

A multi-resolution Gaussian process model for the analysis of large spatial data sets.

Doug Nychka
Soutir Bandyopadhyay
Dorit Hammerling
Finn Lindgren
Stephen Sain

NCAR Technical Notes

National Center for
Atmospheric Research
P. O. Box 3000
Boulder, Colorado
80307-3000
www.ucar.edu

NCAR TECHNICAL NOTES

<http://library.ucar.edu/research/publish-technote>

The Technical Notes series provides an outlet for a variety of NCAR Manuscripts that contribute in specialized ways to the body of scientific knowledge but that are not yet at a point of a formal journal, monograph or book publication. Reports in this series are issued by the NCAR scientific divisions, serviced by OpenSky and operated through the NCAR Library. Designation symbols for the series include:

EDD – Engineering, Design, or Development Reports

Equipment descriptions, test results, instrumentation, and operating and maintenance manuals.

IA – Instructional Aids

Instruction manuals, bibliographies, film supplements, and other research or instructional aids.

PPR – Program Progress Reports

Field program reports, interim and working reports, survey reports, and plans for experiments.

PROC – Proceedings

Documentation or symposia, colloquia, conferences, workshops, and lectures. (Distribution maybe limited to attendees).

STR – Scientific and Technical Reports

Data compilations, theoretical and numerical investigations, and experimental results.

The National Center for Atmospheric Research (NCAR) is operated by the nonprofit University Corporation for Atmospheric Research (UCAR) under the sponsorship of the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

National Center for Atmospheric Research
P. O. Box 3000
Boulder, Colorado 80307-3000

Date 2013-Dec

A multi-resolution Gaussian process model for the analysis of large spatial data sets

Douglas Nychka

Institute for Mathematics Applied to Geosciences
National Center for Atmospheric Research, Boulder, CO

Soutir Bandyopadhyay

Department of Mathematics
Lehigh University, Bethlehem, PA

Dorit Hammerling

Institute for Mathematics Applied to Geosciences
National Center for Atmospheric Research, Boulder, CO

Finn Lindgren

Department of Mathematical Sciences
University of Bath, Bath, UK

Stephan Sain

Institute for Mathematics Applied to Geosciences
National Center for Atmospheric Research, Boulder, CO

**Institute for Mathematics Applied to Geosciences
Computational and Information Systems Laboratory**

NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

P. O. Box 3000

BOULDER, COLORADO 80307-3000

ISSN Print Edition 2153-2397

ISSN Electronic Edition 2153-2400

A multi-resolution Gaussian process model for the analysis of large spatial data sets.

December 2, 2013

Abstract

A multi-resolution basis is developed to predict two-dimensional spatial fields based on irregularly spaced observations. The basis functions at each level of resolution are constructed as radial basis functions using a Wendland compactly supported correlation function with the nodes arranged on a rectangular grid. The grid at each finer level increases by a factor of two and the basis functions are scaled to have a constant overlap. The coefficients associated with the basis functions at each level of resolution are distributed according to a Gaussian Markov random field (GMRF) and take advantage of the fact that the basis is organized as a lattice. Several numerical examples and analytical results establish that this scheme gives a good approximation to standard covariance functions such as the Matérn and also has flexibility to fit more complicated shapes. The other important feature of this model is that it can be applied to statistical inference for large spatial datasets because key matrices in the computations are sparse. The computational efficiency applies to both the evaluation of the likelihood and spatial predictions. Although our framework has similarities to fixed rank Kriging, the model gives a better approximation to situations where the nugget variance is small and the spatial process is close to interpolating the observations.

Keywords: Spatial estimator, Kriging, Fixed Rank Kriging, Sparse Cholesky Decomposition, Multi-resolution

Contents

1	Introduction	4
2	Spatial model	8
2.1	Process and observational models	8
2.2	Spatial estimate	8
2.3	Radial Basis functions (RBF)	10
2.4	Markov Random fields	11
2.5	Multi-resolution	12
2.6	Normalization to approximate stationarity	14
3	Computation	14
3.1	Normalization to constant marginal variance	15
3.2	Estimating the basis coefficients	15
3.3	Evaluating the determinant of the observation covariance matrix	16
3.4	Conditional Simulation	17
3.5	Timing	18
4	Properties of the covariance model	21
4.1	Process convolution	21
4.2	Numerical approximation	23
5	North American summer precipitation	28
6	Discussions and Conclusions	33

List of Figures

1	Timing results for the lattice/basis model and standard Kriging.	20
2	Theoretical approximation of Matérn correlation functions using the multi resolution process based on process convolutions.	24
3	Approximation of Matérn covariances using the lattice/basis model. . .	26
4	Approximation of a mixture of exponential covariances.	27
5	Stereographic projection of precipitation observation locations.	31
6	Correlation models fit to the precipitation data.	32
7	Spatial predictions for mean summer (JJA) precipitation.	33

List of Tables

1	Maximum likelihood estimates of σ and effective degrees of freedom. . .	31
---	--	----

1 Introduction

Statistical methodology for spatial data is a well developed field and has roots in geostatistics and multivariate analysis. More recently the breakthroughs in Bayesian hierarchical models have added rich new classes of models for handling heterogenous spatial data and indirect measurements of spatial processes. This development in spatial statistics is coincident with emerging challenges in the geosciences involving new types of observations and comparisons of such observations to complex numerical models. For example, as attention in climate science shifts to understanding the regional and local changes in future climate there is a need to analyze high resolution simulations from climate models and to compare them to surface and remotely sensed observations at fine levels of details. These kinds of geoscience applications are characterized by large numbers of spatial locations. The application of standard techniques is often not feasible or at least will take an unacceptably long time given typical computational resources. Moreover, geophysical processes tend to have a multi-scale character over space that requires statistical methods that do not assume a simple parametric model for dependence across a region. This work develops a new statistical model that addresses both of these challenges; our model is applicable to large data sets and has a flexible covariance model and so fills a gap in current statistical methodology.

We assume that spatial observations $\{\mathbf{y}_i\}$ are made at unique two-dimensional spatial locations, $\{\mathbf{x}_i\}$, for $1 \leq i \leq n$ according to the additive model:

$$\mathbf{y}_i = Z_i^T \mathbf{d} + g(\mathbf{x}_i) + \boldsymbol{\epsilon}_i \quad (1)$$

where Z is a matrix of covariates and \mathbf{d} a vector of linear parameters, g is a smooth surface and $\boldsymbol{\epsilon}_i$ are mean zero measurement errors. The parameters \mathbf{d} represent fixed effects in this model and g is assumed to be a realization of a random process with the form

$$g(\mathbf{x}) = \sum_{j=1}^m \mathbf{c}_j \phi_j(\mathbf{x}). \quad (2)$$

Here ϕ_j , $1 \leq j \leq m$, is a sequence of fixed basis functions and \mathbf{c} is a vector of coefficients distributed multivariate normal with mean zero and covariance matrix, ρP . The parameter $\rho > 0$ is useful as a leading scaling parameter for the covariance

matrix, P , that may also depend on other parameters. Thus the model for g is a sum of fixed basis functions with stochastic coefficients.

The statistical problem in this setting is to determine g at locations where observations are not available and quantify the uncertainty of the spatial predictions. Given our main goal to develop an acceptable methodology to handle large data sets, perhaps in a nearly interactive mode, we seek to balance the complexity of the models and methodology with feasibility for effective data analysis. We will focus on maximum likelihood estimates of parameters in the covariance and other components from (1) and (2). For inference on the spatial process we will focus on the conditional distribution of g given the data and other statistical parameters. This basic approach to inference can be expanded to a Bayesian hierarchical approach, where our model and the computational scheme can play a central role in the core of any Bayesian computation.

Our approach combines the representation of a field using a multi-resolution basis with statistical models for the coefficients as a process on a lattice. In this sense it is a blending of ideas from Fixed Rank Kriging (Katzfuss and Cressie 2011, Cressie and Johannesson 2008) and stochastic partial differential equations (SPDE) including the work in Lindgren and Rue (2007), Rue and Held (2005) and Lindgren et al. (2011) (LR2011). It is useful to view the unknown spatial process in (2) as a sum of L independent processes, $g_l(\mathbf{x})$ for $1 \leq l \leq L$, each with a different correlation scale and marginal variance, $\{\alpha_l\}$.

$$g(\mathbf{x}) = \sum_{l=1}^L \sqrt{\alpha_l} g_l(\mathbf{x}) \quad (3)$$

In this way the overall spatial dependence of g can be much more complex than the spatial dependence of each of the individual components. The multi-resolution structure has a similarity with wavelet estimators although we do not assume orthogonality of the basis functions or sparsity in the coefficients. Each component, g_l is defined as a basis function expansion as in (2). The main principle is to expand the two-dimensional spatial field in families of radial basis functions that are organized on regular grids of increasing resolution and correspond to the schematic of a sum of L independent processes. The radial basis functions have compact support and like wavelet bases give

computational efficiencies because of this feature. In our treatment, each increase in resolution will be by a factor of two and the levels associated with finer spatial scales will have more basis functions. Conversely the representation has a parsimony in that the coarser scales require fewer basis functions to approximate the stochastic processes.

The spatial dependence among the coefficients for each level of resolution is modeled using a Gaussian Markov random field (GMRF), specifically a spatial autoregressive (SAR) model. The fact that the basis functions are organized on a lattice gives the SAR a simple form along with its precision matrix, which we denote as $Q = P^{-1}$. The benefit of this approach is that Q is sparse even though P itself can be dense. Thus, g can exhibit long range correlations among coefficients widely separated in the lattice even though the precision matrix is sparse. We have found that this combination of a multi-resolution basis with companion GMRFs for the coefficients at each level can approximate standard families of covariance functions such as the Matérn, but also provides a rich model for more general spatial dependence. It should be noted that we make no assumption on the distribution of observation or prediction locations even though the latent components of our model will exploit regular grids. We are also able to give some analytical results that suggest why this model can approximate a range of spatial processes exhibiting different degrees of smoothness.

