Advance Access publication November 24, 2015

Political Analysis (2016) 24:121–129
doi:10.1093/pan/mpv029

Crossing the Boundaries: An Implementation of Two Methods for Projecting Data across Boundary Changes

Max Goplerud

Harvard University, Department of Government
e-mail: goplerud@g.harvard.edu

Edited by R. Michael Alvarez

Much of the data used in social science is aggregated into spatial units, even if the analysis itself does not explicitly incorporate that information. A key concern with such aggregation, however, is that changes in the units of aggregation themselves cause difficulty in comparing data gathered on the old boundaries and the new boundaries. Such changes present serious concerns to researchers who may exclude observations or cases due to a lack of comparable units or omit certain key variables. While geographers have long examined this problem and created methods of projecting data from one spatial unit into another, known as areal interpolation,[1] it is telling that a recent article notes the difficulty for researchers in implementing even the most basic solutions without relying heavily on programming skills or proprietary software.[2]

This note makes three contributions toward resolving this problem; first, its accompanying R program allows researchers to easily implement two simple—but flexible—methods of areal interpolation in any context for which the relevant spatial representations (shapefiles) exist. Second, it applies these methods to a specific set of variables that are widely used in political science (electoral results) and confirms that the methods produce generally accurate results using real data gathered from six heterogeneous countries (Australia, Canada, Great Britain, Germany, New Zealand, and the United States). Third, it runs Monte Carlo simulations to examine the performance of the methods given violations of the assumptions about the underlying distribution of the variables of interest.

1 Areal Interpolation

The problem addressed in this note is well known in the geographical literature as the Modifiable Areal Unit Problem (MAUP) and also relates to the problem of ecological inference.[3] A simple example makes this clear; Fig. 1 illustrates two sets of randomly generated constituencies (hereafter used as a generic term that applies to any spatial unit) that partition a unit square, where the solid gray and dashed black lines demarcate the old and new constituencies, respectively.

If we knew the value of the variable of interest in each area defined by the intersection of the old and new boundaries, there would be no problem of areal interpolation—one could simply examine

Author's note: I thank Catherine de Vries, Andy Eggers, Danny Dorling, Ron Johnston, Iain McLean, Juta Kawalerowicz, the anonymous reviewers, and the editors for helpful comments. Any remaining errors or omissions are my own. This work uses shapefiles based on a number of sources that are acknowledged in full in Appendix A (available on the *Political Analysis* website). Replication materials and the R program itself can be found on the *Political Analysis* Dataverse at http://dx.doi.org/10.7910/DVN/MVIQUD. Supplementary materials for this article are available on the *Political Analysis* Web site.

[1]Openshaw (1984).
[2]Qiu, Zhang, and Zhou (2012).
[3]Gotway Crawford, and Young (2004); Fisher and Langford (1995); Lam (1983).

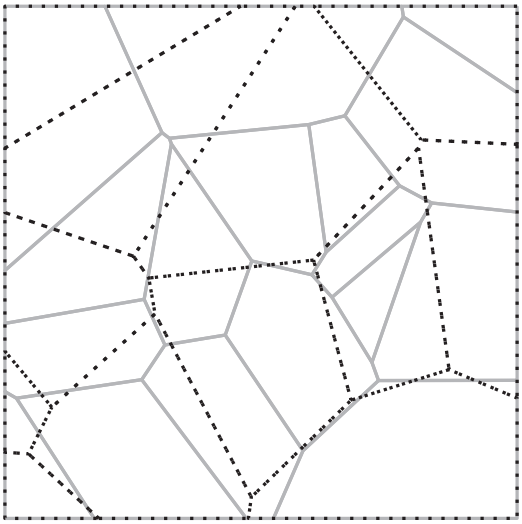

Fig. 1 The problem of areal interpolation.

the sub-constituency variable totals and aggregate them to produce the accurate projections. This "first-best" solution is possible when the variable in question is known at the level of units *at least as small* as the intersection of the old and new boundaries. For example, in the context of elections studied below, this would occur in case where small units, e.g., polling stations or local wards, were used to "construct" the old and new constituencies and where electoral returns were released at this level.

