Published in final edited form as:

J Am Stat Assoc. 2016 ; 111(515): 1111–1118. doi:10.1080/01621459.2016.1200918.

Comment: Getting into Space with a Weight Problem

Jon Wakefield^{1,2}, Daniel Simpson³, and Jessica Godwin¹

¹Department of Statistics, University of Washington

²Department of Biostatistics, University of Washington

³Department of Mathematical Sciences, University of Bath

1 Introduction

It is a pleasure to discuss the interesting and wide-ranging article of Diggle and Giorgi (2016), henceforth referred to as DG. Prevalence mapping in low resource settings is an increasingly important endeavor to guide policy making and to spatially and temporally characterize the burden of disease. We will focus our discussion on consideration of the complex design when analyzing survey data, and on spatial modeling. With respect to the former, we consider two approaches: direct use of the weights and a model-based approach using spatial modeling. The first of these is considered in Section 2. With respect to spatial modeling we describe, in Section 3, the stochastic partial differential equations (SPDEs, Lindgren et al. 2011) approach to modeling. Throughout, we use the integrated nested Laplace approximation (INLA, Rue et al. 2009) to perform computation. In general, a spatial target of interest may be associated with a point or an area, and in Section 4 we describe how inference can be made for area averages and probabilities of exceedance of a threshold, using INLA. A simulation to present the power of the INLA/SPDE approach is provided in Section 5. We conclude with final remarks in Section 6.

2 Surveys with a Complex Design

In the developing world, it is often the case that disease indicators are collected via complex survey designs. For example, Demographic Health Surveys (DHS) are nationally-representative household surveys that are carried out extensively in the developing world and typically use a stratified two- or three-stage cluster design (Corsi et al., 2012). Hence, the data are available with accompanying weights and a randomization-(or design-)based approach to inference is common (for a very readable introduction to the analysis of survey data, including randomization-based inference, see Lohr, 2010). In the model-based approach to the analysis of complex survey data (Gelman, 2007), one accounts for the sampling scheme by including the design (e.g., stratification) variables in a regression model. Unfortunately, it is not uncommon for these variables to be unavailable. An alternative approach (Chen et al., 2014; Mercer et al., 2014, 2016) takes the (asymptotic) sampling distribution of a weighted estimator, such as the Horvitz-Thompson (Horvitz and Thompson, 1952) or Hájek (Hájek, 1971) estimator, as the likelihood and then smooths across space and time. Often the survey design is ignored in prevalence mapping (DHS, 2014; Bhatt et al., 2015).

In DG, data from a number of different surveys are analyzed. The rolling malaria indicator survey (rMIS) has a design in which households are randomly selected with a household sampled with probability proportional to village size (Roca-Feltrer et al., 2012). School survey data (Stevenson et al., 2013) are also analyzed by DG; these data do not arise from a standard design, with an iterative process being used for school selection, to limit the chance of overlapping school catchment areas.

We briefly describe the approach described in Mercer et al. (2014, 2016). Let p_k be the unknown prevalence associated with area k , and let $\hat{p}_k$ be the design-based (weighted) estimator of this prevalence with associated estimated design-based variance estimator $\hat{V}_k^*$, $k = 1, \dots, K$. The summaries $\{\hat{p}_k, \hat{V}_k^*, k = 1, \dots, K\}$ may be obtained using standard software, for example, we use the survey package (Lumley, 2004) in R. The svyby function allows the mean prevalence in a region to be estimated, using all of the data collected in that region.

The “data” are then taken as $y_k = \log\left(\frac{\hat{p}_k}{1 - \hat{p}_k}\right)$ and the asymptotic variance of y_k is obtained, from $\hat{V}_k^*$ via the delta method, and is denoted $\hat{V}_k$. If the weighted prevalence estimates are 0 or 1, a fix is required; for example, empirical Bayes may be used. The first stage of the hierarchy is then taken as

$$y_k | \eta_k \sim \text{iid} N(\eta_k, \hat{V}_k), \quad (1)$$

and smoothing models over space can then be applied to η_k to alleviate instability due to small samples, a standard approach in small area estimation (SAE). The use of a normal likelihood based on the empirical logit was used, in a non-complex survey setting, by Stanton and Diggle (2013), but with a constant variance. This model is straightforward to fit in R, since we can use INLA with a fixed and known variance at Stage 1 of the hierarchy, as in (1). Mercer et al. (2016) demonstrate the use of this model when modeling under-5 mortality in Tanzania and we present an example in Section 5.