Many of the ingredients for this model are not new, however, their particular combination with a view towards efficient computations for large and irregular spatial data sets has not been exploited in previous works. The key is to introduce sparsity into the computations in a way that does not compromise covariance models with long range correlations and models with many degrees of freedom. This is achieved by using compactly supported radial basis functions and computing the *precision* matrix of the basis coefficients directly, not the covariance matrix. In addition we add some normalizations that reduce the degree of artifacts from using a discrete basis. The net result is a flexible covariance model that has rank comparable or greater than the number of spatial locations and where spatial prediction, conditional simulation and evaluation of the likelihood can be done on a modest laptop computer.

Recent work on statistical methods for large spatial data sets has used a fixed rank

Kriging approach to make computations feasible. This can either take the form of a small number of basis functions and an unstructured and dense P matrix such as in Cressie and Johannesson (2008) or large number of basis functions and a sparse model such a Markov random field for Q (Eidsvik et al. 2010). An insightful approach was suggested in Stein (2008) and later in Sang and Huang (2011) where a low rank process was combined with a process that has a compactly supported covariance. This superposition of two processes anticipates our model where we consider a mixture of covariances at multiple scales. Reflecting the fact that the likelihood calculation carries most of the computational cost, there has been work on approximations to the likelihood for spatial models by binning the observations and using spectral methods (Fuentes 2007) or considering a partial likelihood (Michael L. Stein 2004) or pseudo likelihood (Caragea and Smith 2007). Our approach differs from these papers in that we are able to compute the likelihood exactly.

The next section describes the fixed rank Kriging model and its likelihood under a setting where the process and measurement errors have a Gaussian distribution. This section also gives details of the basis function construction and the particular Markov random field used for the basis coefficients. Section 3 outlines the computational algorithm and gives some timing results. The approximation properties of this basis/lattice model are reported in Section 4 with the proofs of the asymptotic results relegated to the Appendix. Section 5 provides an example for a climate precipitation data set for the North American region, which is analyzed using a standard Matérn covariance and the multi-resolution model. Section 6 contains conclusions. Much of the computations in this paper can be reproduced using the `LatticeKrig` package in R, which serves as a supplement for implementing the numerical methods and a ready source for the data set.

2 Spatial model

2.1 Process and observational models

Based on the set up in the introduction g will be a mean zero, Gaussian process with covariance function:

$$COV(g(\mathbf{x}), g(\mathbf{x}')) = \sum_{j,k=1}^m P_{j,k} \phi_j(\mathbf{x}) \phi_k(\mathbf{x}'). \quad (4)$$

With respect to the observation model in (1) we assume that ϵ_i are uncorrelated, normally distributed with mean zero and covariance $\sigma^2 W^{-1}$. Here we assume that σ^2 is a free parameter of the measurement error distribution and W is a known but sparse precision matrix. In most applications W is diagonal or just the identity. Let Φ be the regression matrix with columns indexing the basis functions and rows indexing locations. $\Phi_{i,j} = \phi_j(\mathbf{x}_i)$. With these definitions one can now reexpress (1) in matrix vector notation as

$$\mathbf{y} = Z\mathbf{d} + \Phi\mathbf{c} + \mathbf{e}$$

and collecting the fixed and random components we have

$$\mathbf{y} \sim MN(Z\mathbf{d}, \rho\Phi P\Phi^T + \sigma^2 W^{-1}). \quad (5)$$

As a last step it is useful to reparametrize this model to better mesh with the computations and in some instances to simplify formulas. Let $\lambda = \sigma^2/\rho$ and we reparametrize σ in terms of λ and ρ (i.e. $\sigma^2 = \lambda\rho$). Now set $M_\lambda = (\Phi P\Phi^T + \lambda W^{-1})$ and (5) is the same as

$$\mathbf{y} \sim MN(Z\mathbf{d}, \rho M_\lambda) \quad (6)$$

2.2 Spatial estimate

From (6) we have the log likelihood

$$\ell(\rho, P, \lambda, \mathbf{d}) = (-1/2)(\mathbf{y} - Z\mathbf{d})^T (\rho M_\lambda)^{-1} (\mathbf{y} - Z\mathbf{d}) - (1/2)\log|\rho M_\lambda| + (n/2)\log(\pi)$$

As a component of a full Bayesian model this represents part of the marginal distribution having integrated over \mathbf{c} .

This expression is used to find maximum likelihood estimates (MLEs) of the fixed effects and covariance parameters. In finding these estimates it is convenient to first maximize over the fixed effects and then maximize over some of the covariance parameters. For fixed ρ , P and σ the MLEs for \mathbf{d} are also the generalized least squares (GLS) estimates

$$\hat{\mathbf{d}} = (Z^T M_\lambda^{-1} Z)^{-1} Z^T M_\lambda^{-1} \mathbf{y}. \quad (7)$$

Note this estimate only depends on λ and not on ρ . To maximize ℓ on the remaining parameters set $\mathbf{r} = \mathbf{y} - Z\hat{\mathbf{d}}$ and the profile log likelihood is now

$$\ell(\rho, P, \sigma, \hat{\mathbf{d}}) = (-1/2)(\mathbf{r}^T (\rho M_\lambda)^{-1} \mathbf{r}) - (1/2)\log|\rho M_\lambda| + (n/2)\log(\pi). \quad (8)$$

Finally, the expression given above can be maximized over ρ giving $\hat{\rho} = \mathbf{r}^T M_\lambda^{-1} \mathbf{r} / n$. This estimate can be substituted back into (8) to give a profile log likelihood in λ and in any covariance parameters that contribute to P .

Although in this paper we will focus on maximum likelihood estimation for inference, one could also consider a prior distribution on the covariance parameters and carry out a Bayesian inference for these model components. In doing so and also integrating out the dependence of the posterior on \mathbf{c} , one obtains a log marginal posterior that is the sum of (8) and a log prior density on the covariance parameters and those of \mathbf{d} . Thus except for the influence of the priors, (8) comprises the log posterior for inference on the model parameters. Considering a uniform (and improper) prior on \mathbf{d} , $\hat{\mathbf{d}}$ will also be the posterior mode from a Bayesian perspective. A uniform and improper prior on \mathbf{d} may appear to be strange, but in many geostatistical applications it is a reasonable approximation. Typically the likelihood will be informative for this small set of regression parameters and will dominate any diffuse prior. Thus considering the limiting case as the prior becomes diffuse is a reasonable assumption.

The inference for the basis coefficients depends on the standard results for the conditional normal distribution. Specifically, the conditional distribution of \mathbf{c} given \mathbf{y} and all other parameters in the model at their true values is a multivariate normal

$$[\mathbf{c} | \mathbf{y}, \mathbf{d}, \sigma, \rho, P] \sim MN(\hat{\mathbf{c}}, \rho P - \rho P \Phi^T (M_\lambda)^{-1} \Phi P) \quad (9)$$

with

$$\hat{\mathbf{c}} = P\Phi^T M_\lambda^{-1}(\mathbf{y} - Z\mathbf{d}) \quad (10)$$

This conditional mean, $\hat{\mathbf{c}}$, is taken to be the point estimate of \mathbf{c} and by linearity, the spatial estimate for $g(\mathbf{x})$ at an arbitrary location is $\hat{g}(\mathbf{x}) = \sum_{j=1}^m \phi_j(\mathbf{x})\hat{\mathbf{c}}$. Typically a vector of the spatial covariates, $z(\mathbf{x})$, is also provided at this location. To reproduce the familiar universal Kriging estimator \mathbf{d} is set at the GLS estimate given above and so the full spatial prediction is: $\hat{y}(\mathbf{x}) = z(\mathbf{x})^T \hat{\mathbf{d}} + \hat{g}(\mathbf{x})$.

This estimate is also well-known as the best linear unbiased (BLUE) estimator. The conditional variance of $g(\mathbf{x})$ can be interpreted as the mean squared error from a geostatistical perspective or from a Bayesian perspective as quantifying the uncertainty in $g(\mathbf{x})$ when the other parameters are known. Although there are closed form expressions for the conditional covariance it is usually more efficient computationally to draw samples from the posterior and quantify the uncertainty in the estimate by Monte Carlo sample statistics. In geostatistics this technique is referred to as conditional sampling and an efficient algorithm for the sampling is given in the next section. Typically this conditional sampling is done using the estimated covariance parameters but under a Bayesian analysis one can sample from the marginal posterior for these parameters. By varying the covariance parameters in this manner one also includes the uncertainty in the predictions with respect to estimating the covariance model.

2.3 Radial Basis functions (RBF)

Our full model proposes a multi-resolution basis where each level of resolution takes the same form and so we start with describing a single level of basis functions on a common scale. The basis functions are essentially translations and scalings of a single radial function. Let ϕ be a unimodal, symmetric function in 1-dimension and let $\{\mathbf{u}_j\}$, $1 \leq j \leq m$ be a rectangular grid of points. Here we “unroll” the indices over the two dimensional grid into a single index to make the notation simpler.

Consistent with radial basis function terminology, we will refer to the grid points as *node points* and let θ be a scale parameter. The basis functions are then

$$\phi_j^* = \phi(\|\mathbf{x} - \mathbf{u}_j\|/\theta) \quad (11)$$

Geometrically, the basis will consist of bumps centered at the node points with overlap controlled by the choice of θ . In this work we will take ϕ to be a two-dimensional Wendland covariance (Wendland 1995) that has support on $[0, 1]$. The Wendland functions are polynomials in their argument and can be constructed with different degrees of differentiability at the origin. They are also positive definite, which is an attractive property when the basis is used for interpolation. In this work we use a Wendland function valid up to 3 dimensions and belonging to C^4 :

$$\phi(d) = \begin{cases} (1-d)^6(35d^2 + 18d + 3)/3 & \text{for } 0 \leq d \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

2.4 Markov Random fields

In parallel with the preceding section we describe the stochastic model for the coefficients of a basis constructed at a single level of resolution. The multi-resolution aspect replicates this model at each level. The coefficient vector \mathbf{c} at a single level follows a Gaussian Markov random field (GMRF) and is organized by the node points. We will assume the special case that the coefficients follow a spatial autoregression (SAR). The difference with this model for \mathbf{c} and that in LR2011 is that we define the SAR independently from the choice of basis. An alternative is to find the elements of Q directly as inner products of the basis functions. Specifically if \mathcal{L} is the differential operator motivated from a partial differential equation then

$$[Q]_{jk} = \int \mathcal{L}(\Phi_j)\mathcal{L}(\Phi_k)d\mathbf{x}$$

will provide a finite approximation to the continuous process. We take a more direct approach to specifying Q and this avoids the integrations to determine the inner products. These integrals may become difficult in more than two dimensions or if the lattice is not on a rectangular grid. However, this strategy does have the advantage that there is an unambiguous limiting spatial model as the basis becomes dense. Finally, we note that the discretization of the SPDE based on a particular choice of basis functions especially for a limited basis size, may not remove all of the non-stationary artifacts and so adjustments such as those described in Section 2.6 may still be needed.