Unfortunately, the first-best method of direct tabulation is not always possible for two key reasons. First, governments simply may not release information at a sub-constituency level. Second, even if such data exist, there may be substantial changes in the sub-constituency units between boundary changes that prevent simple aggregation. This problem likely becomes more severe when interpolating temporally distant sets of boundaries.

2 Methods of Areal Interpolation

In suboptimal circumstances, the solution to the MAUP is to make some assumption about the underlying distribution of the variable of interest. My program implements methods that make the simplest assumption possible: the variable of interest is uniformly distributed through the constituency and/or sub-constituency unit. In the absence of sub-constituency information, this seems to be a reasonable prior.

As this assumption is demonstrably false for many phenomena, a large literature exists discussing different methods of interpolation.[4] However, most of the additional methods require detailed, high-quality information about key covariates on a detailed (sub-constituency) scale that simply may not be present for historical data. Some researchers have also noted that using poor-quality ancillary information, i.e., information with a weak relation to the variable of interest, might cause detailed methods to perform *worse* than the simplest methods.[5] Thus, while incorporating detailed variable- and case-specific information would likely improve the method's accuracy, the assumptions allow the program to be used in a wide variety of empirical contexts.[6] As long as the researcher has spatial representations of the old and new boundaries, my program will provide an estimate of the interpolated values.

[4]Gotway Crawford and Young (2004) provide an overview.
[5]For example, Sadahiro (1999).
[6]Further work might also extend this program to include other methods of interpolation, e.g., smoothing the data based on the values of adjacent units. See Tobler (1979); Gotway Crawford and Young (2004) for some alternative methods.

The two methods implemented in the program are "areal weighting" (AW) and "dasymetric interpolation" (DI).[7] The AW method uses the physical overlap of the constituencies to determine the transfers. This is the simplest possible method and can be used in the absence of any ancillary data; one needs only detailed representations of the old and new boundaries. This method, however, assumes a uniform distribution of the variable of interest at the *constituency* level and does not take into account the (certainly) varying population density of each constituency. The DI method addresses this by incorporating population data gathered on a smaller geographic scale, collected approximately at the same time as the event on the old boundaries. It assumes that each geographic sub-unit contains the variable of interest in proportion to its share of the constituency's population—again uniformly distributed inside each sub-unit. Ideally, the sub-units should be chosen such that they are rarely "split" by the new boundaries, as this makes the assumption of sub-unit uniformity less important. If that is not possible, however, one can rely on comparably fine-grained data to estimate the distribution of population in each constituency. As this method seeks only to approximate the underlying distribution of population, the nearest census (even if it slightly pre- or post-dates the event being interpolated) may be an especially useful resource.

To formally outline the interpolation methods,[8] assume there are J old constituencies, K new constituencies, and L wards, i.e., a generic term for units of geography that are more numerous than the electoral districts, indexed by j, k, and l, respectively. Each ward contains some number of individuals (p_l) and each unit of geography has some area, e.g., s_l for a given ward. The implementation of this program maps "smooth" shapefiles onto discrete units of area ("pixels").[9] This implies that each ward has some "density" of individuals per pixel: $\theta_l = \frac{p_l}{s_l}$. Let the notation $s_{x,y}$ indicate the area (in pixels) of the portion of (ward) y inside constituency x. The interpolation program calculates the distribution of pixels ($s_{j,l}$) into old constituency-ward combinations. These results are stacked into a matrix for each new constituency ($\mathbf{\Gamma}^k$ with elements $s^k_{j,l}$) where each old constituency and ward occupy a distinct row and column, respectively. For the distribution of wards into the old constituencies $\mathbf{\Gamma}^O$, one can simply sum the K matrices together ($\mathbf{\Gamma}^O = \sum_k \mathbf{\Gamma}^k$).

