

# Large Multi-scale Spatial Modeling Using Tree Shrinkage Priors

Rajarshi Guhaniyogi and Bruno Sansó

June 22, 2017

## Abstract

We develop a multiscale spatial kernel convolution technique with higher order functions to capture fine scale local features and lower order terms to capture large scale features. To achieve parsimony, the coefficients in the multiscale kernel convolution model is assigned a new class of “tree shrinkage prior” distributions. Tree shrinkage priors exert increasing shrinkage on the coefficients as resolution grows so as to adapt to the necessary degree of resolution at any sub-domain. Our proposed model has a number of significant features over the existing multi-scale spatial models for big data. In contrast to the existing multiscale approaches, the proposed approach auto-tunes the degree of resolution necessary to model a subregion in the domain, achieves scalability by suitable parallelization of local updating of parameters and is buttressed by theoretical support. Excellent empirical performances are illustrated using several simulation experiments and a geostatistical analysis of the sea surface temperature data from the Pacific ocean.

## 1 Introduction

Ubiquity of spatially indexed datasets in various disciplines (Gelfand et al., 2010; Cressie and Wikle, 2015; Banerjee et al., 2014) has motivated researchers to develop a variety of methods and models in spatial statistics. Gaussian processes offer a rich modeling framework and are being widely deployed to help researchers comprehend complex spatial phenomena. However, Gaussian process likelihood computations involve matrix factorizations (e.g., Cholesky) and determinant computations for large spatial covariance matrices that have no computationally exploitable structure. This incurs onerous computational burden and is referred to as the “Big-N” problem in spatial statistics.

There are, broadly speaking, two different premises for modeling large spatial datasets. One of them is “sparsity”, while the other is “dimension-reduction”. Sparse methods include covariance tapering (see, e.g., Furrer et al. (2012); Kaufman et al. (2008); Du et al. (2009); Shaby and Ruppert (2012)), which introduces sparsity in the Gaussian covariance matrix using compactly supported covariance functions. This is effective for fast parameter estimation and interpolation of the response (“kriging”), but is less suited for more general inference on residual or latent processes due to exorbitantly expensive determinant computation of the sparse covariance matrix. An alternative approach introduces sparsity in

the inverse of covariance (precision) matrix using conditional independence assumptions or composite likelihoods (e.g., Vecchia (1988); Rue et al. (2009); Stein et al. (2004); Eidsvik et al. (2014); Datta et al. (2015); Guinness (2016)). In related literature pertaining to computer experiments, localized approximations of Gaussian process models are proposed, see for e.g. Gramacy and Apley (2015), Zhang et al. (2016) and Park and Apley (2017). This literature is overwhelmingly frequentist, less model based and has different goals compared to the Bayesian spatial literature.

Dimension-reduction methods subsume the popular “low-rank” models which express the realizations of the Gaussian process as a linear combination of  $r$  basis functions (see, e.g., Higdon (2002); Stein (2007); Banerjee et al. (2008); Cressie and Johannesson (2008); Crainiceanu et al. (2008); Finley et al. (2009); Lemos and Sansó (2009)), where  $r \ll n$ . This leads to a flexible class of models, but otherwise possess a low-rank structure that enables likelihood evaluation by solving  $r \times r$  linear systems instead of  $n \times n$ . The algorithmic cost for model fitting decreases from  $O(n^3)$  to  $O(nr^2 + r^3)$ . However, when  $n$  is large, empirical investigations suggest that  $r$  must be fairly large to adequately approximate the parent process so that  $nr^2$  flops becomes exorbitant. Furthermore, low rank models perform poorly when neighboring observations are strongly correlated and the spatial signal dominates the noise (Stein, 2014). Improvements on low-rank models with properly designed basis functions (e.g. Guhaniyogi et al. (2011)) have appeared in the recent past. However, these improvements detract from the computational advantages.

Some variants of dimension-reduction methods partition the large spatial data into subsets containing fewer observations, run Gaussian processes in different subsets followed by combining inference from subsets, see e.g. Gramacy and Lee (2012), Guhaniyogi and Banerjee (2017). Another important class of models aimed at modeling the spatial surface at multiple scales; finer variations at the local scale and overall trend at the global scale. Most of the geophysical processes naturally tend to have a multiscale character over space that requires statistical methods to allow for potentially complicated multiscale spatial dependence beyond a simple parametric model. However, literature on Bayesian multiscale spatial models for big data is quite insufficient.