Given an autoregression matrix B and \mathbf{e} , a random vector distributed $N(0, \rho I)$, we construct the distribution of \mathbf{c} according to $\mathbf{c} = B^{-1}\mathbf{e}$. The autoregressive interpretation is that $B\mathbf{c} = \mathbf{e}$. That is, B transforms the correlated field to discrete white noise with variance ρ . For our use we will constrain B to be sparse. Each row has a small number of non-zero elements which we will define based on the nodes of the corresponding basis functions. Although one may generalize the elements to other regular node patterns, here the node points will form a rectangular grid. Let \mathcal{N}_j denote the indices of the nearest neighbors of \mathbf{u}_j . For an interior point this will be four neighbors, but less for the nodes at edges and corners. Following LR2011 for interior lattice points we take $B_{j,j} = 4 + \kappa^2$ with $\kappa \geq 0$ and the off diagonal elements to be -1. Although one can modify the weights at the edges of the lattice to approximate free boundary conditions we have found that adding a buffer and keeping zero boundary conditions provides an easier solution. The boundary effects are also diminished by the normalization discussed in Section 2.6. By linearity \mathbf{c} has covariance matrix $\rho B^{-1}B^{-T}$ and precision matrix given by $Q = (1/\rho)B^T B$. Because B is formulated as unconditional weights on the field, any choice of B will lead to a valid covariance and so Q will be positive definite. It is well known that the SAR weights do not specify the Markov structure directly. For nonzero weights on the four neighbors Q will be a sparse matrix with each row having 12 nonzero elements: the first, second and third order neighbors. Thus, \mathbf{c} will be a GMRF conditional on this larger clique of points. The results in LR2011 provide the connection between this GMRF and approximations to the Matérn family of spatial covariances. In this particular case one expects that the SAR from Table 1 will approximate a Matérn process with scale parameter κ in LR2011 and smoothness $\nu = 1$.

2.5 Multi-resolution

From the previous sections we have developed a basis and covariance for a specific grid. The multi-resolution model extends this idea by successively halving the spacing of the grid points and specifying a GMRF for the coefficients at each level. Between levels we assume coefficients are independent. To make this idea explicit assume that the

spatial domain is the rectangle $[a_1, a_2] \times [b_1, b_2]$ and the initial grid $\{\mathbf{u}_j^1\}$ is laid out with $m_x \times m_y$ grid points with the spacing $\delta \equiv (a_2 - a_1)/(m_x - 1) = (b_2 - b_1)/(m_y - 1)$. Note here the constraint that the spatial domain and numbers of grid points are matched so that the grid spacing is the same in the x and y dimensions. Subsequent grids are defined with spacings $\delta_l = \delta 2^{-(l-1)}$ and yield a sequence of grids, $\{\mathbf{u}_j^l\}$ that increase roughly by a factor of four in size from level l to level $l+1$. To define the basis functions for the l^{th} level we take $\theta_l = \theta/2^{(l-1)}$ and define the radial basis functions as in (11). Let L denote the total number of levels, then the (unnormalized) multi-resolution basis is

$$\phi_{j,l}^* = \phi(\|\mathbf{x} - \mathbf{u}_j^l\|/\theta_l)$$

where $1 \leq l \leq L$ and $1 \leq j \leq m(l)$ and $m(l) = (m_x - 1)(m_y - 1)4^{l-1} + m_x + m_y + 1$

The vectors of coefficients associated with each level will be denoted by \mathbf{c}^l and so the multi-resolution representation for g is

$$g(\mathbf{x}) = \sum_{j=1}^{m(1)} \mathbf{c}_j^1 \phi_{j,1}^*(\mathbf{x}) + \sum_{j=1}^{m(2)} \mathbf{c}_j^2 \phi_{j,2}^*(\mathbf{x}) + \dots + \sum_{j=1}^{m(L)} \mathbf{c}_j^L \phi_{j,L}^*(\mathbf{x}) \quad (12)$$

In general we can stack these coefficients as $\mathbf{C} = (\mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^L)$ and the natural extension of the SAR model is a sparse matrix B such that $B\mathbf{C}$ is $N(0, \rho I)$. Although B can be a general matrix we have found it useful to restrict attention to a block diagonal form. Equation (12) is written in this way and is suggestive of the process being a sum of L processes. We assume that they are also independent. Let $\alpha_1, \alpha_2, \dots, \alpha_L$ be a vector of positive weights and for the l^{th} level we assume \mathbf{c}_l follow a GMRF with a SAR matrix, $(1/\sqrt{\alpha_l})B_l$. Here B_l has the same form as in the single level but with the κ parameter possibly depending on the level. One can interpret $\rho\alpha_l$ as parameterizing the marginal variance of the l^{th} level process and κ_l is an approximate scale parameter. Thus we are lead to a block diagonal form for B and also for the precision matrix:

$$\mathbf{Q} = (1/\rho) \begin{bmatrix} (1/\alpha_1)(B_1)^T B_1 & 0 & \dots & 0 \\ 0 & (1/\alpha_2)(B_2)^T B_2 & \dots & 0 \\ 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & (1/\alpha_L)(B_L)^T B_L \end{bmatrix} \quad (13)$$

2.6 Normalization to approximate stationarity

Based on the specific form for \mathbf{Q} we have found it useful to normalize the basis functions to give a better approximation to stationary covariance functions. It is well known that a GMRF on a finite lattice can exhibit edge effects and other artifacts in the covariance model that are not physical. Moreover the radial basis functions having nodes on a discrete set can also contribute to patterns in the implied covariance matrix.

For both free and zero boundary edge adjustments there is substantial difference in the variance over the domain. One obvious correction for this effect is to weight the basis functions so that when (4) is evaluated one will obtain a constant marginal variance. Accordingly, let $\omega(\mathbf{x}) = \sqrt{COV(g(\mathbf{x}), g(\mathbf{x}))}$ from (4) and normalize the basis functions as $\phi_j(\mathbf{x}) = \phi_j^*(\mathbf{x})/\omega(\mathbf{x})$. Because this normalization is tied to the choice of covariance model it means that the basis is no longer independent of the GMRF and this linkage adds more computational overhead. However, the normalization can take advantage of the overall strategy of using sparse matrix decompositions and we believe reducing edge effects and other artifacts is worth the extra computation.

3 Computation

The estimators defined in the previous section can be found efficiently by a judicious use of sparse matrix decompositions and matrix identities. Most of these computations depend on the constructions of Φ , W and Q to be sparse matrices. As an introduction we use throughout the fact that the sparse Cholesky decomposition has two steps: a symbolic factorization to determine the sparsity pattern of the decomposition and then a numerical factorization to compute the triangular matrix. In several instances the sparsity pattern is fixed and only the values of the entries change. Anticipating this case one stores the symbolic factorization and only recomputes the numerical portion. In our R implementation using the `spam` sparse matrix package (Furrer and Sain 2010), we have found that approximately half of the computational time in the decomposition is taken up by the symbolic factorization.

3.1 Normalization to constant marginal variance

Computing the normalization $\omega(\mathbf{x})$ breaks into the following steps:

1. Form the sparse precision matrix \mathbf{Q} .
2. Compute the sparse Cholesky decomposition $\mathbf{Q} = \mathbf{A}\mathbf{A}^T$.
3. Solve the triangular linear system $\mathbf{A}\mathbf{v} = \mathbf{z}$ with $\mathbf{z}_j = \phi_j(\mathbf{x})$.
4. $\omega(\mathbf{x}) = (\sum_{j=1}^m \mathbf{z}_j^2)^{-1/2}$.

For finding the normalization at more than one value of \mathbf{x} the Cholesky decomposition will not change and so steps 1 and 2 need only be done once. Also for changing values of the GMRF the sparsity pattern does not change and so the symbolic factorization of \mathbf{Q} can be reused for different sets of covariance parameters. For some GMRF weights it is possible to identify this problem with a discretization of the classical Poisson problem in two dimensions from numerical analysis. There exist very efficient fast Poisson solvers that could eliminate the computational overhead in this step.

3.2 Estimating the basis coefficients

A common calculation to find the concentrated likelihood in (5) and $\hat{\mathbf{d}}$ in (7) is to evaluate $M_\lambda^{-1}\mathbf{w}$ for an arbitrary vector \mathbf{w} . Recall that $M_\lambda = \Phi P \Phi^T + \lambda W^{-1}$ and taken at face value M_λ is a dense, potentially large matrix and so difficult to work with directly. The strategy is to transform M_λ using matrix identities to involve the sparse precision matrix. The matrix identities themselves are well known and central to reducing the dimensions of the calculations in fixed rank Kriging. They work by shifting the linear system size from matrices with dimensions of the number of observations (n) to dimensions of the number of basis functions (m). Usual fixed rank Kriging has m much smaller than n . In our application, however, exploiting these identities has the potential effect of greatly *increasing* the nominal dimensions of the problem. When the number of basis functions exceeds that of the observations the matrix sizes will be larger but due to sparsity of the matrices the computations are feasible. In fact as reported in Section 3.5 the computations can be much faster than a direct calculation.

