Appendix 1. Detailed methods for landscape (1km sample square) selection

Mark F. Hulme and Gavin M. Siriwardena

Introduction

A consistent methodology for site selection was required in order for results to be comparable to those obtained from other regions in the European FarmLand BiodivERsA network, which include study regions in Spain, Germany and France, along with a sister study region in Canada. Site selection was based on the work of Pasher et al. (2013) who developed a method for the Canadian study (Ontario), aimed at reducing correlations between dependent variables, reducing spatial autocorrelation and ensuring sufficient variation between sites to maximise the possibility of detecting the effects of heterogeneity.

In 2012, three heterogeneity variables were considered, as follow:

Compositional landscape heterogeneity Shannon diversity (or richness and evenness separately) of production area, i.e. crop diversity

Configurational landscape heterogeneity: Scale of the pattern Total length of all agricultural field edges

Configurational landscape heterogeneity: Interspersion/Contagion Shannon edge diversity of the production area only (i.e., mask out all non-production areas)

In 2013, the third variable was dropped after discussion amongst the wider international consortium to simplify the process and because the first two alone provided axes describing both compositional and configurational heterogeneity.

Methods

The approach taken followed a protocol based on Pasher et al. (2013) for Canada and agreed in principle by the European consortium. The UK work was then undertaken in winter 2011-2012, a year ahead of the other European groups, and subsequent consultation among the European partners led to revisions to the process in winter 2012-2013, ahead of the second UK field season and the start of all fieldwork elsewhere in Europe.

The international standard protocol for landscape selection described a complete process to translate satellite photography (raster) data into information both on field (and other land-use) unit sizes, shapes and distributions and on field contents, the latter being identified remotely and ground-truthed. For the UK version of this work, a limited time was available for the landscape selection process (two months versus more than 12 months elsewhere in

Europe) and the pre-existence of high-quality land-use mapping meant that some parts of the protocol were not needed, although they then made some other tasks necessary.

Agricultural Extent

To select representative agricultural sites with a limited range of area of non-agricultural habitat, which is likely strongly to influence biodiversity independent of variation in heterogeneity within agricultural land, landscapes with a range of 60-90% agriculture were considered for site selection in 2012, following what was done in Canada. However, in the East Anglian landscape, 1km squares selected using this criterion were observed often to be non-typical because most of the landscape features considerably less non-agricultural land and, where is does occur in large quantities, it is generally urban, plantation forest or other anthropogenic cover (e.g. quarry or major motorway), unlike the natural forest habitats found in Canada. As well as being sub-optimal locations in which to study the effects of farming in East Anglia, such areas are often difficult to survey due to access restrictions. Therefore, after discussion with the broader consortium, areas with a larger component of agricultural land (90-100%) were considered in 2013. However, this contrast between years provides an important benefit for the broader international collaboration in that it potentially links the Canadian and European contexts by applying both approaches in a single study region.

The principal source of land-use data used was the Centre for Ecology and Hydrology Land Cover Map 2007 (LCM), a remote-sensed vector data set in which land-use is assigned to one of 32 categories, including “arable & horticulture”, “improved grassland”, “deciduous woodland”, etc.. In addition, some sub-divisions of these “broad habitats” are available, but these are limited to “arable bare” and “arable unknown” in the case of the “arable & horticulture” category, so provide little information using for measuring variation in cropping. In the production of this map, 25m pixels are assigned to cover types and discrete parcels of each type are built from data-driven aggregations of similar pixels. In addition, for the 2007 map, parcels were cut by known land-use boundaries (including field boundaries) as drawn in the Ordnance Survey MasterMap, the UK’s national foundation mapping system. Thus, the LCM consists of parcels of contiguous land-use (within the categories the remote sensing is able to distinguish) whose boundaries are either known field or other boundaries or detected discontinuities in the satellite imagery within fields, woods or towns. In practice, there are many of the latter as the mapping is conservative with respect to presumptions that land-use is constant within fields, so there are many more parcels in the LCM than there are discrete fields or patches of woodland.

