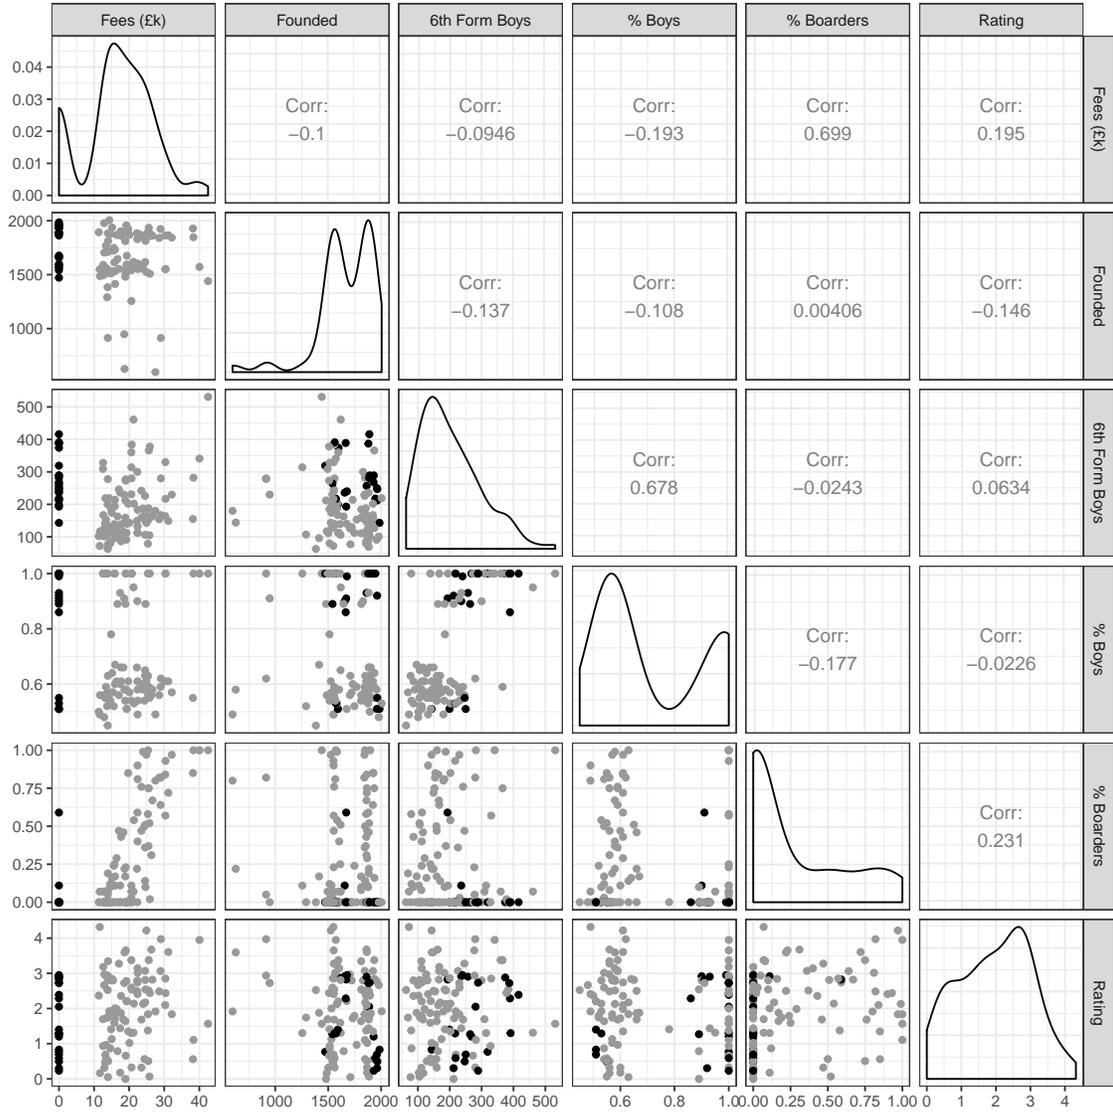


Figure 1: Scatterplots of school covariates. Private schools in grey, state funded schools in black

differentiate state funded and private schools. Apart from the self-explanatory difference in fees, these also show that within this set of schools the state schools are more likely to have high numbers of sixth form boys, to be single sex (or close to), and to be day schools rather than boarding. We might therefore expect some confounding of these factors in later analyses. These charts were also looked at with a differentiation based on the number of terms of rugby played. This highlighted similar features, as one might expect from Table 1, but less strongly.

## 4 Latent Space Model

### 4.1 Model specification

Given the three seasons of fixture data, a binomial latent space model is fitted here, the general form of which is

$$P(A \mid \mathbf{Z}, \mathbf{x}, \boldsymbol{\beta}) = \prod_{i < j} \binom{3}{A_{ij}} \mu_{ij}^{A_{ij}} (1 - \mu_{ij})^{3 - A_{ij}} \quad (1)$$

$$\text{logit}(\mu_{ij}) = \beta_0 + \sum_{k=1}^p x_{ijk} \beta_k - d(\mathbf{Z}_i, \mathbf{Z}_j). \quad (2)$$

Here  $A = [A_{ij}]$  is the symmetric adjacency matrix of fixtures with  $A_{ij}$  equal to the number of matches between teams  $i$  and  $j$  over the three seasons,  $x_{ijk}$  is the  $k$ th pairwise covariate for teams  $i$  and  $j$ ,  $\beta_k$  is the coefficient for the  $k$ th covariate,  $\mathbf{Z}_i$  is the latent position for team  $i$ , and  $d(\mathbf{Z}_i, \mathbf{Z}_j)$  is the Euclidean distance between teams  $i$  and  $j$  in the latent space. This is the binomial version of the model proposed by Hoff et al. (2002).