3 Building Appropriate Spatial Models

We now consider a continuous space Gaussian random field (GRF) model at spatial location $\mathbf{x}_i$ with i indexing points at which responses were measured, $i = 1, \dots, N$, so that N is the number of data points. For the moment we keep things general, and assume a linear predictor of the form

$$\eta_i = \alpha + \mathbf{z}(\mathbf{x}_i)\boldsymbol{\beta} + S(\mathbf{x}_i) + \varepsilon_i,$$

where $\mathbf{z}(\mathbf{x}_i)$ are covariates at the spatial location $\mathbf{x}_i$ with associated regression coefficients $\boldsymbol{\beta}$, $\varepsilon_i \sim \text{iid} N(0, \sigma_\varepsilon^2)$ is measurement error (aka the nugget term), and $S(\mathbf{x}_i)$ represents a spatial GRF.

Many possibilities are available for the form of the covariance function of the GRF but Stein (1999) (amongst others) makes a strong argument for a Matérn function:

$$\text{cov}[S(\mathbf{x}), S(\mathbf{x}^*)] = \frac{\sigma_s^2}{\Gamma(v+1/2)(4\pi)^{1/2} \kappa^{2v} 2^{v-1}} (\kappa \|\mathbf{x} - \mathbf{x}^*\|)^v K_v(\kappa \|\mathbf{x} - \mathbf{x}^*\|)$$

where $K_v(\cdot)$ is a modified Bessel function of the second kind, $\sigma_s^2 > 0$ is a variance parameter, $\kappa > 0$ is a scale parameter and $v > 0$ is a smoothness parameter. When $v + 1$ is an integer, in two spatial dimensions the Matérn fields are Markovian (Rozanov, 1977). Even in this latter case, data analysis that uses the covariance function directly is computationally difficult because of the expensive matrix operations that are required (Rue and Held 2005, Chapter 2).

For modeling the spatial effect, DG use Higdon's convolution kernel approach (Higdon, 1998) in order to control the computational complexity inherent in classical spatial model by replacing a general GRF with a finite dimensional (or low-rank, in their terminology) one:

$$S_n(\mathbf{x}) = \sum_{i=1}^n w_i \phi_i(\mathbf{x}), \quad (2)$$

where the joint distribution of the weights $\mathbf{w} = [w_1, \dots, w_n]^T$ is multivariate Gaussian and the deterministic basis functions $\{\phi_i(\mathbf{x})\}_{i=1}^n$ may depend on some parameters being inferred. The underlying principle is that these finite dimensional random fields will be reasonable proxies for the true latent spatial surface. The advantage of the finite dimensional representation is that inference costs grow like $\mathcal{O}(n^2 N + n^3)$, which, for small enough n , is significantly smaller than the $\mathcal{O}(N^3)$ cost of classical methods. Furthermore, if the basis functions $\phi_i(s)$ are only non-zero in a small part of the domain, the cost is reduced to $\mathcal{O}(N + n^3)$ —or $\mathcal{O}(N + n^{3/2})$ for Markovian models—and the method genuinely grows linearly in the number of basis functions (Simpson et al., 2012a). From this point of view, it is clear that kernel methods (Higdon, 1998), predictive processes (Banerjee et al., 2008), fixed rank Kriging (Cressie and Johannesson, 2008), and the SPDE method (Lindgren et al., 2011) (and for that matter, classical methods like truncated Karhunen-Loève expansions) are all different faces of the same underlying concept. The differences between these methods manifest in the way the basis functions and the weights are chosen. As one would expect, these different choices endow these methods with different sets of advantages and disadvantages (Bradley et al., 2015).

A particular point that we want to emphasize is that the choice of the spatial random field model is not an innocuous one and this choice will filter through into estimates of uncertainty (be they constructed in a Bayesian way or not). In the case where we are interested in predictions at a single unmeasured location, a small forest of results exist on the behavior of spatial point predictions for GRFs under the regime in which the data are very close together (infill asymptotics) or in which the data being collected on an expanding domain (Stein, 1999; Zhang and Zimmerman, 2005). Unfortunately, for the types of models

DG consider, point estimation is not the only summary of interest. In addition, one is interested in estimates of total risk over an area and in locating areas that exceed a threshold; we consider such endeavors in Section 4.