LCM data were obtained and overlaid in ArcMap 10.0 with a 1km x 1km British National Grid for the 100km x 100km squares TF, TG, TL and TM, which cover the East Anglia region in eastern England (NB much of three of these squares falls in the sea, so the total area considered was considerably less than 40,000km^2). The area of LCM parcels labelled with their dominant land cover at Broad Habitat level as “arable or horticulture” or “improved grassland” was summed to give the area of agricultural land and the percentage of the total area of each square calculated. Of those squares in the target range of agricultural cover (60-90% in 2012, 90-100% in 2013), a buffer of 500m around the border of each square was created, creating an area with a maximum width of 2km, and percentage area of agriculture within each square was calculated. This process was repeated for buffers increasing in 500m intervals up to a buffer of 3000m, giving a maximum total width of 7km. Squares were then filtered such that the final selection contained 60-90% or 90-100% agriculture, depending on the year, across all scales from 1km across to 7km across.

Cropping or field contents and field boundaries

In order to calculate the heterogeneity variables, the most accurate information possible was required on field boundary locations and field contents. The LCM provides accurate data on land-use within parcels, but many parcel boundaries do not match real field boundaries and the land-use categories are coarse with respect to the heterogeneity of farmland because crops cannot be distinguished. In order to determine crop composition on agricultural land, additional remote-sensed data were therefore required. The only relevant information available was in the form of Landsat images, which were obtained from USGS (http://glovis.usgs.gov/) for 2011. Ideally, images from multiple dates would have been used to cover different stages of the growing season, in order to differentiate between crops more easily as vegetation cover varied through the spring and summer dependent on the crop. However, for the two Landsat tiles covering the East Anglian region, images from just two dates between April and August, i.e. covering the growing season and avoiding overlap with growth of crops for harvest the following season, had sufficient cloud-free coverage, 23rd April and 26th June. The two tiles for each date were combined and Normalised Difference Vegetation Index (NDVI) values were calculated for each pixel using the following formula:

$$NDVI = \{NIR – R\} / \{NIR + R\}$$

where NIR and R are the Landsat near infrared band 4 and red band 3 brightnesses, respectively. This produced an index of vegetation type for use in identifying a proxy for crop identity (see below).

As noted above, the LCM frequently identifies parcel boundaries within fields that do not represent real crop edges or field boundaries. However, real field boundaries (without information on crop contents) were available in the original Ordnance Survey Master Map, another vector polygon data set. In order to merge adjacent LCM parcels that were actually part of the same field, so that metrics could be calculated based on real field boundary locations, the Master Map polygon data (within which all agricultural land is included as part of a broad land-use category labelled only as “general surface”) were acquired and were overlaid with LCM polygons, labelled with as much field content information as was available in the LCM, i.e. “Broad Habitat” (BH) and Broad Habitat sub-class (BHsub). Adjacent agricultural parcels within the same Master Map polygon for which BH and BHsub were the same were dissolved to approximate individual agricultural fields, but retaining divisions where there was evidence of variation in cropping in the BHsub classification within permanent field boundaries. This produced a map of the spatial distribution of field boundaries and field shapes, but only limited information on field content.

In order to add cropping information, the NDVI raster data were merged with the field polygons. This was achieved by converting the raster 25×25m pixels to a point theme and averaging the NDVI values within each field polygon on merging points and fields. Since the pixels close to parcel boundaries are likely to incorporate some “bleeding” of remote-sensed colour from the boundary vegetation or the adjacent parcel because of the limits of pixel resolution, the matching was done using a version of the field polygon theme in which the fields were buffered by -30m, removing the outer 30m from each parcel. These buffered parcels were then spatially joined to shapefiles consisting of points derived from the NDVI values for pixels in the Landsat rasters for each date. For each “arable or horticulture” parcel, mean NDVI values for each date were calculated in SAS.