It is worth noting that a sociality parameter is not included. In this context a sociality parameter would describe the propensity for a particular school to have matches. It is not included here because the fixtures represented are an incomplete set of fixtures for the participating teams, with teams generally playing additional matches, against teams outside of the tournament, as friendlies or as part of other competitions. While this data is not fully available, the schedule for a sample of teams has been inspected and generally they have played a similar total number of matches, so observing a team to have higher degree within the network is not a reflection of that team having a higher propensity to play matches in general. Including such effects could therefore, for example, misleadingly diminish the extent to which one might infer a lower homophily from the absence of a fixture in the case of teams with lower degree. On the other hand, the current rules of the tournament encourage teams to play as many other tournament participants as possible, since they are awarded bonus points merely for playing matches in the tournament, independent of result. Based on conversations with the tournament organisers and observations of historic results, it is likely that there will be a difference in motivation between teams in

their desire to be competitive in the tournament, and some may actively seek to schedule more tournament matches. However there are only a small number of teams, consistent across years, for whom the tournament is a goal in itself. For the vast majority, they enter simply because they can, as a by-product of their standard fixture list, which is largely similar from year to year. So while this effect could be argued to be a genuine sociality effect, within the context of just this network of fixtures, it is likely to apply to only a small minority of teams. Therefore on balance, the distorting impact of inclusion is considered to be more of a danger than that of exclusion and so no sociality effect is included. All models are fitted using the latentnet package in R (Krivitsky and Handcock, 2008).

## 4.2 Hierarchical Clustering

Initially the parameters of the model represented in equations (1)–(2) are estimated through the method of maximum likelihood.

It was suspected that geographical proximity would be a primary driver of the propensity for a match to occur, and therefore of model distance. Figure 2 plots the pairwise Euclidean distances, calculated using the latent space positions, against the pairwise estimated travel times. This shows a strong relationship, with a Pearson correlation of 74%. Estimated travel time is used here as this would seem to be a more relevant motivating condition for a fixture than actual geographical distance, but substantially similar results are found using geographical distance. The comparison between latent space distance and travel time may also be used as a means of testing the choice to use a two dimensional latent space. When fitting with three dimensions the correlation increased only by 1%, strongly suggesting that modelling in two dimensions, with the representational benefits it brings, is a reasonable choice. As such, all further models will be applied using a two dimensional latent space.

While the fit is reasonable, there seems to be a notable skew to the residuals. Looking at the residual plot in Figure 3 it can be seen that this is driven by two groups. The first has high travel time and considerably lower latent space distance than the linear regression would suggest, the second low travel time but considerably higher than expected latent space distance.

It might be supposed that the former group is likely to be due to the requirement for geographically extreme teams to travel longer distances in order to complete a sufficient number of matches in the tournament, so that travel time for them has a different level of consideration than for teams with greater geographical proximity to other schools. However Figure 4 suggests this does not account for the entire effect. While the most northerly and southerly teams, as well as the single team in East Anglia do stand out in the figure on the left, teams from all over the country are represented and the third highest weighted is King Edward’s, Birmingham in the middle of the country, suggesting there is something else at work here. The other group is made up of geographically proximate teams with a

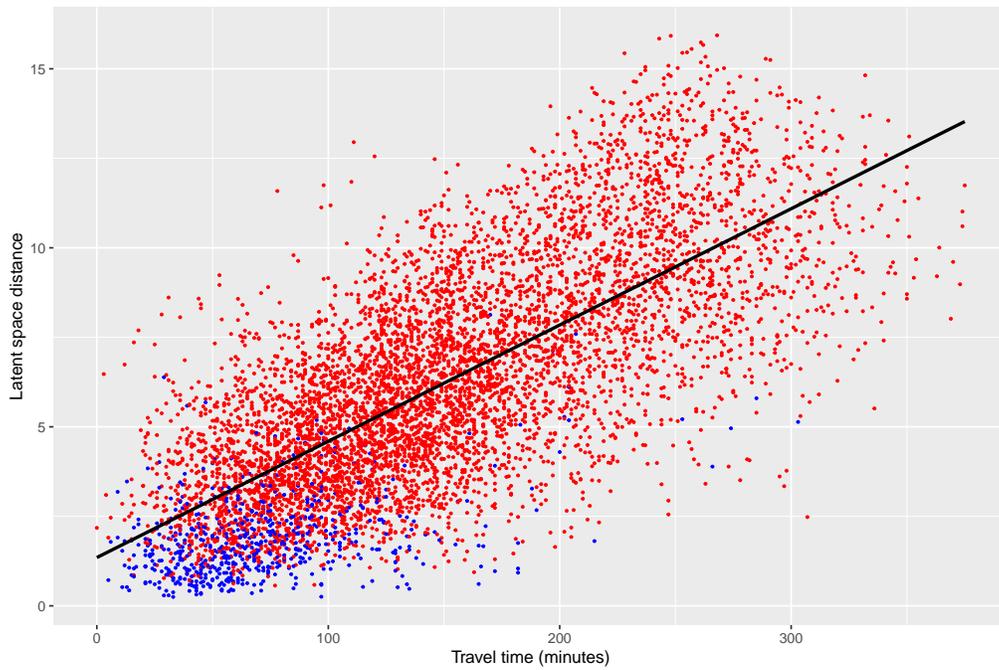


Figure 2: Scatterplot of latent space distance against travel time. School pairs who do not play each other during the three seasons are in red, pairs who play each other at least once are in blue. Linear regression line is shown in black.

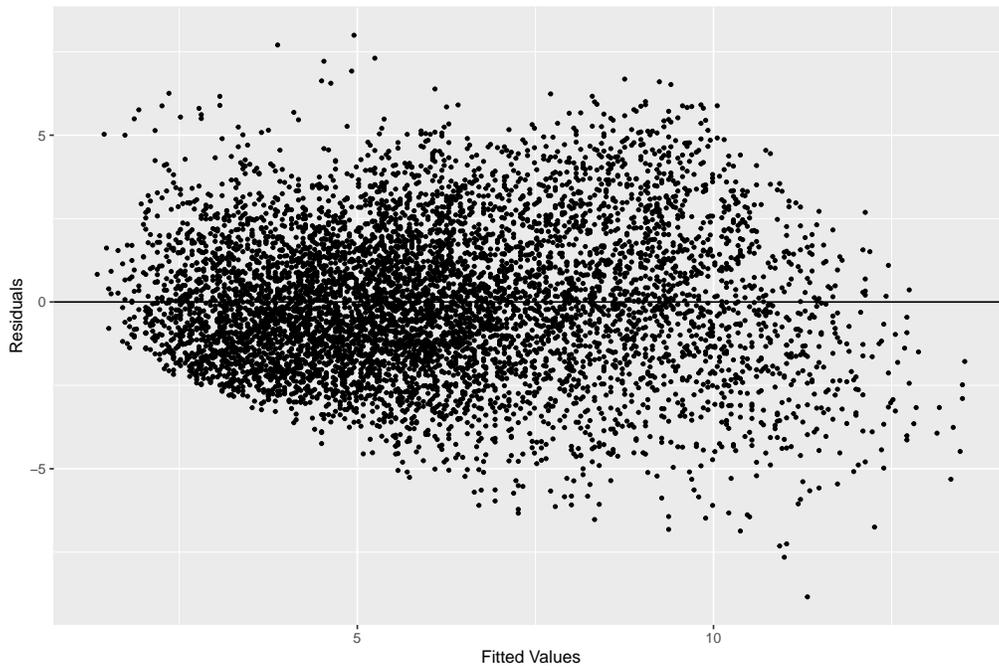


Figure 3: Scatterplot of residuals from fitting latent space distance against travel time