Our approach combines the representation of a random field using a multiresolution basis with coefficients modeled using a newly developed *multiscale tree shrinkage prior*. To be more precise, the spatial surface  $w(\mathbf{s})$  is viewed as the sum of  $R$  independent processes  $w(\mathbf{s}) = \sum_{r=1}^R w_r(\mathbf{s})$ , the  $r$ -th process  $w_r(\mathbf{s})$  corresponds to the  $r$ -th resolution, and each  $w_r(\mathbf{s})$  being modeled using a discrete kernel convolution approach,

$$w_r(\mathbf{s}) = \sum_{j=1}^{J(r)} \kappa_{j,r}(\mathbf{s}) \beta_j^r, \quad (1)$$

where  $\{\kappa_{j,r}(\mathbf{s})\}_{j=1}^{J(r)}$ ,  $r \in \{1, \dots, R\}$ , is a set of basis of functions for the  $r$ th resolution and  $\beta_j^r$ 's are corresponding coefficients. We use families of radial basis functions of minimal order (see Section 2) kept on a regular grid with increasing resolutions. These radial basis functions have compact support that facilitates significant computational gain.  $\beta_j^r$ 's play an important role in determining whether a sub-domain requires modeling at the finest resolution or at coarse resolutions. A new class of *multiscale tree shrinkage prior* is developed to adequately

model  $\beta_j^r$ 's in various sub-domains at different resolutions.

It is noteworthy that a few other important articles on multiscale spatial models for big data have already appeared in the literature, see e.g. Nychka et al. (2015); Katzfuss (2016, 2013) and references therein. We derive a number of very important advantages over the existing literature. Firstly, the most important difference of our approach with `LatticeKrig` model (Nychka et al., 2015) is that the newly proposed multiscale tree shrinkage prior is equipped to impart increasing shrinkage on basis coefficients as resolution increases. This effectively leads to a continuous analogue to selecting the number of resolutions necessary for modeling a sub-domain. The idea of the tree shrinkage prior is novel in its own right with possible applications anticipated in statistical genomics and neuroscience, for example identifying main effects versus interaction effects in genetic studies. Secondly, unlike `LatticeKrig`, the proposed multiscale approach incorporates data dependent choice of the kernel width. As a result, similar performance in terms of surface interpolation is available from these models with the proposed multiscale approach using much less number of knots. Thirdly, the proposed multiscale structure can be naturally embedded in a hierarchical structure to model non-Gaussian data. Fourthly, this article characterizes the function space of the fitted spatial surface and show asymptotic result on consistency of the posterior distribution of the same. Finally, judicious choice of the compact basis functions and the computational strategy described in Section 3.2 evoke extremely rapid Bayesian estimation that only involves inverting a large number of small matrices in parallel.

The remainder of the manuscript evolves as follows. Section 2 outlines the multiscale kernel convolution model development including the choice of knots, basis functions, basis coefficients and priors on them. Section 3 discusses posterior computation strategies and computation complexities. Theoretical insights on asymptotic properties of the posterior distribution of spatial surface is offered in Section 4. Detailed simulation studies are shown in Section 5. Section 6 details out analysis of a massive sea surface temperature data in Pacific ocean. Finally, Section 7 discusses what the newly developed multiscale model achieves, and proposes a number of future directions to explore.

## 2 Multiscale Spatial Kriging

### 2.1 Kernel convolutions as approximations to Gaussian processes

Let  $\{w(\mathbf{s}) : \mathbf{s} \in \mathcal{D}\}$  be a spatial field of interest in the continuous domain  $\mathcal{D} \subseteq \mathbb{R}^d$ ,  $d \in \mathbb{N}^+$ . We assume the true spatial process  $w(\mathbf{s})$  follows a Gaussian process. One may construct a Gaussian process  $w(\mathbf{s})$  over  $\mathcal{D}$  by convolving a continuous white noise process  $u(\mathbf{s})$ ,  $\mathbf{s} \in \mathbb{D}$  with a smoothing kernel  $K(\mathbf{s}, \phi)$  ( $\phi$  might be space varying) so that  $w(\mathbf{s}) = \int K(\mathbf{s} - \mathbf{z}, \phi)u(\mathbf{z})d\mathbf{z}$ , as proposed by Higdon (2002). The resulting covariance function for  $w(\mathbf{s})$  is fully determined by the kernel  $K(\cdot)$  such as

$$\text{cov}(w(\mathbf{s}), w(\mathbf{s}')) = \int K(\mathbf{s} - \mathbf{z}, \phi)K(\mathbf{s}' - \mathbf{z}, \phi)d\mathbf{z}. \quad (2)$$