Qualitative differences in average pixel-specific NDVI between fields (i.e. a proxy for crop type) were derived first using hierarchical cluster analysis (HCA) conducted using the FASTCLUS procedure in SAS and then selecting cluster in ArcMap. HCA was used to produce 100 clusters with 23rd April and 26th June NDVI values as the two input variables, weighted by the number of pixels contributing to the mean NDVI in each parcel. Parcels in clusters made up of 100 or fewer parcels were removed from the dataset, thus removing outliers – there was a natural gap between clusters of <100 parcels and larger ones so this was partially determined by the data. A biplot of April and June NDVI values for the remaining parcels was then created, revealing four clear concentrations of points. All parcels after outlier removal were thus assigned to one of four clusters, each nominally representing a different crop type. For the Shannon index of compositional heterogeneity, small and outlier parcels that had unknown “crops” were omitted from the calculation. All agricultural boundaries were important for the configurational heterogeneity indices, so to maximise the amount of field boundary used in the calculation, the small and outlier parcels that had effectively disappeared from the dataset when 30m internal buffers were calculated were assigned the cluster value for the nearest parcel with the same LCM BH and BHsub values that had a positively identified cluster identity.

The proxy for crop identity described above fundamentally identifies only perceivable differences in the appearance of field contents, not actual differences in crops. In order to calibrate the differences as far as possible, records of field-specific variation in crop type that were contemporaneous with the remote-sensed NDVI values were sought. The BTO/JNCC/RSPB Breeding Bird Survey (BBS) involves volunteer surveys of 1km squares based on the British National Grid and includes habitat monitoring. Data from seven BBS squares in agricultural areas in East Anglia where crop types had been noted during 2011 were obtained in order to ground-truth the cluster results. Mean NDVI values for 23rd April and 26th June were extracted for the agricultural LCM 2007 parcels in each of these BBS squares, which included fields of sugar beet (a spring-sown root crop), oilseed rape (winter-sown broadleaf) and winter cereal crops (Table 1). Comparison with the cluster classifications indicated that spring sugar beet appeared in the cluster in the top left corner of the biplot, but that winter oilseed rape and (winter) cereals were in the top right cluster, not sufficiently separated to be identifiable from each other (Table 1, Figure 1). This indicates that the cluster results represent some, but not all, of the known variation in cropping, but also capture additional variation that could not be tested using the limited ground-truthing information available. It is likely that the crops covered by the latter include sparser winter-sown crops such as field beans and other spring-sown crops, of which potatoes, onions and spring cereals are notable locally.

Overall, the above means that the best possible land-use map available was less than perfect, but that it probably captured some of the ecologically most important factors in crop variation, such as timing of sowing. Heterogeneity variables were calculated from the mapped data as described below. Note that these variables were used only for the selection of study areas; indices were recalculated from field data for use in analyses of relationships between heterogeneity and biodiversity (see Objectives 2 and 4).

Table 1. Mean NDVI values for identified crops in 7 BBS surveys across East Anglia

Crop	NDVI April 23rd 2011			NDVI June 26th 2011			Cluster
	N	Mean	SD	N	Mean	SD	
Spring Sugar Beet	28	0.118	0.039	28	0.570	0.141	Top Left
Winter Oilseed Rape	19	0.448	0.046	19	0.522	0.081	Top Right
Winter Cereal	44	0.552	0.092	44	0.444	0.143	Top Right

Figure 1. Biplot of NDVI values, in which each point represents one LCM parcel. Average pixel values are shown, omitting the outer 30m to avoid "bleeding" of colour from field boundaries into assigned crop values. Concentrations considered as different "crops" are shown as different colours, which ground-truthing identified as winter cereal and winter oilseed rape (blue, top right) and spring sugar beet (red, top left). The other, smaller concentrations were not identified in ground-truthing.