Often, data has both a spatial and a temporal component (such as in the rMIS), and in this case the number of potential asymptotic regimes that we can use to justify our spatial or spatio-temporal model are dizzying. Even more challenging is the idea that for many models the spatial field is designed to model the “residual” effect after the potentially non-linear effects of covariates are taken into account. To the best of our knowledge, the question of how to select the covariance structure of a GRF in the sorts of geostatistical generalized additive mixed models that are increasingly used in practice is completely unstudied.

In our view, the answer is to look for robustness. A little-appreciated fact is that finite-dimensional Gaussian random fields can be spectacularly robust against misspecification. Why? Because the non-robustness in GRFs is driven by the very fine-scale effects, which finite dimensional models necessarily discard. To see this, imagine there is a true value of the underlying spatial field $S^*(\mathbf{x})$, which we can write as

$$S^*(\mathbf{x}) = \sum_{i=1}^n w_i^* \phi_i(\mathbf{x}) + v^\perp(\mathbf{x}),$$

where $v^\perp(\mathbf{x})$ is orthogonal to $\{\phi_i(\mathbf{x})\}_{i=1}^n$. If the basis functions are chosen appropriately, all of the information in the data can then go into estimating the main part of the field, which is modeled by the finite-dimensional GRF, while no assumptions are made about the “fine-scale” effects in $v^\perp(\mathbf{x})$, which are smoothed over. Hence, finite-dimensional GRFs will always get the bulk features right at the expense of the fine-scale ones. This is different to methods like covariance tapering (Furrer et al., 2006), which correctly resolves the fine-scale features necessary for optimal estimation of the field near already observed data points at the expense of resolving the large-scale features (Bolin and Lindgren, 2013).

This discussion gives a lot of insight into how we should choose our basis functions. In the analysis of the rMIS data in DG, for example, the features of interest were village-level prevalence, which suggests that the basis functions $\phi_\lambda(\mathbf{x})$ should be designed to model features on a village scale. We note that this is slightly different from the suggestion of choosing an increasingly large set of basis functions until the inference stabilizes. We are instead suggesting looking at the basis functions themselves to see if they can resolve the types of questions you are interested in. This will lead to very similar results, but is computationally much easier!

For point-referenced data, our preferred modeling strategy is the SPDE approach to spatial modeling, as originally described by Lindgren et al. (2011), and subsequently elaborated upon by Simpson et al. (2012a) and Simpson et al. (2012b). Rather than choosing basis functions according to the convolution square root of the covariance function, as DG do, we instead focus on classes of functions with good approximation properties.

We now describe the approach introduced by Lindgren et al. (2011) to approximate Matérn Markovian Gaussian random fields (MGRFs). The idea is to set up a fine triangular mesh, with m vertices, over the study area. A set of m piecewise linear basis functions $\phi_i(\mathbf{x})$ is then constructed, taking the value 1 at vertex i and 0 at all other vertices, $i = 1, \dots, m$. This gives a set of pyramids that are the building blocks for the approximation. A key point is that these pyramids are non-zero at only a small number of points. The MGRF is again represented by (2), with random Gaussian weights $\mathbf{w} = [w_1, \dots, w_m]^T$. The spatial prior under this model is therefore, in practice, over functions that are linear combinations of the pyramids (i.e., piecewise linear functions over the mesh). The flexibility in choosing the triangular mesh allows careful control of how well the spatial effect is resolved. The general idea is that features that are more than two triangles large are resolved very well, while those that are smaller than the triangle (such as the value of the field at a point) have a bias of the same order as the triangle size. For very precise versions of these results, we refer the interested reader to the technical appendices of Simpson et al. (2016).

The distribution of $\mathbf{w}$ is still required, and is chosen to provide a good approximation to the MGRF. The primary difference between the SPDE approach and the fixed-rank Kriging approach of Cressie and Johannesson (2008), which also recommends using local basis functions chosen for their approximation properties, is the number of parameters that are allowed. While Cressie and Johannesson (2008) aim for a fully flexible model specified with $n(n-1)/2$ parameters, the SPDE approach instead focuses on a more parsimonious specification with, in the simplest case, only 2 parameters: the scale and the range. There are obviously computational advantages to this choice, as well as the parsimony of allowing more straightforward specification of meaningful prior distributions (Fuglstad et al., 2015b). The disadvantage is that the two-parameter model, which essentially corresponds to the assumption that the underlying model is isotropic, is that it may not be flexible enough to correctly model the residual spatial effect.