The Sherman-Morrison-Woodbury formula (Henderson and Searle (1981)) can be applied to give

$$M_\lambda^{-1} = (\Phi(P)\Phi^T + (\lambda W^{-1}))^{-1} = (W - (W\Phi)G^{-1}(\Phi^T W))$$

where $G = \Phi^T W \Phi + \lambda Q$. Using this identity one can now use the Cholesky decomposition for G to solve the linear system $G\mathbf{v} = (\Phi^T W)\mathbf{w}$ for \mathbf{v} and it follows that

$$M_\lambda^{-1}\mathbf{w} = W\mathbf{w} - W\Phi\mathbf{v}$$

Because Φ , W and Q are all sparse, G will also be sparse. Moreover it also follows that G is positive definite. Thus the linear system can be solved using a sparse Cholesky decomposition. Note that an important limitation of this computational strategy is that λ can not be identically zero. In practice this is not an issue as λ can be set to a small enough value to give a result that is close to interpolation and is still numerically stable.

To find $\hat{\mathbf{c}}$ as in (10) one can use the Sherman-Morrison-Woodbury formula or the representation of the conditional multivariate normal in terms of precision matrices to derive the identity

$$P\Phi^T M_\lambda^{-1} = P\Phi^T (W - (W\Phi)G^{-1}(\Phi^T W)) = G^{-1}\Phi^T W$$

Thus $\hat{\mathbf{c}} = G^{-1}\Phi^T W(\mathbf{y} - Z\hat{\mathbf{d}})$ is found by taking advantage of sparsity of Φ and W for multiplication and the sparse Cholesky factorization of G . The evaluation of $\hat{g}(\mathbf{x})$ can also be computed in an efficient manner if the sum is restricted to basis functions that are nonzero at \mathbf{x} .

3.3 Evaluating the determinant of the observation covariance matrix

The other intensive computation occurs in the likelihood as the determinant of M_λ .

Using elementary properties of determinants

$$|M_\lambda| = |\Phi P \Phi^t + \lambda W^{-1}| = |\lambda W^{-1}| |UU^T + I_n| = \frac{\lambda^n |UU^T + I_n|}{|W|} \quad (14)$$

with $U = (\lambda)^{-1/2}W^{1/2}\Phi P^{1/2}$, adopting symmetric square roots of W and P , and I_n being an $n \times n$ identity matrix.

Now one can use a special case of Sylvester's Theorem. For an $n \times m$ matrix U and identity matrices I_n and I_m . $|UU^T + I_n| = |U^TU + I_m|$.

Applying this identity with the choice of U given above

$$\begin{aligned} |U^TU + I_n| = |U^TU + I_m| &= |(1/\lambda)P^{1/2}(\Phi^TW\Phi)P^{1/2} + I_m| \\ &= \lambda^{-m}|P^{-1/2}|\Phi^TW\Phi + \lambda P^{-1}||P^{-1/2}| \quad (15) \\ &= \lambda^{-m}|\Phi^TW\Phi + \lambda P^{-1}||P^{-1}| \end{aligned}$$

substituting for G and Q we have $|U^TU + I_n| = \lambda^{-m}|G|/|Q|$ and so combining with (14)

$$|M_\lambda| = \frac{\lambda^{n-m}|G|}{|Q||W|}$$

The matrices, W , G and Q are all positive definite and sparse so the determinants can be found efficiently from the product of the diagonal elements of the Cholesky decompositions.

3.4 Conditional Simulation

In this section we describe a Monte Carlo algorithm to sample from the conditional distribution of \mathbf{d} and \mathbf{c} (or g) given the observations and the covariance model. From Section 2 the joint distribution of \mathbf{c} and \mathbf{d} conditioned on \mathbf{y} and the covariance parameters is multivariate normal with mean given by $\hat{\mathbf{d}}$ and $\hat{\mathbf{c}}$. We use the following algorithm to generate a sample from the conditional multivariate distribution.

1. Generate \mathbf{c} as $N(0, \rho P)$ by solving the triangular system $A\mathbf{c} = \mathbf{w}$ where $AA^T = 1/\rho Q$ and $\mathbf{w} \sim N(0, I)$
2. Generate synthetic observations $\mathbf{y}^* = \Phi\mathbf{c} + \mathbf{e}^*$ with $\mathbf{e}^* \sim N(0, \sigma^2 W^{-1})$ and compute estimates $\hat{\mathbf{d}}^*$ and $\hat{\mathbf{c}}^*$ based on \mathbf{y}^*
3. $\mathbf{d}^* = \hat{\mathbf{d}} + \hat{\mathbf{d}}^*$ and $\mathbf{c}^* = \hat{\mathbf{c}} + (\hat{\mathbf{c}}^* - \mathbf{c})$ will be a draw from the conditional distribution for (\mathbf{d}, \mathbf{c}) and a draw from the predictive distribution for \mathbf{y} is given by

$$y(\mathbf{x}) = Z(\mathbf{x})^T \mathbf{d}^* + \sum_{j=1}^m \phi_j(\mathbf{x}) \mathbf{c}_j^*$$

3.5 Timing

Here we present some timings for the computations with the main comparison being the dense matrix computations associated with Kriging. The number of spatial locations was varied between 500 and 20000 and spatial predictions were found for an exponential covariance model and several choices of the lattice multi-resolution model. Note that the computation time is essentially independent of the spatial data, the distribution of spatial locations, and the values of the covariance parameters. The timing is done for the function `mKrig` in the R package `fields` (Furrer et al. 2012) implementing standard Kriging and for the function `LKrig` in the R package `LatticeKrig` (Nychka et al. 2012) implementing the multi-resolution basis function model. Times reported are for a Macbook Pro laptop (2.3 Ghz Intel Core i7) and R 2.14 (R Development Core Team 2011). Both of these functions compute the predictions at the observations for a fixed covariance model, evaluate the likelihood, compute the coefficients for predicting the surface at arbitrary points, and an approximate estimate of the generalized cross-validation function. Despite this varied output from the functions, the Cholesky decomposition in both `mKrig` and `LKrig` dominate the time for large n .

Figure 1 reports the total time for these functions using the R utility `system.time`. The dashed line is the time for the standard “Kriging” estimate using `mKrig` up to 10,000 observations and with times extrapolated to 20,000. Thus the time for 20,000 observations and standard Kriging is estimated to be about 1,300 seconds (about 21 minutes). The solid black line is the time for the function `LKrig` with a single level, the number of basis functions chosen to be approximately equal to the sample size, and with the basis functions normalized to have unit marginal variance. The dotted black line is the same scheme but without normalizing the basis functions. For 20000 spatial locations the times for these models are 31 seconds (normalized) and 5.5 seconds (unnormalized).

The grey lines report timing with the number of basis functions kept fixed. The line labeled 10 has four levels of multi resolution and where the coarsest basis has centers on a 10×10 grid and totaling 7159 basis functions. The line labeled 20 is the same specification but with the coarsest grid being a 20×20 grid with a total of 31,259 basis

functions. The solid lines are timings with normalized basis functions and dotted are without normalization.

These results indicate substantial time savings over the dense matrix computations and evaluations of the likelihood are feasible even for 20000 spatial locations. The unnormalized computation times are particularly striking and are largely dominated by the sparse Cholesky decomposition of the matrix G discussed in Section 3.2. For this work we have not exploited more efficient algorithms in the normalization step and there is some difference between the normalized and unnormalized cases. As might be expected the two covariance models with fixed degrees of freedom (7,159 and 31,259) are closer to being linear as a function of sample size.




Figure 1: Timing results for the lattice/basis model and standard Kriging. Plotted are times in seconds as the number of spatial locations is varied between 500 and 20000. Spatial locations are generated from the uniform distribution on the domain $[-1, 1] \times [-1, 1]$. (However, the particular values spatial locations do not effect these timing results.) The dashed line is the time for the `mKrig` function from the `fields` R package that computes the likelihood and related statistics for an exponential covariance model with a fixed set of covariance parameters using a standard dense matrix Cholesky decomposition. Solid and dotted lines are times for the `LKrig` function from the `LatticeKrig` R package that compute the likelihood and related statistics for a multi-resolution lattice covariance with fixed parameters. Solid lines are times with normalization to a constant marginal variance and dotted lines are times without normalization. Among these cases the black lines are for a single level model where the basis functions are chosen to be roughly equal to the number of spatial locations. The orange lines use a fixed number of basis functions comprising four levels and with the coarsest level being either 10×10 or 20×20 . Text labels identify these cases.

4 Properties of the covariance model

In this section we present some theoretical expansions and some exact computations that outline the properties of the covariance family introduced in Section 2. As a foundation, we first consider a convolution approximation to the sum over radial basis functions.

4.1 Process convolution

It is simplest to define a single convolution process and then extend this to an infinite mixture. Let z be a unit variance, isotropic, two dimensional Matérn process with spatial scale parameter κ , smoothness ν , and $C_\nu(\|\mathbf{x} - \mathbf{x}'\|/\kappa) = E(z(\mathbf{x}), z(\mathbf{x}'))$, the corresponding covariance function. Also let ϕ be a compactly supported RBF with $\phi(0) = 1$. For $\theta > 0$ a scale parameter, define the convolution process

$$g(\mathbf{x}) = \int_{\mathbf{R}^2} \frac{1}{\theta^2} \phi(\|\mathbf{x} - \mathbf{u}\|/\theta) z(\mathbf{u}) d\mathbf{u}.$$

This type of process for statistical modeling is well-established (see Higdon 1998) and as written will be Gaussian, mean zero, and have an isotropic covariance function given by

$$\int_{\mathbf{R}^2} \int_{\mathbf{R}^2} \frac{1}{\theta^4} \phi(\|\mathbf{x} - \mathbf{u}\|/\theta) C_\nu(\|\mathbf{u} - \mathbf{v}\|/\kappa) \phi(\|\mathbf{x}' - \mathbf{v}\|/\theta) d\mathbf{u} d\mathbf{v}. \quad (16)$$

Now consider a sequence of independent Matérn processes, $z_l(\mathbf{x})$ with $\{\theta_l\}$ a sequence of scale parameters for the convolution kernel and “hard wire” $\kappa_l = 1/\theta_l$. These define a sequence of convolution processes $g_l(\mathbf{x})$ according to (16) with the same marginal variance. Finally, let k_l denote the covariance function for the l^{th} process.