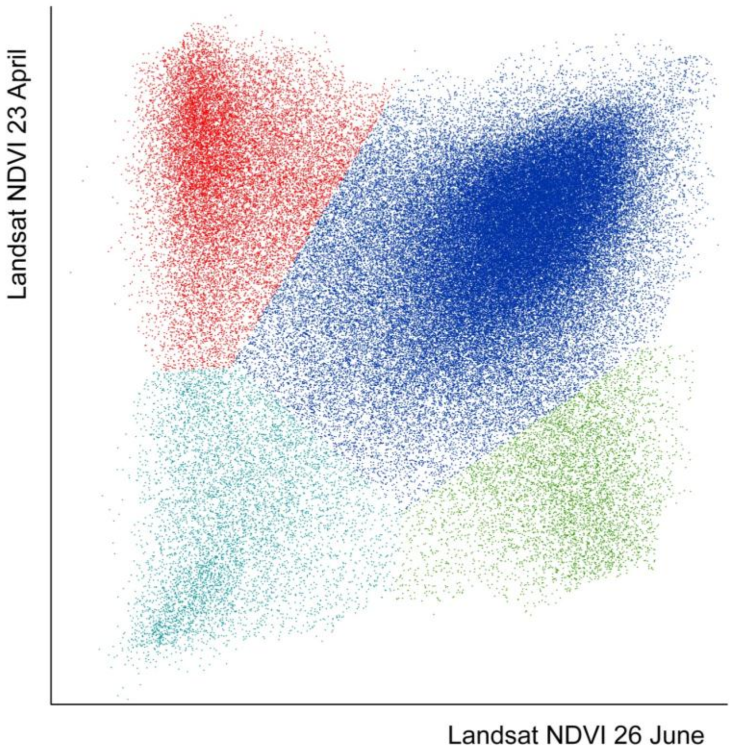

Compositional Heterogeneity

Compositional heterogeneity for each square was calculated as the Shannon Diversity index based on the proportional area within the square of fields of each of the four cluster types, divided into "arable bare" and "arable unknown" arable broad habitat sub-classes from the LCM, and that of improved grassland (i.e. nine "species"). This index was also calculated at each scale from 2km to 7km for all squares with 60-90% agriculture, as selected above.

Equation 1

$$H' = -\sum_{\mathrm{i}=1}^{R} p_i ln\, p_i$$

Where *p*=Proportion of area for cluster type *i* and *R* is the number of "crops" (9).

Configurational Heterogeneity: Total length of all agricultural field edges

The map of agricultural fields described above was used to determine the total length of agricultural field edges in each 1km square and in each buffered area at the larger spatial scales from 2km to 7km. The polygon theme containing all boundaries of arable or agricultural grass fields was converted into a line theme, defining individual line segments as *agricultural boundaries* if they had agricultural habitat LCM parcels on both sides, and the line length summed for each square at each buffer size.

Configurational Heterogeneity: Interspersion/Contagion Shannon edge diversity of the production area only

This measure was considered in landscape selection in 2012, but subsequently dropped for 2013 after agreement across the international consortium. The agricultural field polygon theme was again converted into a line theme, defining individual line segments with respect to the habitats on either side and retaining just those lines that had agricultural habitat LCM parcels on both sides. Different types of boundary were then those that had different combinations of the identified habitat types on either side, considering all available information, i.e. the four NDVI cluster types, improved grassland and the available LCM broad habitat classes and sub-classes. The heterogeneity measure was then the Shannon diversity index, calculated as in equation 1, but with p_i=(boundary length of type i $\times$ number of sections of type i) / (boundary total length $\times$ total number of boundary sections).

<u>Variance</u>

Variability in measured land-use is expected to fall as scale increases, because the finer grain variations are lost. However, it is important to consider a range of scales because the critical scale influencing biodiversity might be large. Moreover, results will have greater generality if the scale at which analyses are conducted is generally representative of other scales. For each of the three heterogeneity variables, the coefficient of variation (CV) was calculated and plotted against buffer scale, from 1km to 6km (Figure 2). CV values fell with scale, but were generally higher than those found by Pasher et al. (2013), so homogenization at large scales was less of an issue in this study. Nevertheless, the configurational heterogeneity variables, in particular, were rather invariable, so it was important to choose a scale for the remainder of the study at which variation was relatively high, but still representative of the variation at broader scales. There was little change in CV once the scale increased above 3km, so this buffer size was taken as a maximum for the calculation of indices used for the rest of the landscape selection process. It is also important to note that larger buffer sizes necessarily lead to less choice for independent areas within a given study region, so the sue of buffers greater than 3km wide would progressively limit selection in practice.

Figure 2. Coefficient of Variation for three heterogeneity variables at scales from 1km x 1km squares to 7km wide buffers. CVs decline at a lesser rate beyond 3km so 1km to 3km scales were used in subsequent investigations.