A discrete approximation of (2) is obtained by sampling the convolved processes on a grid. Letting  $\mathbf{s}_1^*, \dots, \mathbf{s}_J^*$  be a set of knots in  $\mathcal{D}$ , a discrete approximation of  $w(\mathbf{s})$  is given by

$$\theta(\mathbf{s}) = \sum_{j=1}^J K(\mathbf{s} - \mathbf{s}_j^*, \phi) u_j, \quad (3)$$

where  $u_j$ 's are basis coefficients. The  $J$  knots are typically placed in a grid in  $\mathcal{D}$ , though other placements of knots have appeared in the literature. Varying the choice of the kernel functions and coefficients  $u_j$ , a rich variety of processes emerge from (3). Following Lemos and Sansó (2009), we term (3) as Discrete Process Convolutions (DCT). When  $J$  is small DCT provides computationally convenient approximation of the Gaussian process  $w(\mathbf{s})$ . However, Smaller  $J$  would greatly reduce approximation accuracy, while moderately large  $J$  exacerbates computational burden. The computational challenges cannot be solved by brute-force use of high-performance computing systems, and approximations or simplifying assumptions are necessary. One compelling idea to both reduce computation and increase approximation accuracy may come from using DCT at multiple scales with proper choice of kernel functions and basis coefficients. Next few sections carefully develop a multiscale-DCT model.

## 2.2 Partition of domain and choice of knots

To define the multiscale-DCT with resolution  $R$ , we iteratively partition  $\mathcal{D}$  up to level  $R$ . Let at the lowest level, one partitions  $\mathcal{D}$  into  $J(1)$  subsets  $\mathcal{D}_1, \dots, \mathcal{D}_{J(1)}$ . In the second level, each  $\mathcal{D}_i$  undergoes  $P$  partitions so that the total number of partitions in the second level is  $PJ(1)$ . Likewise, let in the  $(r-1)$ th level the set of partitions can be described as  $\{\mathcal{D}_{i_1, \dots, i_{r-1}} : i_1 \in \{1, 2, \dots, J(1)\}, i_2, \dots, i_{r-1} \in \{1, \dots, P\}\}$ . In the  $r$ th level each  $\mathcal{D}_{i_1, \dots, i_{r-1}}$  is partitioned into  $P$  subsets  $\mathcal{D}_{i_1, \dots, i_{r-1}, 1}, \dots, \mathcal{D}_{i_1, \dots, i_{r-1}, P}$ , so that

$$\mathcal{D}_{i_1, \dots, i_{r-1}} = \bigcup_{i_r=1}^P \mathcal{D}_{i_1, \dots, i_{r-1}, i_r}, \quad \mathcal{D}_{i_1, \dots, i_{r-1}, s} \cap \mathcal{D}_{i_1, \dots, i_{r-1}, s'} = \phi, \quad \forall s \neq s'. \quad (4)$$

Therefore, the number of partitions at the  $r$ th resolution is  $J(r) = P^{r-1}J(1)$ . In one dimensional ( $d = 1$ ) case,  $P = 2$ , i.e. bisection method is adopted to partition each subset for the next resolution. This naturally implies that the number of partitions at the  $r$ th level is  $J(r) = 2^{r-1}J(1)$ . In the two dimensional examples, any subset at a resolution is divided into 4 equal subsets, i.e.  $J(r) = 4^{r-1}J(1)$ . Partitioning of the domain can be envisioned as formation of a tree, with sub-domains  $\mathcal{D}_{i_1, \dots, i_r}$ 's as nodes of the tree. Lower and higher resolutions correspond to the upper and lower nodes of this tree.  $\mathcal{D}_1, \dots, \mathcal{D}_{J(1)}$  correspond to uppermost nodes of the tree.  $P$  branches emerge from each of these nodes leading to  $P^2$  nodes in the second level of the tree and this process continues. Indeed, for any  $i_1, \dots, i_r$ ,  $1 \leq r \leq R$ , we define *Subtree*( $\mathcal{D}_{i_1, \dots, i_r}$ ) by