Each interpolation method relies on weights assigned to each pixel; Equation 1 shows how these weights are calculated for each constituency-ward paring in both the AW ($\lambda_{j,l}$) and DI ($\delta_{j,l}$) methods. Areal weighting assumes that each ward has the same density (1). This is equivalent to stating that population is evenly distributed across all pixels inside a constituency, and therefore the transfers capture the physical overlap of the constituencies.

$$\lambda_{j,l} = \frac{1}{\text{Pop}_j} \quad \text{where} \quad \text{Pop}_j = \sum_{l \in L} (\Gamma^O_{j,l}) \tag{1a}$$

$$\delta_{j,l} = \frac{\theta_l}{\text{Pop}_j} \quad \text{where} \quad \text{Pop}_j = \sum_{l \in L} (\theta_l \cdot \mathbf{\Gamma}^O_{j,l}) \tag{1b}$$

From here, one creates the transformation matrix—**A** or **D**, for AW and DI methods, respectively—that generates by simple matrix multiplication (e.g., **DX**) the projection of a set of variables gathered on J old constituencies onto the K new constituencies. The elements of **D** and **A**, subscripted k, j, indicate the proportion of any variable in old constituency j that is included in new constituency k. Equation 2 shows how **D** and **A** are calculated. Each element is simply the weighted sum of the density of each pixel ($\lambda_{j,l}$ or $\delta_{j,l}$) and the number of such pixels in new constituency k.

$$\mathbf{A}_{K,J} = \{a_{k,j}\} \quad \text{where} \quad a_{k,j} = \sum_{l \in L} \lambda_{j,l} \cdot s^k_{j,l} \tag{2a}$$

[7] The first method is regularly used as a "baseline" in articles comparing multiple techniques. Gregory (2002) provides a description of the latter, though the notation differs and the term "dasymetric" refers to many different types of interpolation that incorporate ancillary information.

[8] A more detailed discussion of the program itself is provided in the Supplementary Information.

[9] While an important parameter is the number of pixels, i.e., the level of detail of the "map" of each set of boundaries, the program is fairly robust to changing this parameter—see Appendix C.

$$\mathbf{D}_{K,J} = \{d_{k,j}\} \quad \text{where} \quad d_{k,j} = \sum_{l \in L} \delta_{j,l} \cdot s^{k}_{j,l} \tag{2b}$$

Most of the final matrices are extremely sparse, thus permitting quick calculation of the interpolated values by selecting only the relevant sub-matricies. As the sum of each column in the transformation matrix is one, this method ensures that interpolated variables have the same sum across boundary changes.[10]

3 Empirical Application

An interpolation method's utility can be judged by the extent to which its predictions agree with the "actual" results calculated using the first-best method or other methods that are sensitive to non-uniform distributions of the underlying variables. I test the methods by examining their performance in interpolating electoral results in six countries: Australia, Canada, Germany, Great Britain, New Zealand, and the United States.[11] I chose these countries for two reasons; first, shapefiles exist for old and new boundaries, as well as sub-constituency units. Second, they vary on a number of dimensions, such as electoral system, number of constituencies, geographic heterogeneity of constituencies, and the nature (partisan or non-partisan) of boundary reforms. Unfortunately, data limitations imply that only a select number of elections can be tested here.

In all of the countries included except Great Britain, I evaluate the results against the first-best method using results based on official sub-national electoral returns.[12] As the British government does not release any sub-constituency returns for parliamentary elections, I compare the program's methods against one employed by two British academics. Their method uses detailed information on local election results to estimate the non-uniformity of each major party's vote share in the sub-constituency units; they contend that high support for Party X in a local ward l suggests that its support at parliamentary elections in l is higher than the constituency average. Thus, the transfer of ward l from the old constituency to a new constituency should move proportionally more of the support for party X in the old constituency into the new constituency than a uniform distribution would suggest.[13] This seeks to approximate the first-best method, and it has been applied to the three boundary changes since 1979. It is seen as the "gold standard" in the discussion of British elections, with most broadcasters, government agencies, academics, etc., using their results without modification. Thus, comparing the methods in my program against theirs serves as a useful test of the effects of the uniformity assumption versus a different method of dealing with boundary changes.