The MGRF that is to be approximated arises as the solution to the SPDE

$$(\kappa - \Delta)^{\alpha/2} S(\mathbf{x}) = \sigma W(\mathbf{x}), \quad (3)$$

where $\Delta = \frac{\partial}{\partial x_1^2} + \frac{\partial}{\partial x_2^2}$ is the Laplacian and $W(s)$ is white noise. The solution to (3) corresponds to a stationary GRF and if α is an integer then the GRF is Markovian (Whittle, 1954), which is the key for implementation.

A solution to the SPDE satisfies, for any suitable function $\psi(\mathbf{x})$:

$$\int \psi(\mathbf{x}) (\kappa - \Delta)^{\alpha/2} S(\mathbf{x}) d\mathbf{x} = \sigma \int \psi(\mathbf{x}) W(\mathbf{x}) d\mathbf{x},$$

and, these functions are taken to be $\phi_i(\mathbf{x})$, $i = 1, \dots, m$. The use of these test functions leads to a system of linear equations to solve, and the solution produces the distribution of $\mathbf{w}$ which with a little modification, is a Gaussian Markov random field (GMRF). For the

missing details see (Simpson et al., 2012b). The GMRF that we obtain for the weights comes from two places: the fact that the Matérn form that is used is Markovian and the fact that the basis functions are only non-zero across a small portion of the space.

This prior is combined with the likelihood, with the spatial contribution being evaluated as a piecewise linear function of the MGRF at the data locations. Combining the data $\mathbf{y}$ with the above prior we have the posterior on S being of the form (2) but with $w|\mathbf{y} \sim N(\mathbf{0}, \mathbf{Q}_{S|\mathbf{y}}^{-1})$ being the posterior “weights”. In practice, the evaluation of the likelihood at a particular data location turns out to be a weighted sum of the values of the GMRF on the nearest three vertices. The above strategy can be used with a wide range of likelihoods and the SPDE can also be extended to a variety of non-stationary models (Lindgren et al., 2011; Fuglstad et al., 2015a).

4 Area Averages and Excursions

One of the really enjoyable features of DG’s paper is their use of continuously specified Gaussian random fields even when the quantities of interest are areal averages. We broadly think this is a good idea. One concrete reason is that integrating risk over areas allows one to avoid ecological bias, if covariate information is available within areas (Wakefield, 2008). As the authors point out, however, using such fields is a computational challenge. Our favorite engine for overcoming computational challenges is the R-INLA package (Rue et al., 2009; Martins et al., 2013; Lindgren and Rue, 2015). Of three types of problem considered in the paper—estimating area-level prevalence, computing areas where prevalence is above a prescribed level, and using spatially-varying models of zero-inflation—R-INLA can be used to solve two of them (spatially varying models of zero inflation are beyond the functionality of R-INLA for fundamental software design reasons).

When DG say that INLA does not provide the joint predictive distribution for the latent field, they are both right and wrong. By default, INLA computes the univariate predictive distributions and in a lot of cases this is sufficient. It also produces posterior distributions for linear combinations of the latent field, which means that the distribution of the average of the logit prevalence can be obtained. This is, unfortunately, not enough to compute the joint posterior distribution for area-level prevalences. Thankfully, the R-INLA package provides a mechanism for sampling from the joint posterior distribution, which allows one to estimate the distribution of any functional of the latent field. The sampler works by noting that the posterior for the latent Gaussian component, which we will denote by $\boldsymbol{\eta}$, can be approximated by

$$\tilde{\pi}(\boldsymbol{\eta}|\mathbf{y}) = \int_{\Theta} \pi_G(\boldsymbol{\eta}|\mathbf{y}, \boldsymbol{\theta}) \tilde{\pi}(\boldsymbol{\theta}|\mathbf{y}), \quad (4)$$

where $\pi_G(\boldsymbol{\eta}|\mathbf{y}, \boldsymbol{\theta})$ is the Laplace approximation to the full conditional and $\tilde{\pi}(\boldsymbol{\theta}|\mathbf{y})$ is the INLA approximation to the posterior of the hyperparameters (Rue et al., 2009). This could be called an “integrated Laplace approximation” as the full INLA method proceeds from by using another Laplace approximation to approximate the marginal distributions $\pi(\eta_j|\mathbf{y})$.