$$\text{Subtree}(\mathcal{D}_{i_1, \dots, i_r}) = \{\mathcal{D}_{i_1, \dots, i_r}\} \cup_{j=1}^{R-r-1} \{\mathcal{D}_{i_1, \dots, i_r, i_{r+1}, \dots, i_{r+j}} : i_{r+1}, \dots, i_{r+j} \in \{1, \dots, P\}\} \cup \{\mathcal{D}_{i_1, \dots, i_R}\}. \quad (5)$$

$Subtree(\mathcal{D}_{i_1, \dots, i_r})$  consists of all sub-domains of  $\mathcal{D}_{i_1, \dots, i_r}$  in higher than  $r$ th resolution, including itself. Evidently,  $Subtree(\mathcal{D}_{i_1, \dots, i_R}) = \mathcal{D}_{i_1, \dots, i_R}$ . On a similar note, we also define the *father* node of  $\mathcal{D}_{i_1, \dots, i_r}$  as the node  $\mathcal{D}_{i_1, \dots, i_{r-1}}$ .


Defining multiscale-DCT also requires choosing a set of knot points at every level. We place knots in the center of every partition at a resolution. To be more precise, the knots  $\mathbf{s}_1^1, \dots, \mathbf{s}_{J(1)}^1$  in the first level is placed at the centers of  $\mathcal{D}_1, \dots, \mathcal{D}_{J(1)}$ . Likewise, knots  $\mathbf{s}_1^r, \dots, \mathbf{s}_{J(r)}^r$  are kept at the centers of the partitions at the  $r$ th level. Therefore, there is a one-one correspondence between the set of knots and the set of partitions of  $\mathcal{D}$ . Henceforth, we will interchangeably use *Subtree* and *Father* of a sub-domain with *Subtree* and *Father* of the knot that resides at the midpoint of that sub-domain, e.g. if  $\mathbf{s}_j^r \in \mathcal{D}_{i_1, \dots, i_r}$ , then  $Subtree(\mathbf{s}_j^r)$  and  $Father(\mathbf{s}_j^r)$  are synonymous with  $Subtree(\mathcal{D}_{i_1, \dots, i_r})$  and  $Father(\mathcal{D}_{i_1, \dots, i_r})$  respectively. It needs to be noted that the indexing set of knots is a bit different from the indexing set of partitions and they require reconciliation. The  $j$ th knot at the  $r$ th resolution  $\mathbf{s}_j^r$ ,  $j = 1, \dots, J(r)$ , belongs to the sub-domain  $\mathcal{D}_{i_1, \dots, i_r}$  if  $j = \sum_{l=1}^{r-1} (i_l - 1)P^{r-l} + i_r$ . With this notation,  $\mathbf{s}_k^{r-1}$  is the father node of  $\mathbf{s}_j^r$  iff  $k = \sum_{l=1}^{r-2} (i_l - 1)P^{r-l} + i_{r-1}$ , i.e.  $k = \lfloor \frac{j-1}{P} \rfloor + 1$ , where  $\lfloor x \rfloor$  is the greatest integer less than  $x$ .