Given, non-negative weights $\{\alpha_k\}$ that are summable we are lead to the multi-resolution process that is Gaussian, mean zero and covariance given by

$$k(\mathbf{x}, \mathbf{x}') = \sum_{l=1}^{\infty} \alpha_l k_l(\mathbf{x}, \mathbf{x}').$$

Given this representation, a theoretical question is how the choice of $\{\theta_l\}$ and $\{\alpha_l\}$ influence the properties of k . In particular, is it possible to construct covariances that represent different degrees of smoothness than those implied by the basis functions

and Matérn process used in the convolution? Typically the smoothness of an isotropic, stationary Gaussian process is tied to the differentiability of the covariance function at the origin. An alternative measure is to characterize the tail behavior of the spectral density of the process. Under isotropy the spectral density will be radially symmetric and we focus on the decay rate as r increases. In particular, for spectral densities whose tails are bounded by a fixed polynomial decay we will take the polynomial order as a convenient measure of the process smoothness. For the Matérn family a smoothness of ν and dimension 2 the spectral density will have a tail behavior following $r^{-(2\nu+2)}$ as $r \rightarrow \infty$. For example the exponential covariance ($\nu=1/2$) will have a spectral density that decreases at the polynomial rate r^{-3} . A covariance spectrum with tail behavior of the same order might be expected to provide a process model with similar smoothness to the exponential at small spatial scales. The following theorem reports the tail behavior for the multi-resolution process for different choices of the scale and weight sequences. An interesting result is that the multi-resolution process can reproduce a scale of different decay rates for the tail of the spectral density and can recover the -3 rate of decay for the exponential covariance.

Theorem 4.1. *Assume:*

1. ϕ is a two-dimensional Wendland covariance function of order K
2. The smoothness of the Matérn processes is fixed at $\nu = 1$.
3. $\alpha_l = e^{-2\beta_1 l}$ and $\theta_l = e^{-\beta_2 l}$ with $\beta_1, \beta_2 > 0$ and $(\beta_1/\beta_2 + 1) < (5 + 2K)$

If $S(r)$ denotes the spectral density of g (or k) with respect to the radial coordinate then there are constants independent of r , $0 < A_1, A_2 < \infty$ such that

$$A_1 < S(r)r^{2\mu+2} < A_2$$

with $\mu = \beta_1/\beta_2$.

Corollary 4.2. *Under assumptions 1 and 2 and $\theta_l = 2^{-l}$, $\alpha_l = \theta_l^{2\nu}$ and $(\nu + 1) < (5 + 2K)$, $S(r)$ will have tail behavior with the same polynomial order as a two-dimensional, Matérn process spectrum with smoothness ν .*


The proof of this theorem is given in the Appendix.

4.2 Numerical approximation

Given the theoretical foundation for the multi-resolution covariance it is interesting to explore the approximation for a finite number of levels and for the lattice process. As an example of this property, the setup from Theorem 1 was evaluated to approximate members of the Matérn family. Figure 2 gives the resulting approximations for 2, 4, and 6 multi-resolution levels (k_l , for l being 2, 4, and 6) and for smoothness $\mu = 0.5$ (exponential) and 1.0 (Whittle). The multi-resolution and the Matérn correlation functions are isotropic and so they are plotted against the distance of separation. The weight formulae are taken from the Corollary to Theorem 1. The Matérn covariance used for the convolution has smoothness 1.0 with a scale of 1.0 and the Wendland function, ϕ , has smoothness degree 2.0 with a scale of 2.5 (i.e. identically zero outside a radius of 2.5). The Matérn covariance function for comparison was determined by finding the scale parameter that minimizes the sum of squares (i.e. a discrete L_2 norm) between the weighted mixture given by the multi-resolution model and the Matérn family. Although one might derive this value from the asymptotic formula in practice this would be estimated and so it seems reasonable to choose a scale as the result of an optimization of the norm. This gave 0.159 for $\mu = .5$ and .131 for $\mu = 1.0$.

The approximation improves with increasing level and the benefit from more levels is most apparent near the origin. Moreover the approximation can be understood by the support of the basis functions from the first row of plots. Here we see that finest level of the multi-resolution correlation function has a half height on the order of .05 and this distance is about the order where the six level approximation breaks down for the exponential. In general we have found the approximation to be accurate provided enough levels of resolution are included. As a rule of thumb 6 levels seems to be a practical upper limit to represent members of the Matérn family.

Exponential $\mu = .5$



Whittle $\mu = 1.0$




Figure 2: Theoretical approximation of Matérn correlation functions using the multi resolution process based on process convolutions. The first row plots the approximations to an exponential function (black solid). The grey lines are the approximations for 2 (dotted), 4 (dashed) and 6 (solid) levels of the multi-resolution. On the left side is the full range of this case and the right is a magnification near the origin and with a log scale on the y axis. The second row is organized the same way as the second row but in this case the approximation is to a Whittle covariance function (a Matérn covariance with a smoothness of 1.0).

The theoretical approximation is based on a continuous convolution of the basis functions with the Matérn covariance and does not match the discrete stochastic model used for data analysis. A more practical comparison is how well a multi-resolution

basis can match members of the Matérn family. Here we indicate the quality of the approximation given $\theta_l = 2^{-l}$ but optimizing over $\{\kappa_l\}$ and $\{\alpha_l\}$. Note that this scheme is slightly different than the theoretical setup because κ_l is allowed to vary independently from θ_l and α_l is not constrained to be a power of θ_l . The first row in Figure 3 shows the approximation for an exponential covariance with range parameters .1, .5 and 1.0. using 3 and 4 levels of multi-resolution basis functions. The multi-resolution parameters κ_l and α_l have been found by minimizing the mean squared error between the approximation and the target covariance function on a grid of 200 distances in the interval $[0, 1]$. The coarsest basis function centers are organized on a 10×10 grid on the square $[-1, 1] \times [-1, 1]$ and so with four levels the approximation has $10^2 + 19^2 + 37^2 + 73^2 = 7159$ two-dimensional basis functions. The plots in the left hand column are the target and approximate covariances as a function of distance from the point $(0, 0)$ along the x-axis. The approximation is close to being stationary and isotropic and so this comparison is representative for distances along other orientations. In the plots the solid curve is the covariance, the dotted line is the approximation with 3 levels, and the dashed line is the approximation at 4 levels.

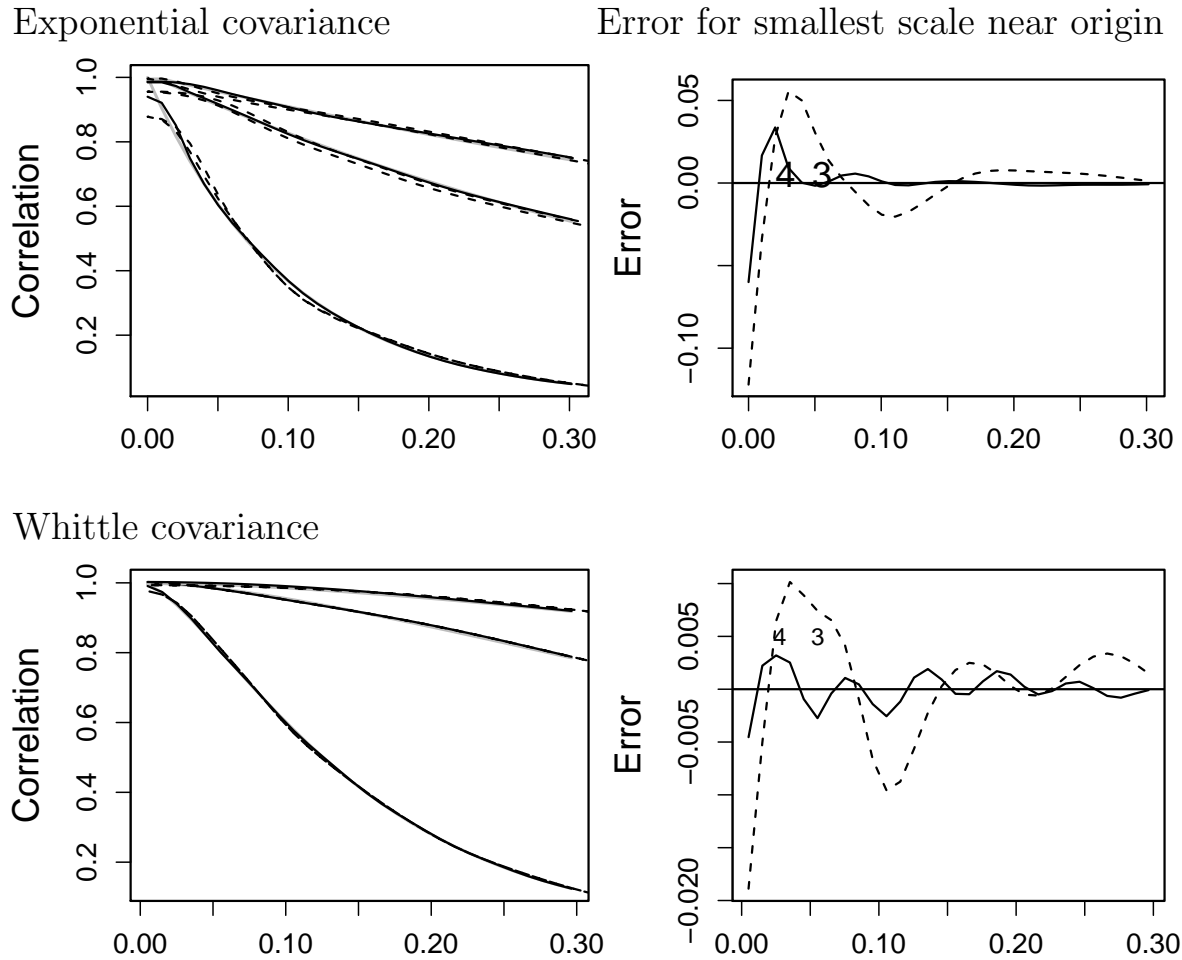


Figure 3: Approximation of Matérn covariances using the lattice/basis model. For the plots on the left hand side columns the solid grey lines are the true correlation functions. First row is an exponential correlation with range parameter (.1, .5 and 1.0) and second row is the Whittle correlation with ranges 0.1, 0.5 and 1.0. Black lines are the approximations to these correlation functions. Approximations are indicated in black and are based on a three level (dashed) or four level (solid) multi-resolution model. The domain for the approximation is a $[-1, 1] \times [-1, 1]$ a square and the initial grid is 10×10 for the range parameters .1, .5. The initial grid for the 1.0 ranges is 5×5 . The left columns are the approximations with the true correlations over the distance limits $[0, .3]$ to highlight errors near the origin even though the approximation is found for a spatial domain $[-1, 1] \times [-1, 1]$. The right columns are the differences between the approximation and the true correlation function for the cases when the range is .1 or for the mixture model. The characters 3 and 4 indicate the support for the basis functions at the third and fourth levels of resolution.