I test my program in the following fashion; I project the constituency vote totals for each major party (as well as the number of valid votes cast and the size of the electorate) onto the new boundaries.[14] I then calculate the projected share of the (valid) vote for each party in each constituency, as this measure is comparable across countries and more directly of interest to researchers. To assess the error in my projections (i.e., my projections against the "first-best" solution—or its approximation in Britain), I rely on statistics used in the existing literature on areal interpolation. Zhang and Qiu provide an extensive discussion of possible measures from

[10]This is because both methods preserve the pyncophylatic property (Tobler 1979) of the data.
[11]A broader issue, especially in non-proportional systems, is whether the strategic behavior of voters and parties would differ under the new boundaries. I set aside this issue as outside the scope of this note.
[12]Appendix A describes the data and it can be downloaded from the *Political Analysis* Dataverse, citing Goplerud (2015). The source of the US data—the Harvard Election Data Archive—lacked consistent, comparable, precinct-level information (linked to the necessary shapefiles) on congressional elections (in 2008 or 2010). Thus, to provide a fuller test of this method by including more states with potentially differing patterns of sub-constituency vote distributions, I collapsed the 2008 presidential vote onto the relevant House boundaries and interpolated this variable. To the extent that voting behavior in presidential elections tracks voting behavior in legislative elections, the gerrymandering and spatial autocorrelation present in congressional districts will appear in the collapsed presidential voting data and thus will provide an adequate test of the accuracy of the interpolation program.
[13]Rallings and Thrasher (2007). Appendix A provides further details and discusses the limitations of their method.
[14]Appendix A contains further information on each party included as well as the party-specific interpolation error.

which I use the RMSE (root mean-squared error) and MAE (mean absolute error).[15] To get a sense of the distribution, I also report the middle 90% of the Absolute Error (AE) distribution of the party shares interpolated. I pool the projection errors for each of the major parliamentary parties when calculating each measure, as this captures the possibility of different parties having different degrees of sub-constituency heterogeneity.[16] Finally, I created a measure of "correct prediction" (CP) by noting the percentage of cases where the interpolation methods and the "first-best" results predict the same first-place party. To provide a hard test of the accuracy of my interpolation methods, I report this measure only for *marginal* seats, i.e., those in which the difference between the first- and second-place party is below 10%.[17] In all of the results reported, I exclude constituencies with only minor changes.[18] As Table 1 illustrates, both the AW and DI methods perform well in all countries. While the DI method generally performs better, both have a mean absolute predictive error of around 2–3%. The rates of "correct prediction" in marginal contests are more variable, though the DI method more clearly outperforms the AW method in some elections. As Appendix A shows, when all constituencies are included (regardless of marginality), both methods "correctly predict" the first-place party in each constituency at least 90% of the time. Overall, the methods perform least well in the United States, but this is perhaps unsurprising due to the existence of gerrymandering—discussed in the next section.

4 Simulations

This section runs Monte Carlo simulations to examine the robustness of the projection methods to variation in the constituencies and the underlying voting population. I examine three key factors that likely affect the accuracy of the program: the homogeneity of the constituencies, the homogeneity of the voters, and the existence of gerrymandering in the boundary change process.[19] The simulations discussed below take a unit square as representing the geographic area of interest.

I first simulate spatially correlated data to represent the proportion of individuals in any given area who vote for the government. This is generated using a Gaussian process, described in Appendix B, and then rescaled to a unit interval. I then independently simulate a second set of spatially correlated data to represent the population in any given area. The expected population value in any given location, in the absence of spatial correlation, is 100.[20] Multiplying the two sets of data generates a smooth underlying field that represents the "exact" distribution of government support. A key parameter here is the extent of spatial autocorrelation in the share of those supporting the government (r_{share}) and in the population (r_{pop}). I vary this by changing the "range" in the variogram that generates the random fields. The range represents the distance at which spatial autocorrelation has decreased to zero; as I specify an exponential variogram, the range is approached asymptotically and thus the "effective range" (approximately three times the "actual" range) is reached when 95% of the spatial autocorrelation has disappeared. Higher ranges create greater autocorrelation, leading to larger "clustering" of support (or lack thereof), while lower ranges approximate data generated in an i.i.d. random fashion at all locations. My expectation is that the interpolation method should perform better in the presence of lower spatial autocorrelation.