To give examples of domain partitioning and knots, consider the one dimensional case with the spatial domain of interest  $[h_1, h_2]$ . In the first resolution, the domain is partitioned into  $J(1)$  intervals of equal length, i.e. each interval is of length  $\delta = (h_2 - h_1)/J(1)$ . The knots in the first resolution are placed at the midpoint of each interval so that the spacing between two successive knots is always  $\delta$ . At the second resolution, each interval is partitioned into equal length intervals with knots in resolution 2 placed right at the midpoints of these intervals. Therefore, the space between two successive knots is reduced to  $\delta/2$ . This process continues iteratively leading to the space between two successive knots at resolution  $r$  as  $\delta/2^{r-1}$ . For  $d = 2$ , let the domain of interest be  $[h_1, h_2] \times [h_3, h_4]$ . The first resolution divides the area into  $h_x \times h_y$  equi-dimensional rectangles with knots placed in the center of each rectangle. Therefore, the number of knots in the first resolution is  $J(1) = h_x \times h_y$ . In resolution 2, every rectangle in the first resolution is divided into four congruent rectangles. Knots in the second resolution are placed at the centers of these new rectangles. Clearly, the distance between two horizontally adjacent knots is  $(h_2 - h_1)/h_x$  and two vertically adjacent knots is  $(h_4 - h_3)/h_y$ , in the first resolution. At the  $r$ th resolution these distances decrease to  $2^{-r+1}(h_2 - h_1)/h_x$  and  $2^{-r+1}(h_4 - h_3)/h_y$  respectively.


Figure 1 shows the domain partitioning and the set of knots for both one dimensional and two dimensional applications. For the visual illustration, we restrict to  $R = 3$ ,  $h_1 = 0, h_2 = 10, h_3 = 0, h_4 = 10$ ,  $J(1) = 5$  in one dimension. For two dimensions, we restrict  $R = 2$ ,  $h_x = 4, h_y = 4$ ,  $J(1) = 16$  for better visualization. Hereon, we fix the template of domain partitions and the placement and number of knots in each partition. We offer some more discussion on them in the conclusion and future work section.

## 2.3 Multiscale spatial process with radial basis functions

We model the spatial effects by a multi-scale DCT with  $R$  resolutions,  $r$ th resolution being modeled by a DCT with kernel  $K(\cdot, \cdot, \phi_r)$ , knots  $\mathbf{s}_1^r, \dots, \mathbf{s}_{J(r)}^r$  and coefficients  $\beta_1^r, \dots, \beta_{J(r)}^r$ ,  $r = 1, \dots, R$ . To elaborate it further, the spatial surface  $w(\mathbf{s})$  is written as  $w(\mathbf{s}) = \sum_{r=1}^R w_r(\mathbf{s})$ ,



(a) knot placement: one dimension



(b) knot placement: resolution 1 (2d) (c) knot placement: resolution 2 (2d)

Figure 1: (a) placement of knots in one dimension for three resolutions. Knots in resolution 1 are shown in the upper level. Middle and lower level show knots in resolutions 2 and 3 respectively. For visualization in one dimension, we fix  $R = 3$ ,  $J(1) = 5$ . (b) Shows placement of knots in resolution 1 for two dimensions; (c) shows placement of knots in resolution 2 for two dimensions. For better visualization, we keep  $R = 2$ ,  $J(1) = 16$  in two dimensions.

such that

$$w_r(\mathbf{s}) = \sum_{j=1}^{J(r)} K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) \beta_j^r. \quad (6)$$

Contrary to the one scale DCT, multiscale DCT captures spatial variability at multiple scales. The lower resolutions capture variability at large distances while the finer local level variabilities are captured by higher resolutions. Intuitively, the implication is that one needs more basis functions in one scale DCT compared to multiscale DCT to draw similar inference. We will formally discuss this issue in simulation studies.

The choice of the kernel function  $K(\cdot, \cdot, \phi_r)$  is crucial for estimating the spatial variability at multiple scales. In the context of the ordinary one resolution kernel convolution literature, Lemos and Sansó (2009); Cressie and Johannesson (2008) proposed to use Bezier kernels

which are continuous but not differentiable. In the multiscale literature, Nychka et al. (2015) uses a Wendland kernel that is four times continuously differentiable. Let  $\kappa$  be a Wendland polynomial functions (Wendland, 2004), supported on  $[0,1]$ , having the form

$$\kappa(z) = (1 - z)_+^{l+1}(1 + (l + 1)z), \quad (7)$$

where  $(1 - z)_+ = (1 - z)$  if  $0 < z < 1$  and  $= 0$  otherwise and  $l = \lfloor d/2 \rfloor + 2$ . For our proposed approach, we choose kernel function  $K$  defined as

$$K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) = \kappa\left(\frac{\|\mathbf{s} - \mathbf{s}_j^r\|}{\phi_r}\right) = \left(1 - \frac{\|\mathbf{s} - \mathbf{s}_j^r\|}{\phi_r}\right)_+^{l+1} \left[1 + (l + 1)\frac{\|\mathbf{s} - \mathbf{s}_j^r\|}{\phi_r}\right]. \quad (8)$$