Figure 4: Location maps for the teams featuring in the extreme residuals. Size of dot represents the number of times that the team appears in the relevant set of pairs. Left hand chart includes pairs where  $d < \hat{d} - 4$ , and right hand where  $d > \hat{d} + 4$ , where  $d$  is the latent space distance and  $\hat{d}$  the expected latent space distance based on the linear regression with travel time.

large latent space distance. Again it is perhaps to be expected that many of these are in the more densely geographically clustered south east, but some of the most westerly teams also feature strongly. Even in the case of the south east teams, it remains unexplained as to why these particular teams have this greater latent space distance. This will be examined further in Section 5 when the influence of the other pairwise covariates is investigated.

The geographical implications of model distance may also be considered by investigating community detection in relation to the geographical location of the schools. Numerous methods of clustering could be applied given the latent space distance matrix. As an example, by applying a hierarchical clustering with complete linkage the dendrogram presented in Figure 5 is obtained.

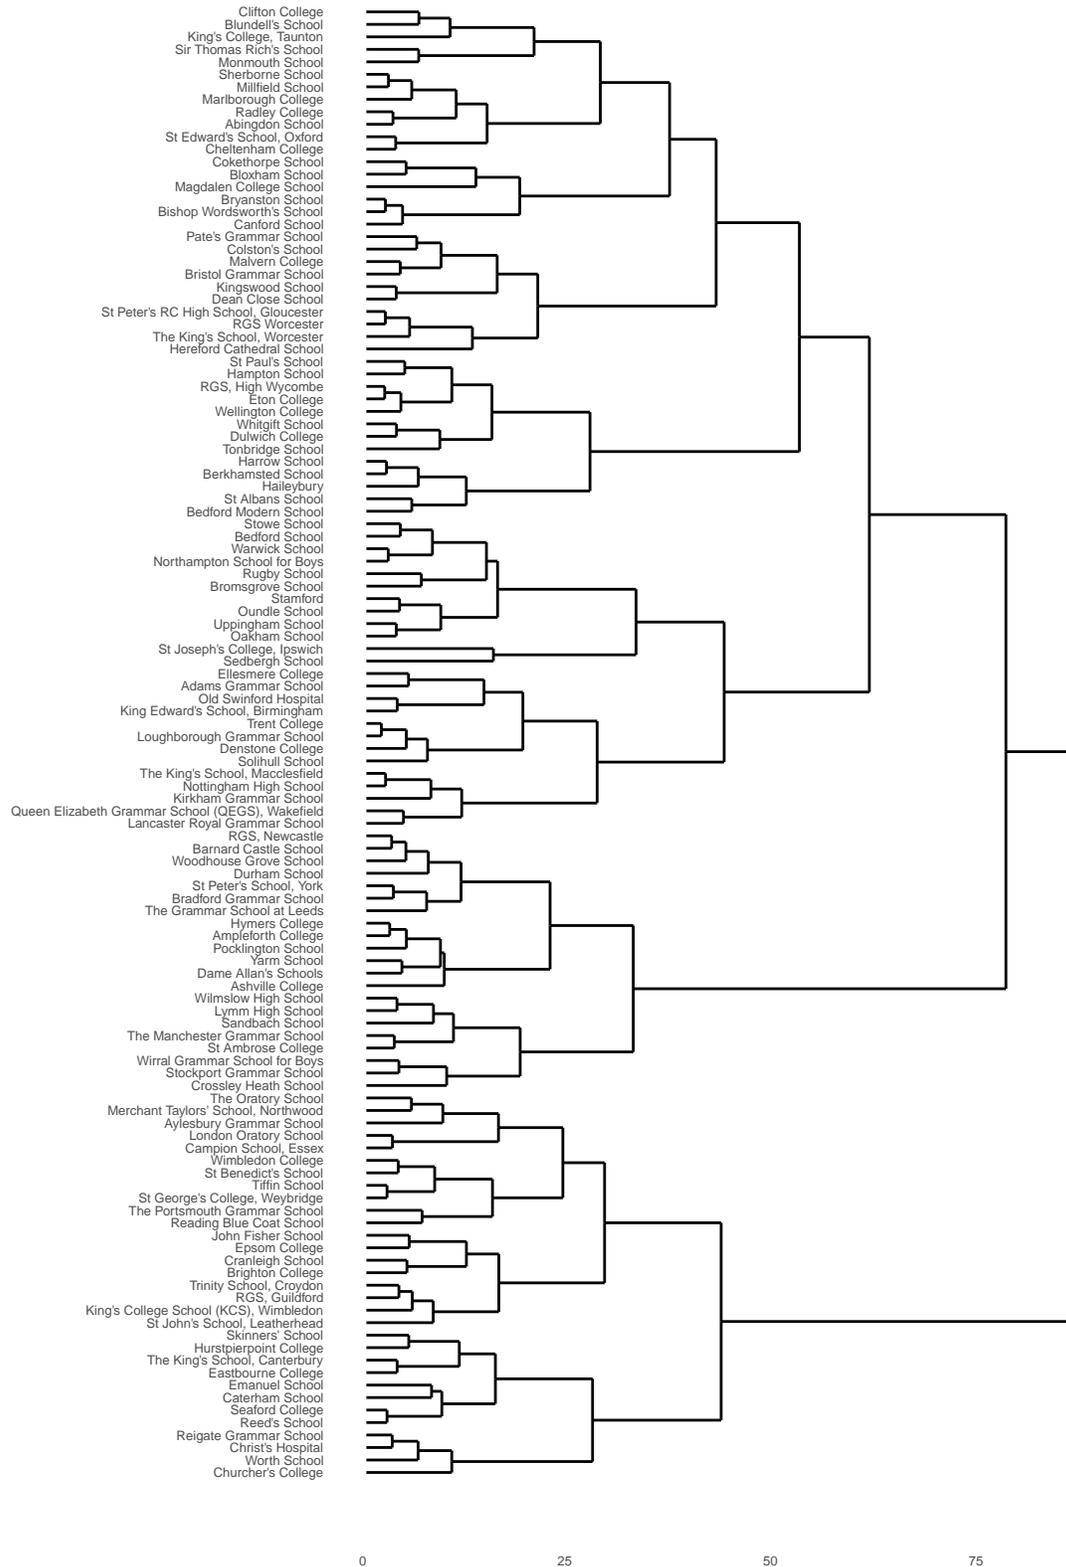


Figure 5: Dendrogram of hierarchical clustering of schools by latent space