Empirically, the marginals computed from $\tilde{\pi}(\boldsymbol{\eta} | \mathbf{y})$ are quite close to the marginals produced by the full INLA approximation, although there may be some errors in the higher-order moment. The implementation of this method approximates the integral in (4) making the final approximation to the joint posterior a mixture of multivariate normals.

The second inferential target that DM consider are excursion sets $\mathcal{A}_u^+ = \{\mathbf{x}: p(\mathbf{x}) > u\}$, where $p(\mathbf{x})$ is the probability of a case and u is some fixed threshold. Excursion sets are subtle and difficult beasts that have been studied extensively in both the probability (Adler, 1981; Adler and Taylor, 2007) and statistics (Bolin and Lindgren, 2015; French and Hoeting, 2016) literature. The reason these objects are so hard to study is straightforward: regardless of the statistical philosophy that is being used, the function $p(\mathbf{x})$ is random and hence the set $\mathcal{A}_u^+$ is a random variable. We would therefore want to find a quantile of $\mathcal{A}_u^+$, for example, the set A_u^+ such that

$$\Pr(p(\mathbf{x}) > u \text{ for all } \mathbf{x} \in A_u^+) > 1 - \alpha,$$

for some prescribed level α . DG's approach to estimation is to construct the set

$\tilde{A}_u^+ = \{\mathbf{x}: \Pr(p(\mathbf{x}) > u) > 1 - \alpha\}$. Since the set is constructed pointwise, this is clearly a multiple testing problem and, in general, the set $\tilde{A}_u^+$ will be too big. This is because the pointwise tests do not take into account the fact that $p(\mathbf{x})$ is a continuous function and hence there is strong dependence between nearby tests (points): in order for a function value to be above a threshold with high probability (which is the usual case), all of the surrounding points also need to be above that threshold with high probability. This is similar to the reason that care must be taken when simultaneous bands are calculated for unknown functions using splines see, for example, Wakefield (2013, Section 11.2.7).

The situation is even more challenging in the cases that DG have considered due to the complicated sampling design, i.e., where the data are spatially located. In order to say with high probability that a point is above a given threshold, there needs to be a sufficiently large number of observations nearby to narrow down the pointwise uncertainty. Hence, when you have an inhomogeneous sampling design an excursion set isn't really enough to convey the full information about whether or not you are above a specific threshold. It would be more useful to divide the study area into three distinct regions: the upward excursion set A_u^+ ; the downward excursion set A_u^- , which is the set of all points such that $p(\mathbf{x}) < u$ with high probability; and the set of points that are in neither the upward or downward sets. This then acknowledges that under imperfect information, there are some areas of the space that you cannot with any certainty say are above or below the threshold. This type of target cannot be directly computed in R-INLA. Fortunately, David Bolin has written the excellent excursions package for R (Bolin and Lindgren, 2015), which contains a function for computing these regions using output from the R-INLA package. In the next section we illustrate the calculation of both area averages and exceedence probabilities.

5 Simulation

We now demonstrate the power of the SPDE approach as implemented within INLA. We simulate data within the geography of Kenya, using spatial locations for sampling that correspond to 400 points (enumeration areas) in the 2003 DHS.

For the simulation, we mimic some aspects of the DHS design with enumeration areas assumed to be sampled (as first stage clusters) and then households sampled within areas (as second second stage clusters). The number of positive responses is denoted Y_{ij} , with $i = 1, \dots, n = 400$ indexing the first stage clusters, and $j = 1, \dots, m_i$ representing households sampled within clusters so that m_i is the number of households in first stage cluster i . The sampling model is

$$Y_{ij}|p_{ij} \sim \text{Binomial}(N_{ij}, p_{ij}),$$

with

$$\text{logit } p_{ij} = \beta_0 + S_i + \varepsilon_j,$$

where N_{ij} is the number of individuals in household j of first stage cluster i , β_0 is the intercept (which relates to the overall log odds of prevalence), $S_i = S_i(\mathbf{x}_i)$ arises from a spatial model (which we take as an MGRF) and $\varepsilon_j \sim_{iid} N(0, \sigma_2^2)$ is a random effect that induces dependence between individuals in the same household. In the results we show below, we do not display/include the ε_j terms in prevalence surfaces, as these are assumed to be household specific “noise”.