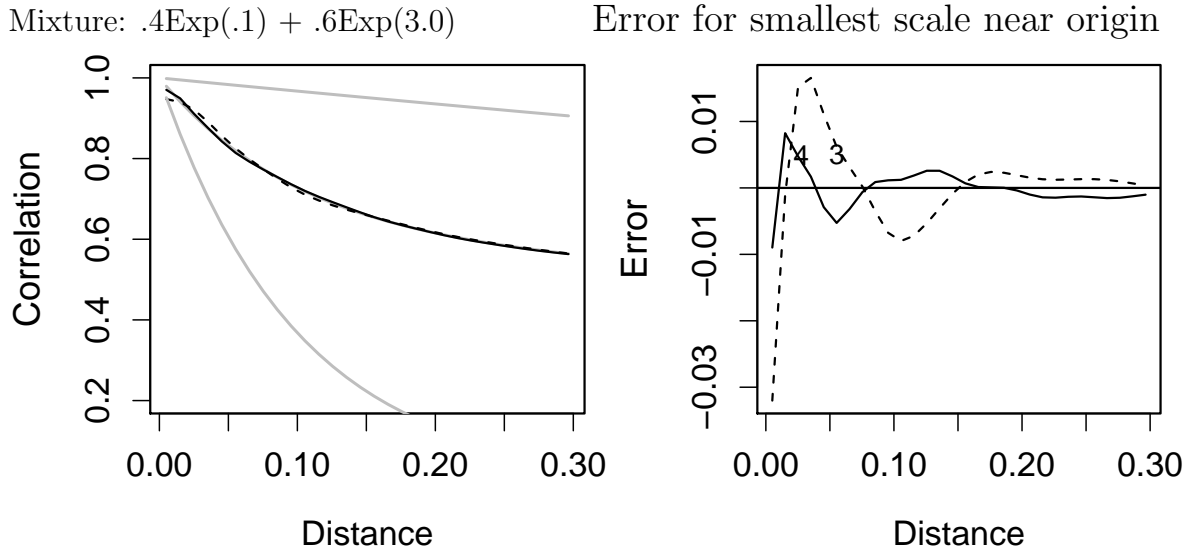


Figure 4: Approximation of a mixture of exponential covariances. The figure is organized the same as a row in the previous but with addition grey lines in the left plot indicating the two separate exponential covariances comprising the mixture.

Not surprisingly the approximation breaks down at small distances that are below the resolution of the finest basis functions. This feature is highlighted by the plots in the right column where the approximation is given for points in a range close to zero. The characters “3” and “4” indicate the smallest scale of the basis functions and thus indicate the limits of the multi-resolution for the 3 and 4 level choices. A similar approximation is made for the Whittle covariance ($\mu = 1$) except for the largest range parameter the coarsest basis has centers on a 5×5 grid (giving a total of 1484 basis functions). Note that in the error plot there is also a small artifact, a rippling feature that is from the discrete spacing of the basis functions. Figure 4 is an example of the ability of the multi-resolution to approximate more general correlation functions. This is perhaps the most striking example of the flexibility of this model. Here the target is a mixture of exponentials: $.4 \exp(-d/.1) + .6 \exp(-d/3)$. For reference the individual exponential correlation functions are plotted as grey solid lines. The approximation is also accurate with the error localized near the origin and being large below the smallest scale of the multi-resolution.

5 North American summer precipitation

The multi-resolution lattice model was applied to a substantive climate data set in order to test its practical value and compare it to standard Kriging. The goal is to estimate the average summer rainfall on a fine grid for North America based on high quality surface observations. These types of fields are an important reference in studying the Earth's climate system. These observational climate data can be used to assess the validity of global numerical models that attempt to simulate climate. Typically one requires that irregularly spaced station data be transformed to a regular grid that is comparable to numerical climate model output. Some standard gridded climate datasets are those from the Hadley Center (Jones et al. 2012) and the climate group at the University of Delaware (Willmott and Robeson 1995). However, we suggest that these data products can be improved using spatial methods to quantify the uncertainty and also to facilitate the use of orography (e.g. elevation) for spatial prediction.

The monthly data used to construct this example are freely available as the product Global Historical Climate Network (GHCN) Version 2 Precipitation (NOAA/NCDC 2011). GHCN data is quality controlled, curated and served by the US National Climatic Data Center. The *adjusted* data from this archive has been modified from its raw form to make the record more homogenous. Heterogeneities can come from a variety of sources such as moving the station a short distance or changes in instruments (see <http://www.ncdc.noaa.gov/ghcnm>). The data used here consists of 1720 stations. 1595 stations from the adjusted subset and 125 unadjusted stations in Northern Canada, Alaska and Mexico to fill out the edges of the study domain. For each station, a least squares, straight line was fit to the summer precipitation totals (June, July, August) for the period 1950-2010 and the trend line was evaluated at the midpoint time (1980.5). The trend line for each station at this time is taken to represent the mean climate for this reference period. Note that with complete observations this is just the sample mean however, 75% of the adjusted stations are missing at least 10 values in this period.

For reference the version of the climate data used is the R data set `NorthAmericanRainfall` in the `LatticeKrig` package. (Much of the basic analysis is included as source code

in the help file for this data set.) Elevation is based on a gridded field at 4km resolution derived from GTOPO30, a global digital elevation model created by the U.S. Geological Survey (see <http://eros.usgs.gov>). The spatial model was fit using stereographic map coordinates for the station locations (the option `stereographic` from the R `mapproj` package). This projection gave spatial coordinates whose euclidean distances were similar to great circle distance (see Figure 5). Thus this transformation of coordinates avoided specializing the basis for the sphere. With larger regions this will not be effective and one would have to construct a different set of basis functions and a GMRF better suited to spherical geometry.

The spatial model was fit to the log of mean precipitation with the spatial coordinates and elevation included as linear fixed effects. Three correlation models were considered:

Matern (2 parameters) A stationary, isotropic Matern with range and smoothness parameters.

Matern-like (2 parameters) A three-level, multi-resolution covariance with coarsest level having a lattice of 16×13 within the rectangular spatial domain amounting to approximately 4000 basis functions. A common value for κ was used to control the range at all levels. The first multi-resolution model constrains $\{\alpha_1, \dots, \alpha_3\}$ $\alpha_k \sim 2^{-2\nu}$ with the additional constraint that $\sum \alpha_k = 1$.

Multi-resolution (3 parameters) The same three-level structure as the Matern-like model with κ a common parameter. However, α_k are only constrained as $\alpha_k > 0$ and $\sum \alpha_k = 1$.

All three covariance functions include the variance parameter, ρ being the marginal variance of the spatial process and the parameter, σ^2 that is the measurement (or nugget) variance.

The covariance parameters were estimated by maximum likelihood and confidence regions for the parameters were derived using the large sample chi squared approximation to -2 times the log likelihood. Based on a 95 % confidence set the range parameter for the Matern model was not constrained from above and so a thin-plate spline model, i.e. a limiting process as the range becomes large, is not ruled out. The smoothness

parameter however has an MLE of .64. Figure 6 compares the correlation functions for these three different models based on the confidence sets for the parameters. Here the 95% confidence set for the models parameters was translated into a confidence band of the corresponding correlation functions. The multi-resolution models have the flexibility to have long range correlations and it is interesting for these data that their shape is different than the Matern family. The spatial predictions given by all three models are similar, within the prediction uncertainty measures. The estimated measurement error (σ^2) and effective degrees of freedom for the spatial models are reported in Table 1. The measurement error variance is smaller for the Matern compared to the lattice models and this is consistent with the Matern representing a slightly rougher process than the multi-resolution models and so more of the fine scale variance is captured by the process.

Figure 7 is an example of the expected precipitation surface for a subregion over the Rocky Mountains centered on Colorado. The multi-resolution covariance with the MLE parameters reported in Table 1 is used for these estimates, which are evaluated on a 200×200 grid. 200 conditional fields were simulated and to increase the accuracy of this sample the realizations were centered so that their mean matched their conditional expected mean, which can be computed exactly. Although the spatial model was estimated on a log scale of precipitation, the conditional samples were transformed to the raw scale to represent the distribution for unlogged values. Specifically the surface in (a) is the mean of the exponentiated conditional fields. Here the elevation covariate explains a large amount of the spatial structure but this component is modified by the smooth nonparametric component based on the location. Plot b) is the estimated prediction standard error as a percentage of the mean predicted field.

	MLE σ	Model degrees of freedom
Matern	.1084	943
Matern-like	.1402	489.4
Multi-resolution	.1353	550.6

Table 1: Maximum likelihood estimates of σ and the effective degrees of freedom for the three covariance models described in Section 5.