Geometrically, the kernel function consists of bumps centered at the node points with interpolation of the spatial surface at  $\mathbf{s}$  in the  $r$ th resolution is governed by knots located in  $B_{\phi_r}(\mathbf{s})$ , where  $B_\nu(\mathbf{s})$  is the Euclidean ball of radius  $\nu$  around  $\mathbf{s}$ . Section 3.2 describes computational advantages derived from the compact support of this kernel.

Note that  $\kappa$  is a Wendland polynomial function supported on  $[0, 1]$ . Wendland (2004) ensures that  $\kappa$  is positive definite which seems to be an attractive feature when this function is used for interpolation. Further,  $\kappa(z) \in C^2$  and it is the positive definite compactly supported polynomial of minimal degree for a given dimension  $d$  that possesses continuous derivatives up to second order. Theorem 2.1 characterizes the space of functions of the form  $w_r(\mathbf{s})$  spanned by the basis functions  $K(\mathbf{s}, \cdot, \phi_r)$ . Proof of the Theorem 2.1 is given in the Appendix.

**Theorem 2.1** *Consider the Reproducing Kernel Hilbert Space (RKHS) of the space of functions*

$$\mathcal{H}_r = \text{Span}\{K(\mathbf{s}, \cdot, \phi_r)\}$$

*spanned by the kernel at the  $r$ th resolution. Then  $\mathcal{H}_r = \mathcal{S}^{d/2+3/2}(\mathcal{R}^d)$ , where*

$$\mathcal{S}^{d/2+3/2}(\mathcal{R}^d) = \{f \in L_2(\mathcal{R}^d) \cap C(\mathcal{R}^d) : \hat{f}(\cdot)(1 + \|\cdot\|^2)^{(d+3)/4} \in L_2(\mathcal{R}^d)\},$$

*is the Sobolev space of order  $d/2 + 3/2$ , and  $\hat{f}(\cdot)$  is the Fourier transform of  $f(\cdot)$ .*

**Remark:** The result establishes that the sample paths of  $w_r(\mathbf{s})$  belong to  $\mathcal{S}^{d/2+3/2}(\mathcal{R}^d)$ . Roughly speaking, for integer  $\zeta$ ,  $\mathcal{S}^\zeta$  contains functions whose derivatives upto  $\zeta$ -th order are continuously differentiable. Thus,  $w_r(\mathbf{s})$  constructed in this way ought to provide continuously differentiable realizations of the spatial surface a priori.

The choice of the scale parameter  $\phi_r$  for the  $r$ th resolution follows from several considerations. Since the kernels in lower resolutions are meant to capture long range variabilities, one naturally imposes the constraint

$$\phi_1 > \phi_2 > \dots > \phi_R > 0. \quad (9)$$

Secondly, given  $\beta_j^r$ ,  $j = 1, \dots, J(r)$ ,  $r = 1, \dots, R$ ,  $\phi_r$  determines the set of knots in the neighborhood of  $\mathbf{s}$  for interpolating the spatial surface at  $\mathbf{s}$ . One could possibly keep  $\phi_r$

as a parameter and update them using MCMC. Our detailed investigation reveals that such an approach adds unnecessary computational burden with no apparent inferential advantage. Therefore, this article fixes  $\phi_r = \eta \|\mathbf{s}_j^r - \mathbf{s}_{j-1}^r\|$ ,  $\eta > 0$ .  $\eta$  is a tuning parameter that determines the computational advantage vis a vis long range spatial dependence between observations. Since  $\eta$  is not a parameter of interest, the present article does not attempt to make full scale Bayesian inference on  $\eta$ . Rather at each step of the MCMC iteration posterior likelihood is maximized over a grid of  $\eta$ . We further elaborate it in Section 3.1.