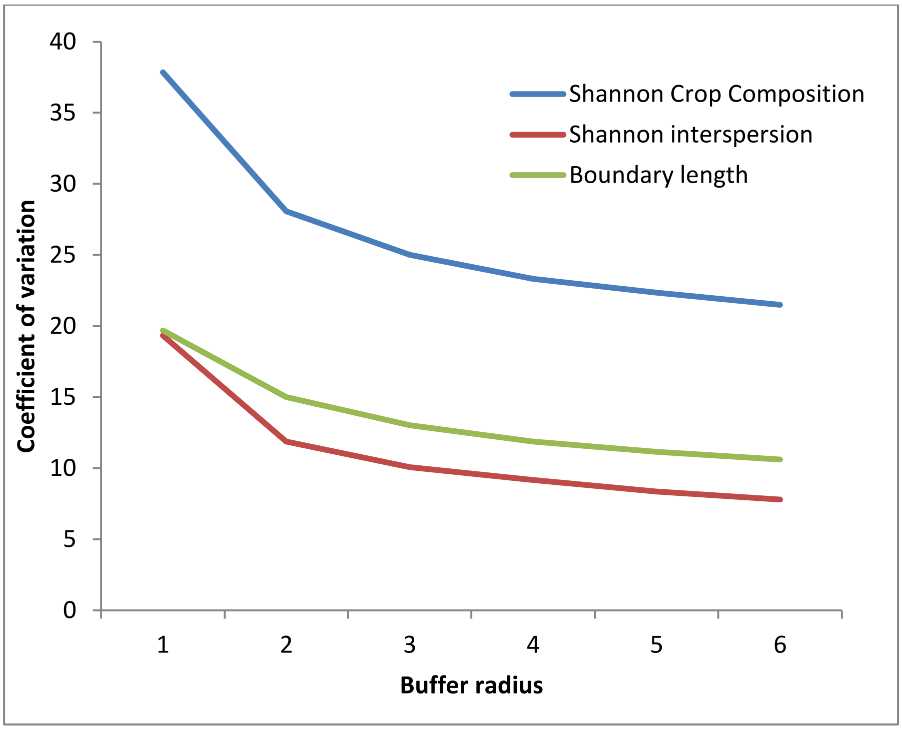

Spatial Autocorrelation

Moran's I measure of spatial autocorrelation was calculated for each heterogeneity variable for calculated for 1km squares, 2km buffers and 3km buffers. This was using lag distances between the central x and y coordinates of the 1×1km squares on which each study landscape was based, from a lag distance of 1km, increasing by 1km up to a lag of 30km. These values of Moran's I were plotted against lag distance (Figure 3). The decrease in Moran's I starts to level off at around 8km so, in order to minimise spatial autocorrelation, an ideal minimum distance between the central points of selected 1km^2 study landscapes was set at 8km.

Figure 3. Moran's I values for three heterogeneity values and three scales with lag distances from 1km to 30km from the centre of the 1km x 1km squares. Spatial autocorrelation levels out between a lag of 5km and 10km so a lag of 8km was selected as an ideal minimum distance between selected squares.