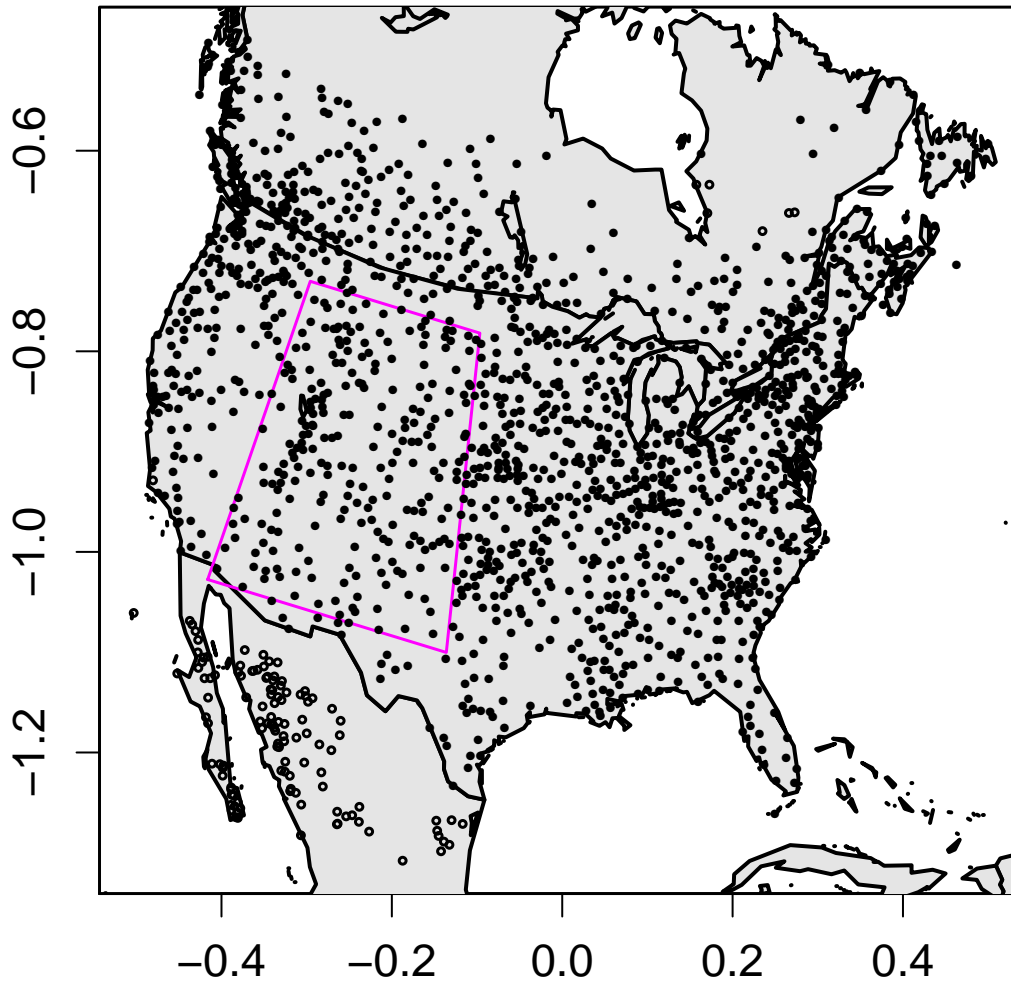


Figure 5: Stereographic projection of precipitation observation locations. Solid points show observation stations where the records have been adjusted and open dots those stations with unadjusted records.




Figure 6: Correlation models fit to the precipitation data. Dashed line is the Matérn correlation function found by maximum likelihood and the light grey shading is an approximate 95% uncertainty region based on a confidence set for the range and smoothness parameters. Note that the range is not constrained in its upper limit. Dotted red line is the estimated correlation and uncertainty (red shading) for the Matern-like covariance model. Solid green line with darker grey shading is a similar summary of the correlation for the three level multi-resolution model. Note that distances are in stereographic map coordinates and their scale can be inferred from Figure 5.

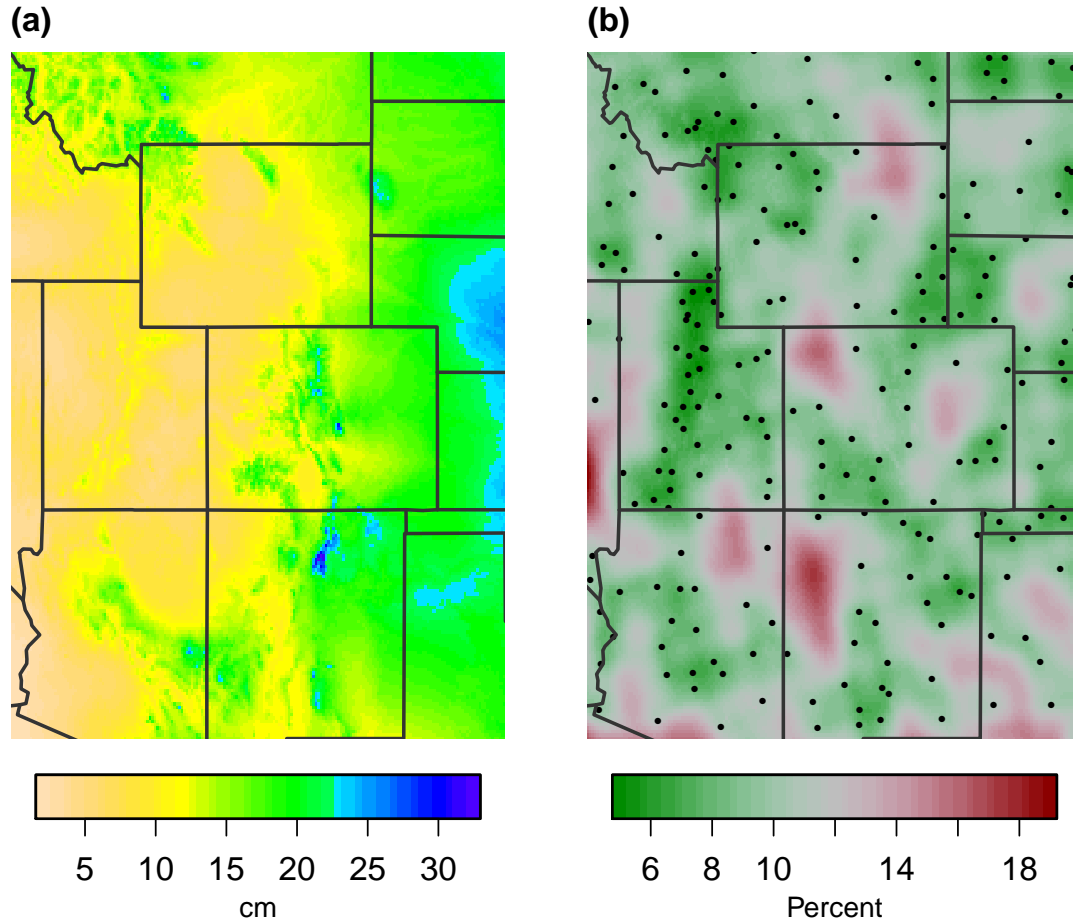


Figure 7: Plot (a) reports the spatial predictions for mean summer (JJA) precipitation in cm and includes elevation as a fixed linear covariate. This subregion is outlined in Figure 5 and units are centimeters of total rainfall. The spatial covariance function is the three level multi-resolution model described in Figure 6. (b) reports approximate prediction standard errors for this surface as a percentage of the predicted mean field. Solid points show observation stations.

6 Discussions and Conclusions

This work has developed a new model for a spatial process: a lattice/basis model that builds on ideas from fixed rank Kriging and the computational efficiencies that are inherited from Markov random fields. The key contribution is that an independent sum of the processes at different scales can approximate a larger family of processes

not limited to the properties of the covariance at each resolution level. One advantage of our model is numerical evidence that it can accurately reproduce the Matérn family of covariances. Also we give some asymptotic results based on a theoretical convolution model that indicate that a range of smoothness properties can be achieved. This result is unexpected given that the lattice/basis process has a fixed smoothness controlled by the choice of basis functions. It is still an open task to determine the precise connection between the convolution model and the discrete sum suggested in Section 4. To this extent, the theoretical results are at best suggestive of the lattice/basis model properties. We have found it advantageous to normalize the basis functions to give constant marginal variances for a given GMRF, which comes at some computational costs. The obvious advantage here is that, albeit achieved in a fairly heavy-handed way, the discrete model and the continuous convolution model agree exactly in terms of the process marginal variance.

Besides the value of the lattice/basis formulation as a new covariance model there is an equally important contribution in computational efficiency for large data sets. In fact it is our perspective that more complex covariance models can only be exploited when a large number of observation locations allow for accurate estimation of covariance parameters. Thus the computational aspects are intrinsic to applying new spatial models. We have been successful in identifying algorithms that allow for computing the restricted likelihood to estimate covariance parameters and the prediction of the spatial field using very large data sets. In addition, conditional simulation of the spatial fields is a simple and computationally feasible method to incorporate uncertainty from both the covariance parameters and also the spatial interpolation or smoothing. Although we have not developed a fully Bayesian algorithm for inference we emphasize that the main computational ingredients are present to apply integrated nested Laplace approximations (Rue et al. 2009) or to generate complete Monte Carlo samples from the full posterior distribution.

Because of the description of the stochastic spatial elements in terms of a SAR, it is straightforward to propose a non stationary extension to the lattice basis model. One would allow both the κ_l and α_l to vary over the lattice at each level. An additional

refinement would allow the SAR weights between the neighboring lattice points to be directionally dependent. The spatial variation in these parameters could be modeled by a set of covariates and fixed effects or one could include a spatial process prior on these parameter fields. The advantage of our approach and also of the related SPDE and process convolution models is that one will always obtain a valid covariance function because the model focuses on a process level description. It is still an open question how much non-stationary structure can be estimated from data sets currently encountered in geophysical applications. However, we believe that the lattice/basis formulation will provide competitive computational efficiency and model parsimony compared to other approaches.

Finally, we note that the lattice/basis model can be implemented using a collection of simple numerical algorithms and readily available software. We avoid the triangulation of the spatial locations or computation of basis function inner products and the sequence of regular lattices for the basis function centers simplifies coding the algorithms. An R implementation is available with amply documented and commented source code and uses the generic sparse matrix R package `spam`. The `LatticeKrig` source code is largely written in the R language with limited use of lower level C or FORTRAN functions and is hence easy to modify. The maximization of the likelihood (or more generally a posterior) can be done using packages and tools in R which makes this model easy to apply for spatial data analysis and inference.