The prevalences were generated from a GRF with mean prevalence of 7%, so that $\beta_0 = \log(0.07/0.93)$. This mean prevalence was chosen based on the national prevalence of HIV in Kenya estimated in the Kenya DHS 2003. The other parameters of the GRF were taken as $\sigma_1^2 = \sqrt{e}$ and $\kappa^2 = \sqrt{e}$ with noise variance $\sigma_2^2 = 0.01$, in order to produce a prevalence field that matched empirical HIV prevalence estimates from the Kenya DHS 2003 AIDS recode. The number of households in first stage clusters, m_i were taken from the set $(2, 3, \dots, 16)$ which matches the range in the Kenya 2013 DHS. Denominators (household sizes) N_{ij} were sampled from a discrete distribution on $(1, 2, \dots, 12)$ also determined by the empirical distribution of the number of people tested per household in the Kenya DHS 2003 AIDS recode.

We let π_i be the probability that cluster i is selected, and π_{ji} the probability that household j is selected, given PSU i was selected, with $i = 1, \dots, I$ and $j = 1, \dots, m_i$ so that I is the total number of PSUs and m_i is the number of SSUs in PSU i . The design weight for all individuals within household j of cluster i are taken as the reciprocal of the selection probabilities which are

$$\pi_{ij} = \pi_i \times \pi_{j|i} = \frac{400}{46034} \times \frac{m_i}{100}$$

where m_i is assumed to be the pre-chosen number of households to select from the 100 households in cluster i .

As an illustration, we make inference for the prevalence at the level of ADM1 in Kenya, whose areas we index by k , $k = 1, \dots, 47$. For comparison, and to link with Section 2, we also fit the model in which we obtain a design-based estimate, logit p_k , for ADM1 area k , with an associated variance. Letting y_k represent the logit of the weighted (Hájek) prevalence, we have $y_k | \eta_k \sim \text{iid } N(\eta_k, \hat{V}_k)$, and

$$\eta_k = \beta_0^* + S_k + \varepsilon_k, \quad (5)$$

where β_0^* is the area-level intercept, $\varepsilon_k \sim \text{iid } N(0, \sigma_\varepsilon^2)$ are unstructured random effects and S_k are ICAR random effects with variance σ_s^2 . Hence, we use the popular BYM model Besag et al. (1991)

The SPDE mesh is shown in Figure 1 and in Figure 2 we display the posterior median of $S(\mathbf{x})/y$, along with the locations at which samples were obtained. We calculate area-wide summaries of the area level averages,

$$T_k = \int_{A_k} \frac{\exp[\beta_0 + S(\mathbf{x})]}{1 + \exp[\beta_0 + S(\mathbf{x})]} d\mathbf{x}$$

where A_k represent the areas in ADM1, $k = 1, \dots, K$. Posterior means of T_k were constructed in R-INLA using Monte Carlo integration with points $\mathbf{x}_{kj}$, $j = 1, \dots, 100$, simulated in area A_k .

Figure 3(a) displays the true values of T_k and panel (b) the posterior mean estimates obtained from the SPDE model with the Monte Carlo calculation. We emphasize that the latter do not include the household random effects contribution, e_j . The posterior mean estimates from the design-based approach (including the independent and ICAR random effects) are displayed in Figure 3(c). Overall, the SPDE and smoothed design-based (BYM) estimates are quite similar, and display some attenuation as compared to the truth.

Figure 4 shows the estimate of the 95% excursion sets at the 7% prevalence level calculated using the `excursions.inla` function from the `excursions` package. The large blue area is such that $\Pr(\mathbf{p}(\mathbf{x}) < 0.07)$ for every blue point) > 0.95 , while the red areas show the points that simultaneously exceed the threshold. An interesting features of this figure are the black areas, in which there is not enough information to determine with 95% confidence whether the field is above or below the threshold.