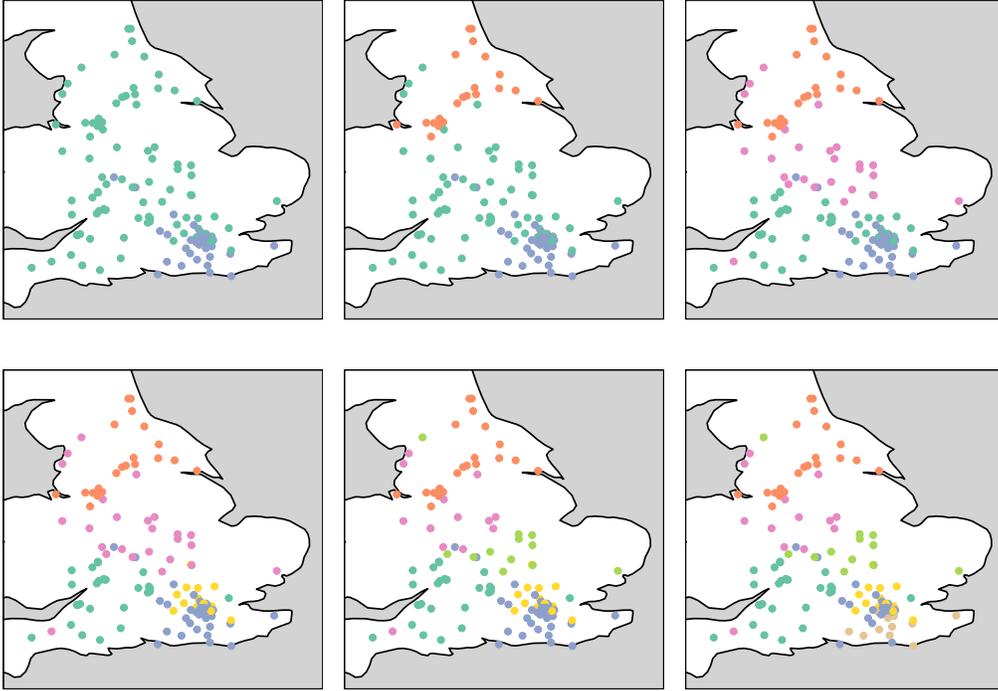


Figure 6: Communities of size  $G=2,3,4,5,6,7$  based on complete linkage hierarchical clustering of latent space distances

Inspection of the dendrogram suggests that five communities could be an appropriate partition with a number of communities converging at a distance of around 45. In Figure 6 communities are plotted for  $G = 2, 3, 4, 5, 6, 7$ , where  $G$  is the number of groups, in order to show the geographical detection ability with different numbers of groups.

### 4.3 Latent Position Cluster Model

Handcock et al. (2007) provided a model with the extended feature that the latent positions are drawn from a finite mixture of multivariate normal distributions. That is,

$$\mathbf{z}_i \stackrel{i.i.d.}{\sim} \sum_{g=1}^G \lambda_g \text{MVN}_d(\mu_g, \sigma_g^2 I_d) \quad .$$

Two methods of estimation were proposed in Handcock et al. (2007). The first method is a two-stage maximum likelihood procedure, and the second a fully Bayesian estimation using Markov chain Monte Carlo sampling. In this paper, this is fitted based on a burn-in

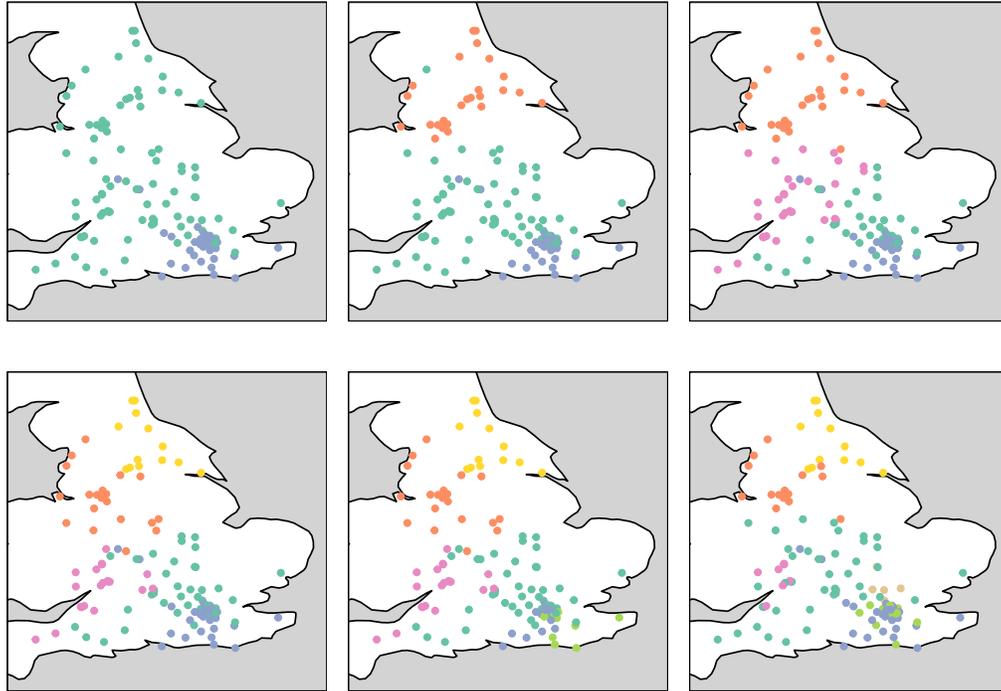


Figure 7: Communities of size  $n = 2, 3, 4, 5, 6, 7$  based on two stage MLE

period of 10,000 iterations and a sample run of 1,000,000 iterations of which every fiftieth was sampled, giving a sample size of 20,000.