I randomly create constituencies in this square by drawing N i.i.d. points from inside this square using a truncated normal distribution with mean μ and standard deviation σ. I use those points to

[15]Zhang and Qiu (2011).
[16]Party-specific errors are reported in Appendix A.
[17]Appendix A varies this threshold to show how the CP rate changes (rises) as more "safe" contests are included.
[18]Appendix A outlines the exact rule and perform sensitivity analysis by varying the threshold at which a constituency is sufficiently "changed." The rule used here is approximately equivalent to excluding constituencies where a single transfer from an old constituency made up 90% of the new constituency.
[19]Existing simulations of areal interpolations tend to rely on creating random combinations of "real" underlying units for which a variable is known to create "constituencies"; Fisher and Langford (1995) is an early and influential example. Others, notably Sadahiro (1999, 2000), have employed a related strategy to the one used here.
[20]Other parameters are specified in Appendix B but are held constant across all simulations.

Table 1 Error of interpolation program

Country	*Year*	*Method*	*N*	*RMSE*	*MAE*	*Middle 90% of AE*	*CP (Marg)*
Australia	2010	DI	156	2.366	1.550	[0.055, 5.660]	0.833
		AW	144	2.729	1.730	[0.139, 5.398]	0.778
	2013	DI	45	1.426	1.039	[0.047, 2.964]	1.000
		AW	48	2.440	1.470	[0.086, 5.267]	1.000
Canada	2004	DI	611	1.594	0.944	[0.032, 3.700]	0.892
		AW	589	2.417	1.446	[0.057, 5.262]	0.838
	2013	DI	428	2.456	1.662	[0.098, 4.534]	0.816
		AW	444	3.542	2.265	[0.083, 6.766]	0.784
Germany	2005	DI	147	0.495	0.332	[0.013, 1.280]	1.000
		AW	136	0.597	0.332	[0.011, 1.475]	1.000
	2009	DI	288	0.702	0.460	[0.026, 1.546]	0.971
		AW	288	0.689	0.454	[0.017, 1.528]	0.971
	2013	DI	80	0.465	0.334	[0.015, 0.937]	1.000
		AW	70	0.409	0.279	[0.017, 0.800]	1.000
Great Britain	1983	DI	1278	3.099	1.922	[0.048, 6.559]	0.777
		AW	1394	3.589	2.363	[0.075, 7.668]	0.699
	1997	DI	785	3.171	2.288	[0.133, 6.772]	0.766
		AW	826	3.325	2.421	[0.122, 6.929]	0.721
	2005	DI	216	2.840	1.735	[0.064, 5.659]	0.800
		AW	192	3.395	2.224	[0.274, 6.651]	0.889
	2010	DI	546	3.051	2.181	[0.119, 6.244]	0.619
		AW	588	3.241	2.290	[0.148, 6.992]	0.702
New Zealand	2014	DI	220	1.720	0.854	[0.003, 4.282]	0.833
		AW	236	2.825	1.160	[0.003, 5.497]	0.857
United States	2012	DI	656	4.075	2.913	[0.197, 7.494]	0.810
		AW	620	4.441	3.352	[0.212, 9.147]	0.718

Note. Methods: DI - Dasymetric Interpolation; AW - Areal Weighting. *N* indicates the number of observations, i.e., the total number of major party candidates in the constituencies that were changed. In Germany and New Zealand, this pools the constituency-level results from both tiers of the electoral system. The RMSE and MAE are calculated with respect to the vote shares of each party. RMSE - Root Mean-Squared Error; MAE - Mean Absolute Error; Middle 90% of AE - Middle 90% of the Absolute Error Distribution; CP (Marg) - Proportion Correctly Predicted in Marginal (Majority below 10%) Seats.

create the Voronoi partition of the unit square; this creates N polygons such that for any polygon i corresponding to point n, all points inside i are closer to point n than any other point $n' \in N$.[21] The draws from the truncated normal distribution set $\mu = 0.5$ and vary σ; a low σ will tend to generate a number of small tightly clustered constituencies at the center of the square as well as larger constituencies towards the "outskirts" of the square. As σ increases, the truncated normal distribution approaches the uniform distribution (see Appendix B) and thus generates constituencies that are roughly equally sized. Old and new constituency boundaries are independently generated using this procedure. My expectation is that the interpolation methods should perform better with more homogeneous constituencies.