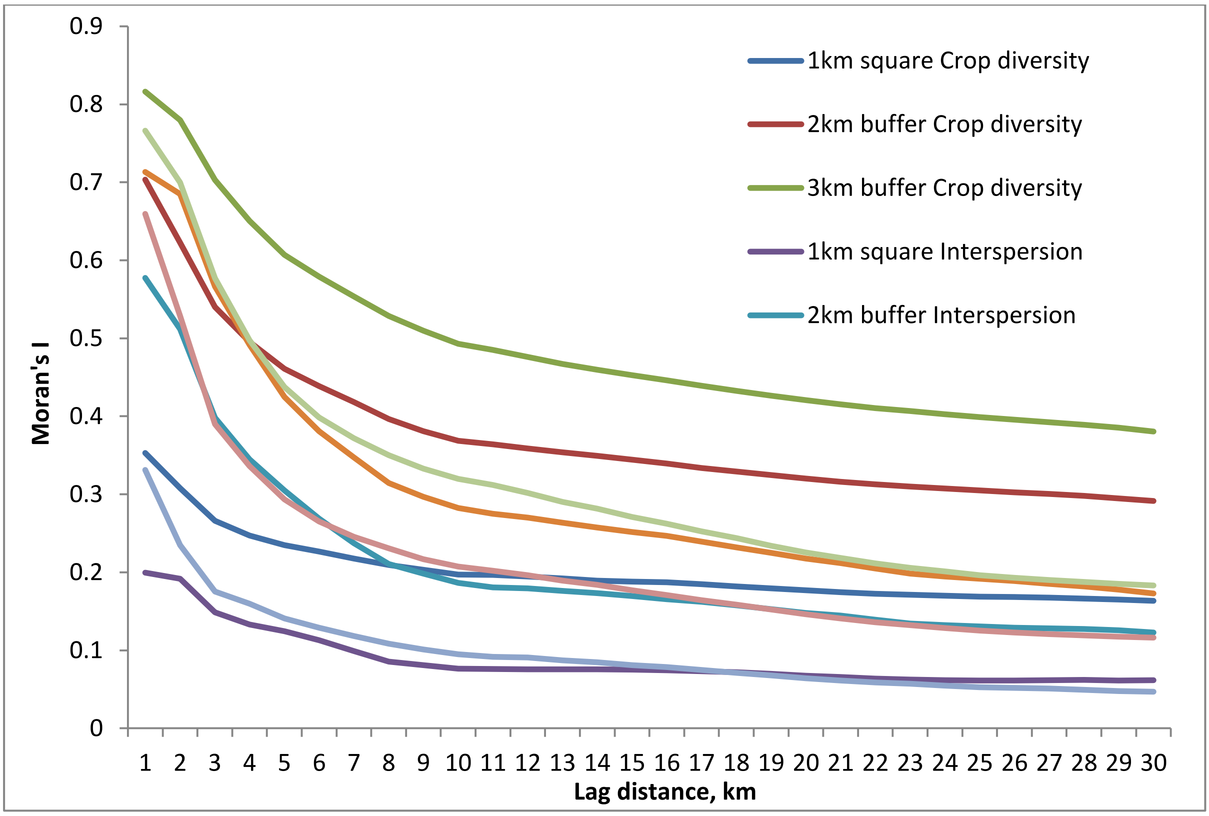

Minimising Correlation between Variables

When selecting squares to cover a wide range of values in the heterogeneity variables, it was also important to minimise correlation between these variables to allow separation of the influences of the different forms of heterogeneity on biodiversity. This was done by selecting squares from the outer corners of a tri-plot (in 2012) and a bi-plot (in 2013) of the heterogeneity variables by square, i.e. avoiding squares along the diagonal of the plot, where their use would tend to increase the correlation between variables (after Pasher et al. 2013). Correlations between analogous variables at different spatial scales were high (>0.890), so the rest of the process used the largest extent (3km wide buffer), from which squares in the lowest 2.5% and the highest 2.5% of the range of each variable were removed (leaving values between the 2.5 and 97.5 percentile) to remove outliers and minimise the inclusion of non-representative sites. The values for the remaining squares (95% of the original ranges) were split into 20 bins of equal size (each representing 5% of the remaining ranges). Correlations between variables describing different forms of heterogeneity were low (<0.23 in 2012, <0.42 in 2013). In 2012, values in the eight different combinations of high and low ranges, or corners of a tri-plot, for the three variables were investigated. "Cubes" of increasing size were selected in each corner, starting with the smallest, 5% x 5% x 5%, then increasing in size until between 20 and 30 values were present in each selection. From this selection, sites were selected such that roughly equal numbers of sites from each of the eight corners, totalling 30 sites, were included, whilst checking that the correlations between forms of heterogeneity remained low in the sample selection. In 2013, the same approach was taken, but using a bi-plot of two variables, with four corners.

The 30 squares from the corners of the tri-plot or bi-plot were treated as candidate study areas each year and initial contact made with known landowners, followed with contact with other landowners as necessary. If agreement for fieldwork access was not forthcoming for a significant part of the farmland in the square, an alternative square was selected from the same corner of the tri-plot/bi-plot, although the paucity of possible squares in some corners meant that selection criteria (size of the bin describing the corner and distance from other survey squares) had to be relaxed in a few instances. Some squares were also rejected in 2012 because the high proportion of non-agricultural land (relative to the norm for the landscape) required for site selection meant that these potential squares had large areas of quarries or motorways that made access for survey work impractical.