Figures 7 and 8 present the results from the two methods. In Figures 6, 7, and 8, communities are identified such that the number of schools remaining in the same community as in the previous clustering (as represented by a particular colour) is maximised. The different colours representing different regions in the three Figures is thus a result of different evolutions of the community detection in each case. The fittings using the Gaussian clustering, both based on the two step MLE and the MCMC, appear to show better geographic separation than did the hierarchical clustering. The MCMC clustering arguably shows a better separation in the London area, though there is substantial agreement in the community membership up to  $G = 5$ .

For the purposes of further investigation it is useful to select a single number of communities,  $G$ , with which to work. Graphical inspection of Figures 7 and 8 suggests that between three and five communities may fit best, given the clear geographical separation they evidence. Handcock et al. (2007) suggest the use of a Bayesian Information Criterion (BIC) for the purpose of selecting the number of communities when using the MCMC fitting.

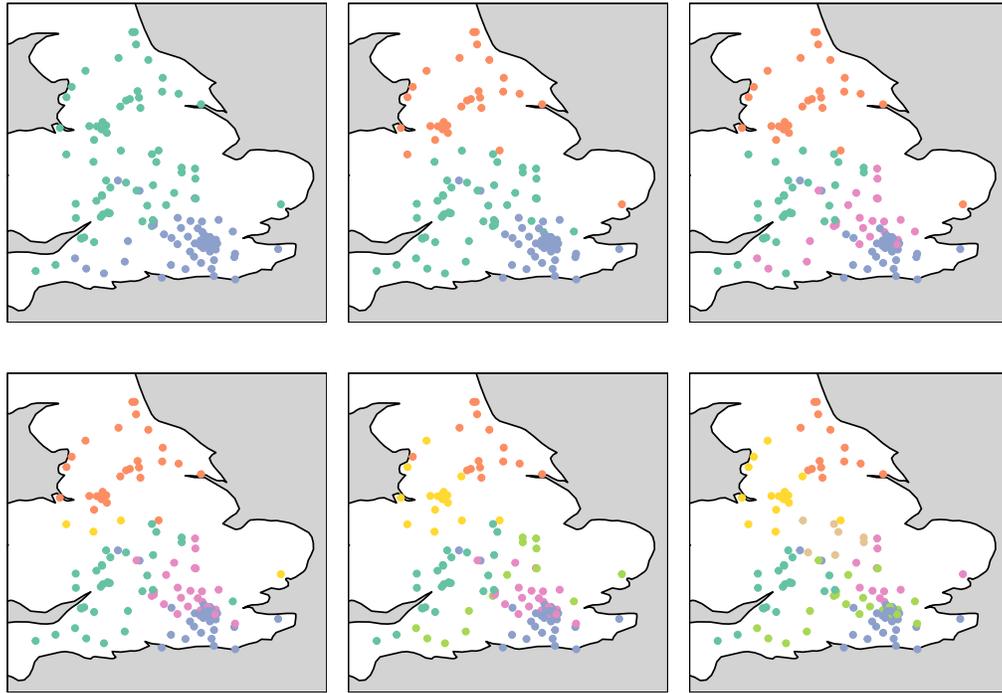


Figure 8: Communities of size  $n=2,3,4,5,6,7$  based on MCMC parameter estimation

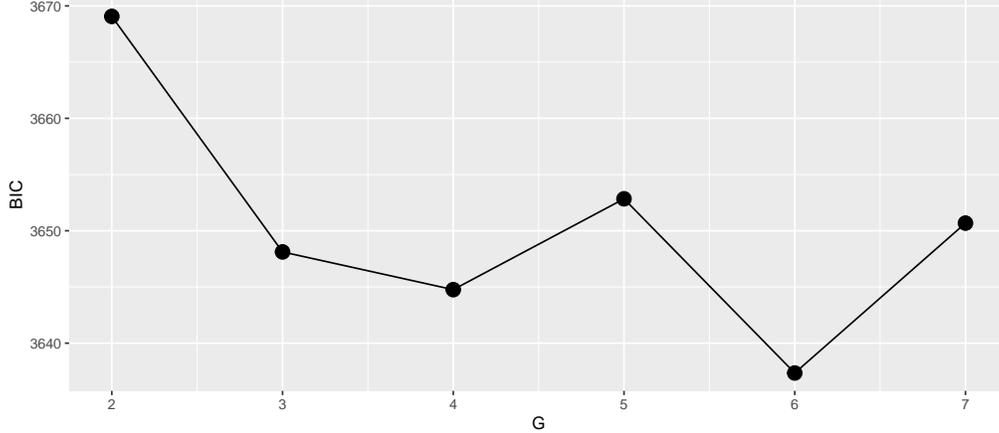


Figure 9: Bayesian Information Criterion for different values of  $G$

Figure 9 shows this BIC for each value of  $G$ . Experimentation with different specifications of the algorithm showed the BIC value to be somewhat unstable. As such this is taken to be indicative rather than definitive and so taking into account both the BIC and the geographical separation  $G = 4$  is chosen.

In Figure 10 the travel time in minutes is plotted against the latent space distance. This is shown side by side for the two stage MLE and the MCMC estimation to aid comparison. The most notable feature is that the effect of applying the simultaneous Bayesian estimator obtained from the MCMC fit is to reduce and also compress latent space distances particularly for the large latent space distances. The interquartile range of the residuals, for example, is thus more than 25% lower when compared to the MLE fit. The Pearson correlation is slightly higher at 78% compared to 74% from the MLE fit, and the intercept is closer to zero. This suggests that the assumption of community membership usefully constrains the fit. The community memberships are substantially similar with 108 of the 118 schools belonging to equivalent groups. However based on the seemingly improved latent space fit, the MCMC estimation with  $G = 4$  is the model selected for further analysis.