While gerrymandering is a complex process, I create a simple version to implement in the simulations. The program randomly select some proportion (g) of the sub-constituency units (i.e., the set of polygons defined by the intersection of the old and new boundaries) and then reassigns them, if possible, to an adjacent constituency if doing so would decrease the government's share in the gerrymandered new constituency.[22] In the simulations presented below, the gerrymandering procedure causes the median of the mean vote share to rise slightly (by around 2.5% from $g = 0$ to $g = 0.7$) but also causes a noticeable increase in the number of constituencies the

[21]One of the founding works on the MAUP also uses this method: Openshaw (1984, 20).
[22]Appendix B outlines the exact procedure and visualizes the results of the procedure on government vote share and number of seats won.

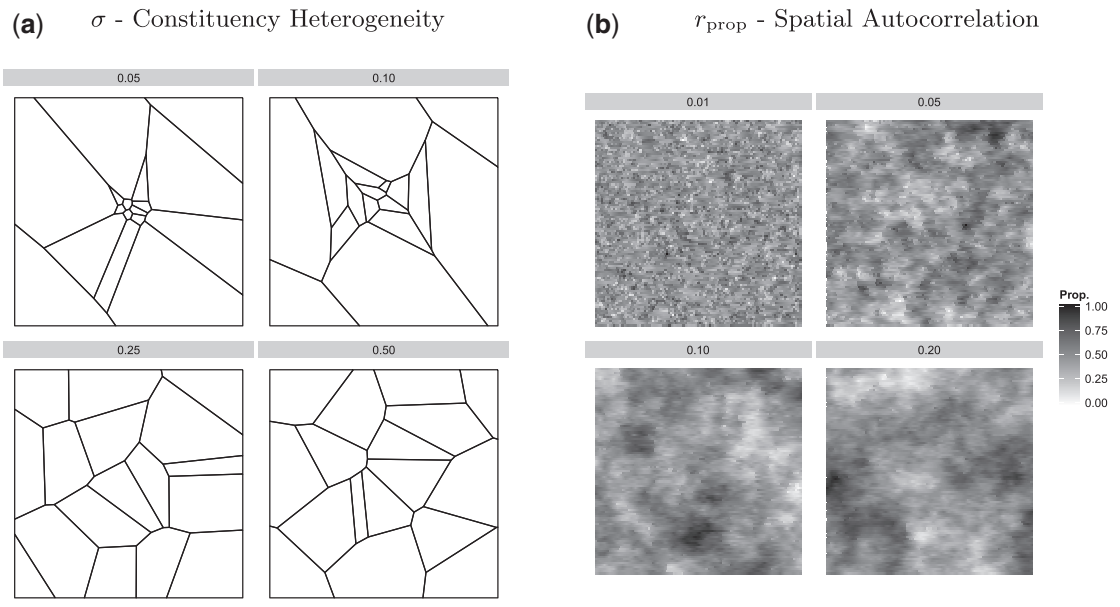

Fig. 2 Illustrations of σ and r_{prop}.

government "wins"; the median government when $g = 0.7$ wins 64% of contests versus 50% without gerrymandering. My expectation is that the interpolation method should perform better with lower levels of gerrymandering.

Figure 2 presents a number of figures for selected values of σ and r_{prop} to give a visual sense of the possible distributions of support and the constituencies. The simulations set σ and r_{prop} as follows to create twenty-five different combinations.

$$r_{\text{range}} \in \{0.01, 0.05, 0.10, 0.20, 0.25\} \quad \sigma \in \{0.05, 0.10, 0.25, 0.50, 1\}$$

I draw fifteen points to form the old constituency boundaries and twenty to form the new set; existing research suggests that methods of interpolation perform less well when the number of units is increased and thus this is a harder test of the methods.[23]

For each simulation, I create a set of boundaries and an accompanying vote distribution. I run the interpolation program and calculate the RMSE and the MAE at each level of gerrymandering $g \in \{0, 0.1, 0.3, 0.5, 0.7\}$. I do this 750 times for each combination of r_{range}, σ. To concisely analyze the results, I present a heat map of the average RMSE and MAE for each combination of the three simulation parameters and interpolation method.