## 2.4 Mutiscale spatial regression model

Our proposed mutiscale spatial model typically assumes, at location  $\mathbf{s} \in \mathcal{D}$ , a response variable  $y(\mathbf{s}) \in \mathcal{R}$  along with a  $p \times 1$  vector of spatially referenced predictors  $\mathbf{x}(\mathbf{s})$  which are associated through a spatial regression model as

$$y(\mathbf{s}) = \mathbf{x}(\mathbf{s})' \boldsymbol{\gamma} + \sum_{r=1}^R \sum_{j=1}^{J(r)} K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) \beta_j^r + \epsilon(\mathbf{s}), \quad \epsilon(\mathbf{s}) \sim N(0, \sigma^2), \quad (10)$$

where  $\boldsymbol{\gamma}$  is the  $p \times 1$  vector of regression coefficient. The medium and short range spatial variability of  $y(\mathbf{s})$  is determined by the multiscale DCT term, while  $\epsilon(\mathbf{s})$  adds a jitter that corresponds to unexplained micro-scale variability or measurement error, with  $\sigma^2$  as the error variance.

## 2.5 Mutiscale shrinkage prior on $\beta_j^r$

Once the model formulation is complete, attention turns to assigning prior distributions on  $\beta_j^r, \boldsymbol{\gamma}, \sigma^2$ . While the prior specification on  $\boldsymbol{\gamma}$  and  $\sigma^2$  is straightforward, specifically  $\boldsymbol{\gamma}$  is assigned a noninformative prior and  $\sigma^2 \sim IG(c, d)$ , constructing a prior distribution on  $\beta_j^r$  requires a bit of reflection. Note that the local variability within the spatial domain varies in relation to the sub-domains. There are regions within the domain which exhibit small scale spatial variability, thereby required to be modeled by the higher resolutions. On the other hand, spatial variability is less prominent in some regions, which practically do not require higher resolutions for modeling. Mathematically, this amounts to setting  $\beta_j^r = 0$  corresponding to the knots  $\mathbf{s}_j^r$  located in the latter regions. It is also natural to assume that if  $r$ th resolution deemed unnecessary to model the surface in a sub-domain, any  $l$ th resolution for  $l > r$  should be unnecessary too to model the same subregion. For  $\mathbf{s}_j^r \in \mathcal{D}_{i_1, \dots, i_r}$ , define,

$$\mathcal{B}_{j,r}^{Subtree} = \{\beta_k^l : l \geq r, \mathbf{s}_k^l \in Subtree(\mathcal{D}_{i_1, \dots, i_r})\}.$$

Thus  $\mathcal{B}_{j,r}^{Subtree}$  is the set of coefficients corresponding to basis functions centered at knots in  $Subtree(\mathcal{D}_{i_1, \dots, i_r})$ . As per our discussion, if modeling small scale spatial variation in the subregion  $\mathcal{D}_{i_1, \dots, i_r}$  does not require resolution  $r$ , any higher resolution deem unnecessary too. This leads to condition C

**Condition C:**  $\beta_j^r = 0$  implies  $\beta_k^l = 0$ , where  $\beta_k^l \in \mathcal{B}_{j,r}^{Subtree}$ .

The problem of estimating  $\beta_j^r$ 's finds equivalence in the variable selection literature in high dimensional regression. The goal in variable selection literature lies in identifying predictors



not related to the response, equivalently the predictors having zero coefficients. A rich variety of methods have been proposed ranging from penalized optimization methods, such as Lasso (Tibshirani, 1996) and elastic net (Zou and Hastie, 2005), to Bayesian variable selection or shrinkage methods. Penalized optimization is computationally efficient in identifying unimportant predictors even in the presence of large number of predictors, but suffers due to their inability to produce accurate characterization of uncertainty. Besides, penalized optimization results are highly sensitive to the adhoc choice of tuning parameters. The Bayesian approach is attractive in its probabilistic characterization of uncertainty for regression coefficients in high dimensions and for the resulting predictions. The most popular artifact employed in the Bayesian high dimensional regression for variable selection is the wide class of spike and slab prior distributions on predictor coefficients. The widespread usage of spike and slab priors is a consequence of its easy interpretability and relatively simple computation. It is be noted that **Condition C** hinders standard application of a spike and slab prior on  $\beta_j^r$ . One can possibly build a new class of spike-and-slab prior over the traditional spike-slab prior for selective predictor inclusion (George and McCulloch, 1993; Geweke, 2004; Clyde et al., 1996) respecting the constraints imposed by **Condition C**. However, spike and slab prior faces notorious mixing issues and consequently inaccurate inference for more than a few hundred predictors. This has motivated us to derive a continuous shrinkage prior that does not set  $\beta_j^r = 0$ , but imposes a stochastic ordering between  $\beta_j^r$  along resolutions a priori. The literature on high dimensional regression have consistently found that shrinkage-promoting priors are more effective in practice on real (natural) data than exact-sparsity-promoting models, like the spike-slab setup.