A feature of the mixture model is that it allows us to view the probability of each community membership for each node. Figure 11 provides a graphical interpretation of these probabilistic community memberships. Perhaps the most notable feature is the difference in delineation of the communities. What might be referred to as the northern community is most clearly delineated, with all but two members having greater than 75% probability of being part of that community. On the other hand the two central communities are substantially more indeterminate with the southernmost (that represented in the chart in pink) having no member with a greater than 75% probability of being part of the community. The schools in and around London show perhaps more delineation than one might expect

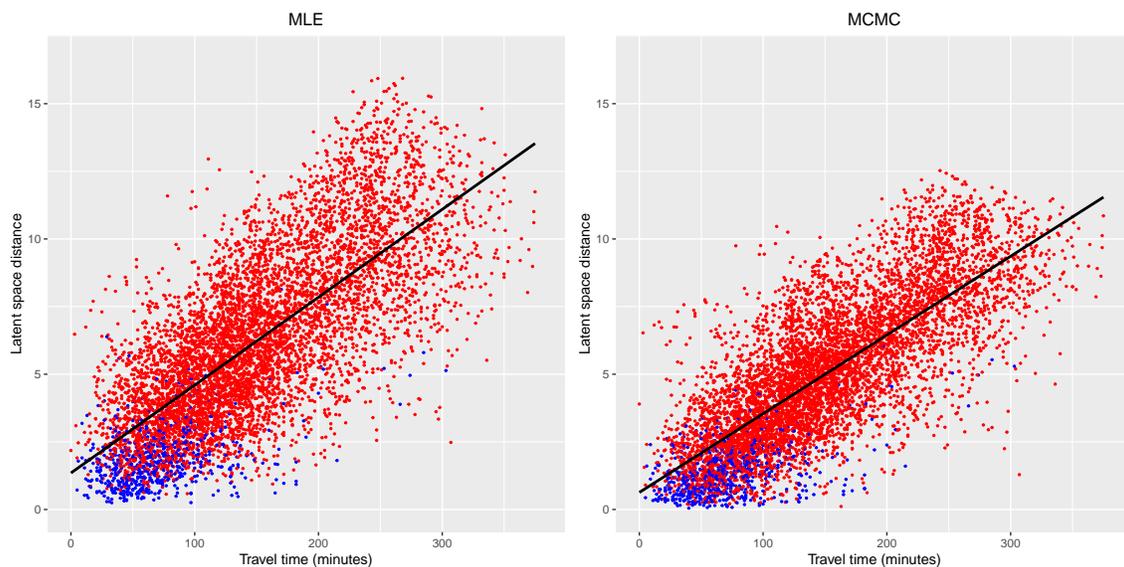


Figure 10: Scatterplot of latent space distance against travel time for two stage MLE and MCMC fits with  $G = 4$ . School pairs who do not play each other during the three seasons are in red, pairs who play each other at least once are in blue. Linear regression line is shown in black.



Figure 11: Probability of membership of each community. Left figure shows all teams. Right figure shows those in and around London

given the short distances. Due to the geographical proximity of a number of schools, some are obscured in the chart. This does not alter the interpretation when the distributions of these schools are known. It could also be noted that the most geographically isolated school, St Joseph’s Ipswich, has a substantial probability of being in three of the four different groups.

An alternative way of considering the link between the latent and geographical spaces is presented in Figure 12. This suggests that the two latent space dimensions are picking up something like a polar coordinate system with a radial coordinate emanating from a centre in the south-east (as seen on the size scale), and an angular coordinate going from the south to the north east (as seen in the colour scale), an observation which is consistent with the probabilistic clustering, as well as the evolution of the community detection seen in Figures 7 and 8.

## 5 A consideration of covariates

### 5.1 Relative Importance

The analysis so far has shown a clear link between the travel time and the propensity for the existence of fixtures, but it is desirable to consider this more directly, and how it compares to the other covariates of interest. Grömping (2015) provides a thorough overview of various

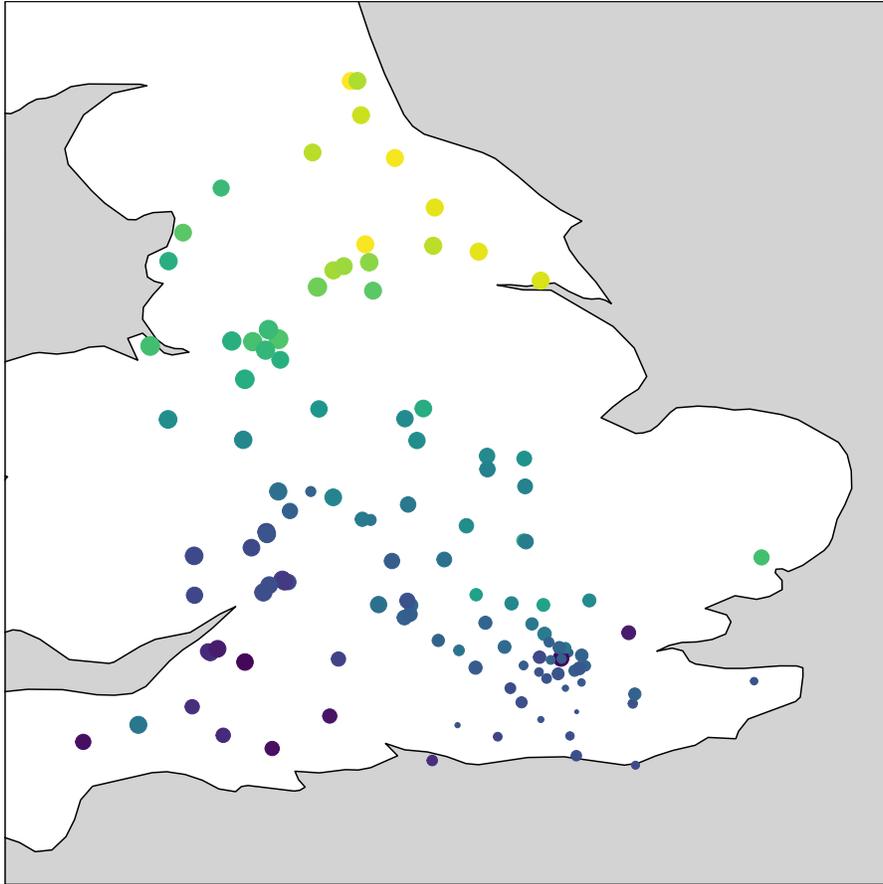


Figure 12: Map of schools with colour and size representing first and second latent space coordinate respectively for each school.

measures of relative importance, and the factors to consider in selecting one. These are presented in the context of linear regression. Some of these measures may be extended to wider families of models, and in particular, Thomas et al. (2008) have proposed relative importance measures for logistic regressions consistent with some of the features of the measure axiomatically justified by Pratt (1987) for linear regression.