An initial point to note in Fig. 3 is that the interpolation method is a substantively small predictor of either measure of error. The clearer improvement when using actual data suggests that while using dasymetric interpolation is preferable, reliance on simple areal weighting is likely acceptable. The results also confirm the three expectations; increasing the homogeneity of the constituencies (increasing σ) and decreasing the spatial autocorrelation (decreasing r_{range}) corresponds to a lower RMSE and MAE. The effect appears larger for spatial autocorrelation, though this may be a function of moving from very limited spatial autocorrelation ($r_{\text{range}} = 0.01$) to more substantial levels. Despite the simulation's reliance on an "automated" gerrymandering procedure, it still worsens the performance of both interpolation methods. The limited effect of gerrymandering in these simulations should be interpreted cautiously; one could imagine forms of gerrymandering that would cause large

[23] Sadahiro (2000, 81). I add points at the upper-right and bottom-left corners of the square to coerce the Voronoi polygons to a unit square; this means that there are 17 and 22 constituencies in the simulation. One could think of these simulations as focusing on a sub-national region of a country. The results here are likely robust to an increased number of constituencies/zones.

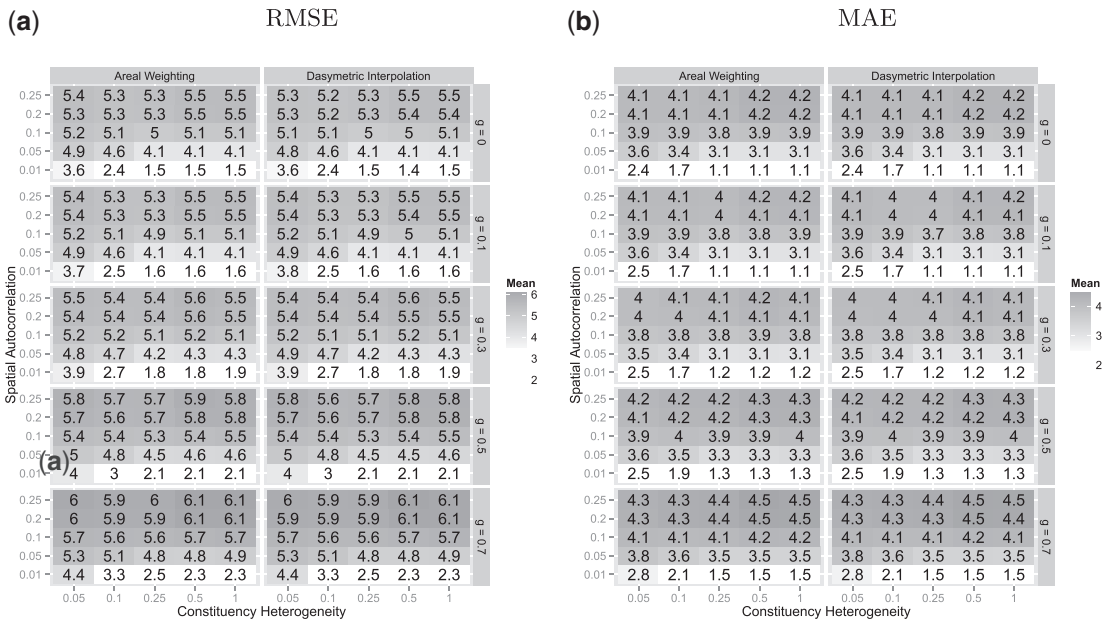

Fig. 3 Heat maps of projection errors.
Note. This figure presents the mean value of the RMSE and MAE for each combination of the simulation parameters. The results are horizontally faceted by interpolation method and vertically by the extent of gerrymandering, i.e., blocks of cells closer to the bottom of the figure report results from simulations with more extensive gerrymandering.

interpolation errors. Rather, the clearest message from the simulations is that high-spatial autocorrelation in party vote share will probably correspond to larger errors in the interpolation program. A plausible explanation for this result is that increasing spatial autocorrelation likely leads to more severe violations of the key assumption of a uniform underlying distribution of support, and thus greater interpolation error will arise.

While it is difficult to directly compare the simulation parameters to the electoral data analyzed in the previous section,[24] examining a rough ordinal ranking of the six countries is illustrative. The country with the worst results for the interpolation program (the United States) has the highest spatial autocorrelation, the most extensive gerrymandering, and high heterogeneity in constituency sizes. The country with the best performance (Germany) has the lowest heterogeneity in constituency size, non-partisan redistricting, but the second highest spatial autocorrelation (albeit only around one-third of that in the United States).[25] As other countries have somewhat less spatial autocorrelation (but more extensive boundary reforms and more heterogeneous constituencies), this could suggest that interactive effect suggested in the heat maps—more homogeneous constituencies (higher σ) correspond to lower interpolation error when spatial autocorrelation is also lower—operates in more subtle ways than the existing simulations captured.