References

- Caragea, P. C. and R. L. Smith (2007), “Asymptotic properties of computationally efficient alternative estimators for a class of multivariate normal models.” *Journal of Multivariate Analysis*, 98, 1417–1440.
- Cressie, Noel A. C. and Gardar Johannesson (2008), “Fixed rank kriging for very large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70, 209–226. Additional Info: Blackwell Publishing; 20080201; DOI: 10.1111/j.1467-9868.2007.00633.x.

- Eidsvik, Jo, Andrew O Finley, Sudipto Banerjee, and Havard Rue (2010), “Approximate Bayesian inference for large spatial datasets using predictive process models.” *Computational Statistics & Data Analysis*, 56, 1362–1380.
- Fuentes, Montserrat (2007), “Approximate likelihood for large irregularly spaced spatial data.” *Journal of the American Statistical Association*, 102, 321–331, URL <http://dx.doi.org/10.1198/016214506000000852>. M3: doi:10.1198/016214506000000852.
- Furrer, Reinhard, Douglas Nychka, and Stephen Sain (2012), *fields: Tools for spatial data*. URL <http://www.image.ucar.edu/Software/Fields>. R package version 6.6.4.
- Furrer, Reinhard and Stephan R. Sain (2010), “spam: A sparse matrix r package with emphasis on mcmc methods for gaussian markov random fields.” *Journal of Statistical Software*, 36, 1–25, URL <http://www.jstatsoft.org/v36/i10>.
- Henderson, H.V. and S. R. Searle (1981), “On deriving the inverse of a sum of matrices.” *SIAM Review*, 23, 53–60.
- Higdon, David M. (1998), “A process-convolution approach to modelling temperatures in the North Atlantic Ocean.” *Environmental and Ecological Statistics*, 5, 173–190, URL <http://dx.doi.org/10.1023/A:1009666805688>. M3: 10.1023/A:1009666805688.
- Jones, P. D., D.H. Lister, T.J. Osborn, C. Harpham, M. Salmon, and C.P. Morice (2012), “Hemispheric and large-scale land surface air temperature variations: an extensive revision and an update to 2010.” *Journal of Geophysical Research*.
- Katzfuss, Matthias and Noel Cressie (2011), “Spatio-temporal smoothing and EM estimation for massive remote-sensing data sets.” *Journal of Time Series Analysis*, 32, 430–446.
- Lindgren, Finn and Håvard Rue (2007), “Explicit construction of gmrf approximations

to generalized matern fields on irregular grids.” Technical report, Lund Institute of Technology.

Lindgren, Finn, Håvard Rue, and Johan Lindström (2011), “An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73, 423–498.

Michael L. Stein, Leah J. Welty, Zhiyi Chi (2004), “Approximating likelihoods for large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66, 275–296.

NOAA/NCDC (2011). URL <http://www.ncdc.noaa.gov/ghcnm>.

Nychka, Douglas, Dorit Hammerling, Stephen Sain, and Tia Lerud (2012), *LatticeKrig: Multiresolution Kriging based on Markov random fields*. URL <http://www.image.ucar.edu/Software/MRKriging>. R package version 2.3.

R Development Core Team (2011), *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, URL <http://www.R-project.org/>. ISBN 3-900051-07-0.

Rue, H., S. Martino, and N. Chopin (2009), “Approximate Bayesian inference for latent gaussian models by using integrated nested Laplace approximations.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71, 319–392.

Rue, Håvard and Leonhard Held (2005), *Gaussian Markov random fields : theory and applications*, volume 104. Chapman & Hall/CRC, Boca Raton. Håvard Rue, Leonhard Held.; : ill. ; 24 cm; Includes bibliographical references (p. 237-253) and indexes.

Sang, H. and J.Z. Huang (2011), “A full scale approximation of covariance functions for large spatial data sets.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74, 111–132.

Stein, Michael L. (2008), “A modeling approach for large spatial datasets.” *Journal of the Korean Statistical Society*, 37, 3–10.

Wendland, H. (1998), “Error estimates for interpolation by compactly supported radial basis functions of minimal degree.” *Journal of Approximation Theory*, 93, 258–272.

Wendland, Holger (1995), “Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree.” *AICM*, 4, 389–396.

Willmott, C.J. and S.M. Robeson (1995), “Climatologically aided interpolation (CAI) of terrestrial air temperature.” *International Journal of Climatology*, 15, 221–229.

Acknowledgements

This work supported in part by National Science Foundation grant DMS-0707069 and the National Center for Atmospheric Research. S. Bandyopadhyay is partially supported by the Reidler Foundation of Lehigh University.

Appendix

Outline of proof

Let $\tilde{\phi}_k$ be the spectral density for ϕ and \tilde{C}_ν the spectral density of a Matérn field with $\nu = 1$, unit variance and unit spatial scale parameter. Including the scale parameter for the radial basis function kernel and using elementary properties of convolution.

$$\tilde{S}(r) = \sum_{l=1}^{\infty} \alpha_l \left[\theta_l^2 \tilde{C}_\nu(\theta_l r) \right] \left[\tilde{\phi}_k(\theta_l r) \right]^2$$

The Matérn spectral density is

$$\tilde{C}_\nu(r) = \frac{1}{(2\pi)} \frac{1}{(1+r^2)^2}$$

For the Wendland spectral density there are constants C_1 and C_2 depending only on K such that for all ω

$$C_1 \leq \tilde{\phi}_k(\omega)(1 + \|\omega\|^2)^{3/2+K} \leq C_2$$

(Wendland (1998)). Using the upper bound on $\tilde{\phi}$, substituting the expressions for θ_l and α_l and finally combining terms gives the upper bound

$$\tilde{S}(r) < C' \sum_{l=1}^{\infty} \alpha_l \frac{\theta_l^2}{(1 + (r\theta_l)^2)^\eta} = C' \sum_{l=1}^{\infty} \frac{e^{-2\beta_1 l} e^{-2\beta_2 l}}{(1 + (r e^{-\beta_2 l})^2)^\eta} = C' \sum_{l=1}^{\infty} \frac{e^{-(2\beta_1 + 2\beta_2)l}}{(1 + r^2 e^{-2\beta_2 l})^\eta}$$

with $\eta = 2 + 2(3/2 + K) = 5 + 2K$.

Now apply the useful lemma given below with the identifications $a = 2\beta_1 + 2\beta_2$, $b = 2\beta_2$, $c = \eta$, and $s = r^2$. We have the rate given by $r^{-2(a/b)}$ and with $2a/b = 2(2\beta_1 + 2\beta_2)/2\beta_2 = 2\beta_2/\beta_1 + 2$. The result for the upper bound now follows and the rate for the lower bound is proved in a similar manner.

Two Useful Lemmas

Lemma 6.1. *Let H be a continuous and integrable function on $[1, \infty]$. Also assume that H is positive and unimodal with maximum at u^* .*

$$\left| \sum_{l=1}^{\infty} H(l) - \int_1^{\infty} H(u) du \right| < H(u^*)$$

Proof Let L be the integer so that $H(L) = \max_l H(l)$ also let $I_l = \int_l^{l+1} H(u) du$ then by elementary properties of the integral and the unimodality of H

$$\begin{aligned} I_l &> H(l), & 1 \leq l \leq (L-1) \\ I_{l-1} &> H(l), & (L+1) \leq l \leq \infty \end{aligned} \tag{17}$$

summing over l gives

$$\sum_{l=1}^{\infty} I_l > \sum_{l \neq L} H(l)$$

Simplifying and rearranging terms

$$\int_1^{\infty} H(u) du - \sum_{l=1}^{\infty} H(l) > -H(L)$$

Again by properties of the integral and H

$$\begin{aligned} I_{l-1} &< H(l), & 2 \leq l \leq L \\ I_l &< H(l), & (L+1) \leq l \leq \infty \end{aligned} \tag{18}$$

summing over l gives

$$\int_1^\infty H(u)du < \sum_{l \neq L} H(l)$$

or

$$\int_1^\infty H(u)du - \sum_{l=1}^\infty H(l) < H(L)$$

Noting that $H(L) < H(u^*)$ the result now follows.

Lemma 6.2. For $a, b, c, s > 0$ and for $(a/b) - c < 0$ there are constants $0 < C_1, C_2 < \infty$

$$C_1 s^{-a/b} < \sum_{l=1}^\infty \frac{e^{-al}}{(1 + se^{-bl})^c} < C_2 s^{-a/b}$$

Proof Based on Lemma 6.1 let $H(u) = e^{-au}/(1 + se^{-bu})^c$. H is unimodal. From basic calculus the maximum of H is $H(u^*) = Cs^{-a/b}$ for $0 < C < \infty$ and C depending only on a, b, c . We now evaluate the approximating integral from Lemma 6.1 as a function of s .

$$\int_1^\infty H(u)du = \int_1^\infty \frac{e^{-au} du}{(1 + se^{-bu})^c} = \int_1^\infty \frac{(e^{-bu})^{a/b} du}{(1 + se^{-bu})^c}$$

Now make the substitution $q = e^{-bu}$ giving $dq = -b(e^{bu})du$ or $du = \frac{-dq}{bq}$ and with limits of integration, e^{-b} and 0. One obtains

$$b \int_0^{e^{-b}} \frac{q^{(a/b)-1} dq}{(1 + sq)^c} \tag{19}$$

Since $(a/b) > 0$ the pole at zero is integrable and the integral is finite. Now make the substitution $p = sq$ giving $dp = sdq$ and

$$b \int_0^{se^{-b}} \frac{(p/s)^{(a/b)-1} dp}{s(1 + p)^c} = bs^{-a/b} \int_0^{se^{-b}} \frac{p^{(a/b)-1} dp}{(1 + p)^c} \tag{20}$$

Under the assumption that $a/b - c < 0$ the integral will be finite in the limit as $s \rightarrow \infty$. Thus $\int_1^\infty H(u)du$ and $H(u^*)$ converge to zero at the polynomial rate $s^{-a/b}$ and the result follows from application of Lemma 6.1.