6 Concluding Remarks

In Section 5 we considered a very simple situation in which the design was cluster sampling only. Often, stratification is present also (for example, typically in the DHS there is stratification by urban/rural and perhaps on other variables). If the design is stratified cluster sampling, then one may add fixed effects for each of the stratification levels. Post-stratification can be addressed in the model-based framework (Gelman 2007, Gelman and Hill, 2007, Chapter 14). Covariates can also be included, those these need to be known at all locations (at least up to the resolution of the grid) for prediction. Code to reproduce the example in Section 5 can be found at <http://faculty.washington.edu/jonno/cv.html>.

References

- Adler, R. The Geometry of Random Fields. Wiley; New York: 1981.
- Adler, RJ., Taylor, J. Springer Monographs in Mathematics. Springer; 2007. Random Fields and Geometry.
- Banerjee S, Gelfand AE, Finley AO, Sang H. Gaussian predictive process models for large spatial datasets. *Journal of the Royal Statistical Society, Series B.* 2008; 70:825–848.
- Besag J, York J, Mollié A. Bayesian image restoration with two applications in spatial statistics. *Annals of the Institute of Statistics and Mathematics.* 1991; 43:1–59.
- Bhatt S, Weiss D, Cameron E, Bisanzio D, Mappin B, Dalrymple U, Battle K, Moyes C, Henry A, Eckhoff P, et al. The effect of malaria control on plasmodium falciparum in Africa between 2000 and 2015. *Nature.* 2015; 526:207–211. [PubMed: 26375008]
- Bolin D, Lindgren F. A comparison between Markov approximations and other methods for large spatial data sets. *Computational Statistics and Data Analysis.* 2013; 61:7–32.
- Bolin D, Lindgren F. Excursion and contour uncertainty regions for latent Gaussian models. *Journal of the Royal Statistical Society: Series B.* 2015; 77:85–106.
- Bradley JR, Cressie N, Shi T. Comparing and selecting spatial predictors using local criteria. *TEST.* 2015; 24:1–28.
- Chen C, Wakefield J, Lumley T. The use of sample weights in Bayesian hierarchical models for small area estimation. *Spatial and Spatio-Temporal Epidemiology.* 2014; 11:33–43. [PubMed: 25457595]
- Corsi DJ, Neuman M, Finlay JE, Subramanian S. Demographic and health surveys: a profile. *International Journal of Epidemiology.* 2012; 41:1602–1613. [PubMed: 23148108]
- Cressie NAC, Johannesson G. Fixed rank Kriging for very large spatial data sets. *Journal of the Royal Statistical Society, Series B.* 2008; 70:209–226.
- DHS. Technical report. ICF International; Rockville, Maryland, USA: 2014. Spatial interpolation with Demographic and Health Survey data: Key considerations.
- Diggle P, Giorgi E. Model-based geostatistics for prevalence mapping in low-resource settings. *Journal of the American Statistical Association.* 2016
- French JP, Hoeting JA. Credible regions for exceedance sets of geostatistical data. *Environmetrics.* 2016; 27:4–14.
- Fuglstad G-A, Simpson D, Lindgren F, Rue H. Does non-stationary spatial data always require non-stationary random fields? *Spatial Statistics.* 2015a; 14:505–531.
- Fuglstad G-A, Simpson D, Lindgren F, Rue H. Interpretable priors for hyperparameters for Gaussian random fields. 2015b arXiv preprint arXiv:1503.00256.
- Furrer R, Genton MG, Nychka D. Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics.* 2006; 15:502–523.
- Gelman A. Struggles with survey weighting and regression modeling. *Statistical Science.* 2007; 22:153–164.
- Gelman, A., Hill, J. Data Analysis using Regression and Multilevel/Hierarchical Models. Cambridge University Press; 2007.