5 Conclusion

This note began by noting that boundary changes complicate the systematic study of many key topics in social science. It implemented and evaluated a systematic, transparent solution to this problem where shapefiles of the aggregate units (and perhaps sub-unit population data) exist. The

[24]Specifically, fitting a variogram to empirical data is quite a complicated process and doing so combines the effects of r_{prop} and r_{pop}. When fitting the variograms using the centroids of each constituency, some models for specific party vote shares did not converge. These are excluded from the following discussion. Secondly, heterogeneity in constituency size must be assed via a different metric; I look at the standard deviation of constituency sizes where each size is normalized to be a proportion of the total area of all constituencies.
[25]It is also worth noting, however, that boundary changes in Germany are fairly limited at each electoral cycle.

R program associated with this note allows researchers to easily and quickly implement these methods for any arbitrary combination of shapefiles and data. Further research in this area might examine the implementation of more complex methods or calculating a measure of uncertainty of the projections.

After outlining their critical simplifying assumptions, this note tested the performance of these methods on electoral data from six countries as well as using Monte Carlo simulations. It showed that the methods perform well—with a mean absolute error of no more than 2–3%. It relied on simulations to further clarify some conditions under which the method should perform less well (heterogeneous constituencies, spatial autocorrelation in the distribution of the vote, and gerrymandering).

Researchers who use this tool should be mindful of the counterfactual it is implicitly assuming; it is important to consider the prevalence of the aforementioned conditions as well as the extent to which the variables of interest are uniformly distributed inside each constituency. However, at its core, the methods implemented here may not be so dissimilar from what researchers already assume. If one believes that using the N^{th} lag of, say, a party's share is an acceptable covariate across boundaries that *have not* changed, then the program used here attempts to produce the same—given the inherent constraints on the data.

References

Fisher, P. F., and M. Langford. 1995. Modelling the errors in areal interpolation between zonal systems by Monte Carlo simulation. *Environment and Planning A* 27(2):211–24.

Goplerud, M. 2015. Replication data for: Crossing the boundaries: An implementation of two methods for projecting data across boundary changes. http://dx.doi.org/10.7910/DVN/MVIQUD, Harvard Dataverse, V1.

Gotway Crawford, C. A., and L. J. Young. 2004. A spatial view of the ecological inference problem. In *Ecological inference: New methodological strategies*, eds. G. King, O. Rosen, and M. A. Tanner. Cambridge: Cambridge University Press.

Gregory, I. 2002. The accuracy of areal interpolation techniques: standardising 19th- and 20th-century census data to allow long-term comparisons. *Computers, Environment and Urban Systems* 26(4):293–314.

Lam, N. S.-N. 1983. Spatial interpolation methods: A review. *American Cartographer* 10(2):129–49.

Openshaw, S. 1984. *The modifiable areal unit problem*. Number No. 28 in concepts and techniques in modern geography. Norwich: Geo Books.

Qiu, F., C. Zhang, and Y. Zhou. 2012. The development of an areal interpolation ArcGIS extension and a comparative study. *GIScience & Remote Sensing* 49(5):644–63.

Rallings, C., and M. Thrasher. 2007. *Media guide to the new parliamentary constituencies*. Plymouth: Local Government Chronicle Elections Centre.

Sadahiro, Y. 1999. Accuracy of areal interpolation: A comparison of alternative methods. *Journal of Geographical Systems* 1(4):323–46.

——— 2000. Accuracy of count data estimated by the point-in-polygon method. *Geographical Analysis* 32(1):64–89.

Tobler, W. R. 1979. Smooth pycnophylactic interpolation for geographical regions. *Journal of the American Statistical Association* 74(367):519–30.

Zhang, C., and F. Qiu. 2011. A point-based intelligent approach to areal interpolation. *Professional Geographer* 63(2): 262–76.