Results

After the selection process, 30 squares had been identified in each year that met the land-use criteria and had sufficient access permissions agreed to allow the various fieldwork tasks to be conducted. These squares are shown in Figure 4. Inter-correlations between the heterogeneity variables for these samples were as shown in Table 2. The inter-relationships between the variables are then shown in more detail in Figure 5.

Table 2. Correlations between heterogeneity variables used in study area selection for the final selections of squares in each year. Note that the diversity of agricultural boundaries was not used in 2013.

Variable 1	Variable 2	Correlations	
		2012	2013
Total boundary length	Shannon diversity of agricultural land-use	0.330	0.459
Total boundary length	Shannon diversity of boundaries	0.205	-
Shannon diversity of agricultural land-use	Shannon diversity of boundaries	0.241	-

Figure 4. Locations of survey 1km squares in East Anglia: 2012 (blue) and 2013 (brown).

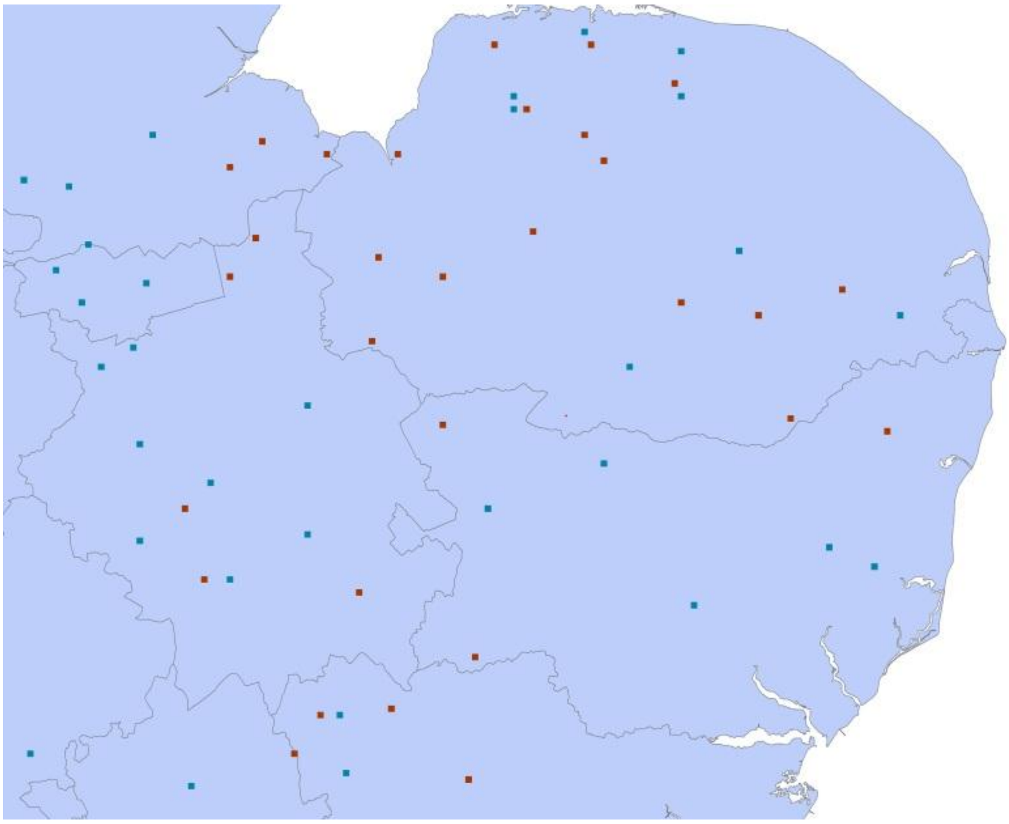

Figure 5. Graphs showing the relationships between the heterogeneity variables used in study area selection for the final selections of squares in each year: (a) Shannon index of land-use composition versus total boundary length index of configuration in 2012, (b) Shannon index of boundary diversity (also configuration) versus total boundary length index of configuration for 2012, (c) Shannon index of boundary diversity versus Shannon index of land-use composition for 2012, (d) Shannon index of land-use composition versus total boundary length index of configuration in 2013.