- Hájek, J. Discussion of, “An essay on the logical foundations of survey sampling, part I”. In: Basu, D.Godambe, V., Sprott, D., editors. Foundations of Statistical Inference. Toronto: Holt, Rinehart and Winston; 1971.
- Higdon D. A process-convolution approach to modelling temperatures in the North Atlantic Ocean. *Environmental and Ecological Statistics*. 1998; 5:173–190.
- Horvitz D, Thompson D. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*. 1952; 47:663–685.
- Lindgren F, Rue H. Bayesian spatial and spatiotemporal modelling with R-INLA. *Journal of Statistical Software*. 2015; 63
- Lindgren F, Rue H, Linström J. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic differential equation approach (with discussion). *Journal of the Royal Statistical Society, Series B*. 2011; 73:423–498.
- Lohr, S. Sampling: Design and Analysis. 2. Boston: Brooks/Cole Cengage Learning; 2010.
- Lumley T. Analysis of complex survey samples. *Journal of Statistical Software*. 2004; 9
- Martins TG, Simpson D, Lindgren F, Rue H. Bayesian computing with INLA: new features. *Computational Statistics and Data Analysis*. 2013; 67:68–83.
- Mercer L, Wakefield J, Chen C, Lumley T. A comparison of spatial smoothing methods for small area estimation with sampling weights. *Spatial Statistics*. 2014; 8:69–85. [PubMed: 24959396]
- Mercer L, Wakefield J, Pantazis A, Lutambi A, Mosanja H, Clark S. Small area estimation of childhood of childhood mortality in the absence of vital registration. *Annals of Applied Statistics*. 2016 To appear.
- Roca-Feltrer A, Lalloo D, Phiri K, Terlouw D. Rolling malaria indicator surveys (rmis): a potential district-level malaria monitoring and evaluation (M&E) tool for program managers. *The American Journal of Tropical Medicine and Hygiene*. 2012; 86:96–98. [PubMed: 22232457]
- Rozanov JA. Markov random fields and stochastic partial differential equations. *Math USSR Sb*. 1977; 32:515–534.
- Rue, H., Held, L. Gaussian Markov Random Fields: Theory and Application. Boca Raton: Chapman and Hall/CRC Press; 2005.
- Rue H, Martino S, Chopin N. Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society, Series B*. 2009; 71:319–392.
- Simpson D, Illian J, Lindgren F, Sørbye S, Rue H. Going off grid: Computationally efficient inference for log-Gaussian Cox processes. *Biometrika*. 2016; 103:49–70.
- Simpson D, Lindgren F, Rue H. In order to make spatial statistics computationally feasible, we need to forget about the covariance function. *Environmetrics*. 2012a; 23:65–74.
- Simpson D, Lindgren F, Rue H. Think continuous: Markovian Gaussian models in spatial statistics. *Spatial Statistics*. 2012b; 1:16–29.
- Stanton M, Diggle P. Statistical analysis of binomial data: generalised linear or transformed Gaussian modeling. *Environmetrics*. 2013; 24:158–171.
- Stein, M. Interpolation of Spatial Data: Some Theory for Kriging. Springer; 1999.
- Stevenson J, Stresman G, Gitonga C, Gillig J, Owaga C, Marube E, Odongo W, Okoth A, China P, Oriango R, et al. Reliability of school surveys in estimating geographic variation in malaria transmission in the Western Kenyan highlands. *PloS one*. 2013; 8:e77641. [PubMed: 24143250]
- Wakefield J. Ecologic studies revisited. *Annual Review of Public Health*. 2008; 29:75–90.
- Wakefield, J. Bayesian and Frequentist Regression Methods. New York: Springer; 2013.
- Whittle P. On stationary processes in the plane. *Biometrika*. 1954; 41:434–449.
- Zhang H, Zimmerman DL. Towards reconciling two asymptotic frameworks in spatial statistics. *Biometrika*. 2005; 92:921–936.

Figure 1.

The mesh has two main features. The first is an inner section, in which the triangles are relatively fine. This is the area that we are most interested in. Outside of this inner area, the triangle rapidly become much larger as they get further from the area of interest. This structure mostly eliminates the boundary effects naturally associated with Markovian models.

Figure 2.

The median posterior spatial effect $S(x)$. The smoothness of the field reflects the relatively low amount of information in the data. In this simulation, points more than $\sqrt{8e} \approx 4.6$ units apart are essentially uncorrelated. Hence, in a large part of the space, the estimated spatial contribution is approximately equal to the prior median 0. At these points, the estimated prevalence is driven entirely by the country-level mean.

(a) True area average prevalence

(b) Monte Carlo estimate of average, SPDE

(c) Design-based estimate, ICAR

Figure 3.

(a) True area-level prevalence averages, (b) estimated area-level prevalences averages from SPDE, (c) estimated area-level prevalences averages from smoothed design BYM model.

Figure 4.
The joint posterior regions where the prevalence is simultaneously estimated to be below (blue) or above (red) 7% at a 95% confidence level. In the black areas, the results are indeterminate.