In the present case, the unreliability and the computational expense of the coefficient estimates obtained when the model is fitted with all covariates simultaneously are considerable hindrances to applying relative importance measures directly to the model. This is especially so since a number of the measures require at least  $p!$  computations, where  $p$  is the number of regressor variables, in order to average over the permutations of regressor variables. One alternative is to instead use the pairwise latent space distances estimated by the selected model. A linear regression of these distances against the pairwise covariates can then be considered. Doing so provides an estimate for the relative importance of the covariates in the determination of pairwise latent space distance. These therefore may be viewed as a measure of relative importance of the covariates in the propensity for a fixture to exist only in so much as latent space distance is a good measure of the propensity for a fixture to exist.

Two measures are considered here. Both have the desirable and intuitive properties of providing: independence from the order of the regressors in the model; scale invariance; a decomposition of the model variance ( $R^2$ ); and a proper decomposition of the model variance for any orthogonal regressor subgroups. The first (which we will refer to as ‘Pratt’) is the one advocated by Pratt (1987), which provides an axiomatic justification for using  $b_k \rho_k$  to assess relative importance, where  $b_k$  is the standardised coefficient with respect to the  $k$ th regressor variable  $x_k$ , and  $\rho_k$  is the correlation of the  $k$ th regressor variable  $x_k$  with the independent variable  $y$ , in this case the latent space distance. The principal objection to this approach is that it can produce a negative measure, which is not clearly interpretable. The second (LMG, after the original authors) is the method originally due to Lindeman et al. (1980), which was also influentially and independently advocated by Kruskal (1987). For each regressor, this takes the marginal increase of explained variance from the addition of the regressor to the model, and takes the mean of these values over all regressor order permutations of the model. Taking the notation of Grömping (2015), denote the explained variance and the sequential additional variance respectively as

$$\begin{aligned} \text{evar}(S) &= \text{var}(y) - \text{var}(y \mid x_j; j \in S), \\ \text{svar}(M) &= \text{evar}(M \cup S) - \text{evar}(S). \end{aligned}$$

Then we can define, without loss of generality, the measure for the first regressor as

$$\text{LMG}(1) = \frac{1}{p!} \sum_{\pi} \text{svar}(\{1\} \mid S_1(\pi)),$$

where  $\pi$  are the regressor permutations and  $S_1(\pi)$  the set of regressors preceding regressor 1 in permutation  $\pi$ . The model has an  $R^2$  of 62% with the contribution of each covariate shown in Table 2.

Variable	Pratt	LMG
Travel time	98.09%	97.77%
Percent Boarder	0.88%	0.88%
Fees	0.72%	0.63%
Percent Boys	0.09%	0.08%
Boys	0.08%	0.12%
School type	0.05%	0.30%
Rating	0.04%	0.14%
Term type	0.02%	0.02%
Founded	0.01%	0.04%

Table 2: Proportion of explained variance ( $R^2$ ) attributable to each covariate in linear regression of pairwise latent space distance against covariates. Ordered in descending order of relative importance under the Pratt measure.

It is perhaps unsurprising based on the previous analysis that travel time is dominant, but the degree to which this is the case is nevertheless notable. The other covariates have negligible relative importance in comparison to travel time, but relatively they suggest that fees and the proportion of boarders may have greater influence. The two measures are very substantially consistent, with Pratt giving slightly higher values for travel time and fees and lower for school type and rating.

## 5.2 Covariate inclusion

An alternative means to investigate the influence of the different covariates is by considering models with each covariate included individually,

$$\eta_{ij}(k) = \beta_0 + x_{ijk}\beta_k - d(\mathbf{Z}_i, \mathbf{Z}_j), \quad (k = 1, \dots, p),$$

where the  $k$  represent the dyadic covariates. Table 3 presents the parameter estimates based on the posterior mean of the coefficient, as well as a statistic representing  $q = 2 \times \min(P(\beta_k > 0), P(\beta_k < 0))$ , where  $\beta_k$  is the coefficient under consideration, and the Bayesian Information Criterion for each model, which may be compared with a BIC of 3647 for the model with no covariates included.

The directions of the coefficient estimates are as one would expect, so that a smaller covariate distance implies a higher probability of a fixture for all covariates. The BIC estimates

are consistent with what was seen in the relative importance analysis in highlighting travel time as the dominant covariate. The estimate of the coefficient implies that the odds of a fixture decrease by 78% for every additional hour of travel between schools. In that fees and the proportion of boarders are the only others that have a notably lower BIC and also  $q$ -value, these are also somewhat consistent with the previous results, with the odds of a match decreasing by 31% for every £10,000 difference in fees, and 58% for a fully boarding school (100% boarders) playing a fully day school (0% boarders) as compared to a match between two with the same proportion of boarders. With the exception of rating, the  $q$ -values are all low, indicating that it is likely that these factors are influencing the propensity for a fixture to exist. This is perhaps surprising in the case of year of foundation, though the  $q$ -value there is less conclusive.

Another method of investigation that the model allows is by inspection of the latent space graphically. Given the dominance of travel time, it is useful to control for it in considering the other covariates. In Figure 13 the model with

$$\eta_{ij} = \beta_0 + x_{ij\text{Travel Time}}\beta_{\text{Travel Time}} - d(\mathbf{Z}_i, \mathbf{Z}_j)$$

is used with  $G=3$  based on BIC values. The colours are used to show the fees and proportion of boarders for each school. As was noted in Section 3 these covariates are strongly linked and this is evident in the plots. It can be seen that while schools with lower fees and low proportions of boarders are distributed quite evenly throughout the latent space, those with higher fees and a high proportion of boarders are to be found disproportionately often in a cluster in the lower right of the latent space. This effect is perhaps clearer for the proportion of boarders. This suggests that there might be a small community of some of the more expensive, primarily boarding, schools who are more likely to play each other.

	$\hat{\beta}_k$	$q$	BIC
Travel Time	-1.523	$< 10^{-15}$	3316
Fees	-0.368	$< 10^{-15}$	3618
Percent boarder	-0.860	$< 10^{-15}$	3630
Term type	-0.309	0.0016	3642
6th Form boys	-0.199	0.0039	3642
School type	-0.370	0.0017	3643
Founded	-0.038	0.0473	3646
Percent boys	-0.756	0.0070	3649
Rating	-0.057	0.3239	3651

Table 3: Coefficient,  $q$ -value and BIC when model fitted with individual additional covariates. Ordered in increasing BIC.

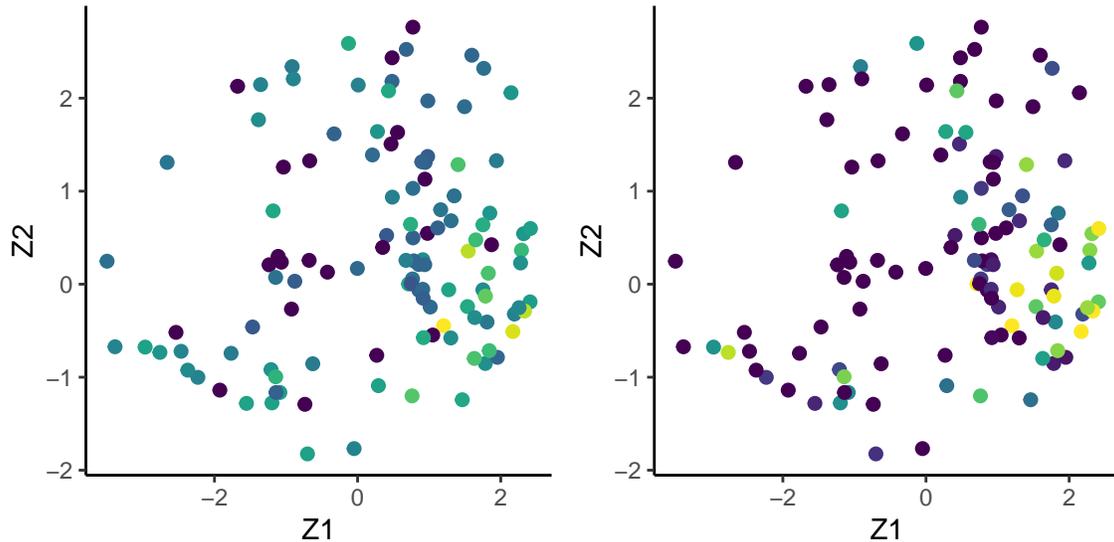


Figure 13: Latent space plots based on model with Travel Time controlled for and with  $G=4$ . Colour scale representing Fees on left hand side, and % Boarders on right hand side

## 6 Concluding Remarks

Perhaps unsurprisingly travel time has been shown to be the dominant factor in contributing to the prevalence of fixtures. The degree of this dominance is notable however, representing 98% of  $R^2$  based on the two relative importance measures employed. The nature of the travel time dependency seemed to be of a community based in the south east with other communities tiered based on latitude. In addition to the travel time dependency, fees and the proportion of boarders were identified as the next most significant factors, with the dependency driven by a community of expensive schools with a high proportion of boarders that showed an increased tendency to play each other. It should be noted however that the data represent a self-selecting group of the top rugby-playing schools. It seems not unreasonable to expect that were a full set of data available then the impact of some of the other covariates would take on a greater significance. For example over a wider data set with more state schools and a greater range of abilities, perhaps fees, school type or rating would take on greater relative importance.

The latent space models demonstrated a number of desirable features in the context of this question. They showed a good ability to detect geographical communities in the data set, and the spatial nature of the output provided a natural means of interpreting the result in terms of pairwise distances. It was possible to employ a variety of mechanisms in order to diagnose the relative importance of the dyadic covariates in the propensity for a fixture to

exist, producing a consistent interpretation in the methods employed. The latent position cluster model of Handcock et al. (2007) usefully constrained the latent space such that it appeared to provide a better fit. It also usefully facilitated a means of assessing the strength of community attribution at a school level by producing the probabilities of each school being a member of each community as seen in Figure 11. It is notable that, while a relatively high number of iterations was chosen to be run based on the diagnostics, most of these conclusions were substantially unchanged when running off the default settings of burn-in of 10,000, with 40,000 sampling iterations of which every tenth was sampled giving a sample size of 4,000. The BIC values and coefficient ranges were however notably more variable to changes in other algorithm specifications when run on a lower number of simulations. These factors suggest that latent space models provide a highly useful and flexible tool in such contexts, providing quantitative and qualitative insights to the nature of communities and the influential elements of relationships in the network.

## References

- Cooley, D. (2017). googleway: Accesses google maps apis to retrieve data and plot maps. *R package version, 2*(0).
- Department for Education (2019). Schools, pupils and their characteristics: January 2019. <https://www.gov.uk/government/statistics/schools-pupils-and-their-characteristics-january-2019>.
- Grömping, U. (2015). Variable importance in regression models. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7(2):137–152.
- Handcock, M. S., Raftery, A. E., and Tantrum, J. M. (2007). Model-based clustering for social networks. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 170(2):301–354.
- Hoff, P. D., Raftery, A. E., and Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098.
- Krivitsky, P. N. and Handcock, M. S. (2008). Fitting position latent cluster models for social networks with latentnet. *Journal of Statistical Software*, 24.
- Kruskal, W. (1987). Relative importance by averaging over orderings. *The American Statistician*, 41(1):6–10.
- Lindeman, R. H., Merenda, P., and Gold, R. Z. (1980). Introduction to bivariate and multivariate analysis, glenview, il. *Scott: Foresman and company*.

- Pratt, J. W. (1987). Dividing the indivisible: using simple symmetry to partition variance explained. In Pukkila, T. and Puntanen, S., editors, *Proceedings of the Second Tampere International Conference in Statistics*, pages 245–260. University of Tampere, Finland.
- Thomas, D. R., Zhu, P., Zumbo, B. D., and Dutta, S. (2008). On measuring the relative importance of explanatory variables in a logistic regression. *Journal of Modern Applied Statistical Methods*, 7(1):4.