An impressive variety of Bayesian shrinkage priors for ordinary high dimensional regressions with scalar/vector response on high dimensional vector predictors has been proposed in recent times, see for example Armagan et al. (2013); Hans (2009); Park and Casella (2008); Polson and Scott (2012); Carvalho et al. (2009) and references therein. Shrinkage priors are based on the principle of artfully shrinking predictor coefficients of unimportant predictors to zero, while maintaining proper estimation and uncertainty of the important predictor coefficients. However, the literature on shrinkage priors that impose increasing shrinkage along resolutions is quite insufficient. This article introduces a *multiscale tree shrinkage prior* to achieve this objective. It proposes

$$\begin{aligned} \beta_j^r &\sim N(0, \alpha_j^r) \\ \alpha_j^1 &= \delta_1^{-1}, \alpha_j^2 = \delta_1^{-1} \delta_{j,2}^{-1}, \alpha_j^r = \alpha_{\lfloor \frac{j-1}{P} \rfloor + 1}^{r-1} \delta_{j,r}^{-1} \\ \delta_1 &\sim \text{Gamma}(2, 1), \delta_{j,r} \sim \text{Gamma}(c, 1), c > 2. \end{aligned} \tag{11}$$

$\delta_{j,r}^{-1}$ 's are stochastically smaller than 1 implying increasing shrinkage apriori along a branch. In fact,  $E[\beta_j^r] = 0$  and  $\text{Var}[\beta_j^r] = \frac{1}{(c-1)^{r-1}} \rightarrow 0$ , as  $r \rightarrow \infty$ , apriori. Thus the prior distribution imposes strong apriori belief of having a parsimonious model with small number of resolutions. The proposed prior offers easy posterior updating with closed form conditional posterior distributions for all the parameters, as is discussed in the next section.

### 3 Posterior computation and inference

This section describes posterior computation and inference for multiscale DCT. The main task for inference remains that of obtaining the posterior distribution of the unknown coefficients  $\beta_j^r$  and  $\delta_{j,r}$   $j = 1, \dots, J(r)$  and  $r = 1, \dots, R$ ,  $\gamma$  and  $\sigma^2$ . The formulation of multiscale DCT is simple, so that all the parameters allow simple Gibbs sampling updates. Once posterior distributions of the parameters are available, they are employed to estimate interpolation of the residual surface and perform spatial predictions. By crucially exploiting the conditional independence among several parameters and multi-resolution structure of the problem, we obtain inference with excellent time and memory complexity (Section 3.2 and 5), can take full advantage of distributed-memory systems with a large number of nodes (Section 3.2), and is thus scalable to large spatial datasets.

#### 3.1 Posterior computation

We proceed to do parametric inference with data  $(y(\mathbf{s}_i), \mathbf{x}(\mathbf{s}_i))_{i=1}^n$  at locations  $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ . Stacking responses and predictors across locations we obtain,  $\mathbf{y} = (y(\mathbf{s}_1), \dots, y(\mathbf{s}_n))'$ ,  $\mathbf{X} = [\mathbf{x}(\mathbf{s}_1) : \dots : \mathbf{x}(\mathbf{s}_n)]'$ . Let  $\mathbf{K}$  be an  $n \times (J(1) + \dots + J(R))$  matrix whose  $i$ th row is given by  $(K(\mathbf{s}_i, \mathbf{s}_1^1, \phi_1), \dots, K(\mathbf{s}_i, \mathbf{s}_R^{J(R)}, \phi_R))'$ . Further assume  $\boldsymbol{\beta}^r = (\beta_1^r, \dots, \beta_{J(r)}^r)'$  and  $\boldsymbol{\beta} = (\boldsymbol{\beta}^1, \dots, \boldsymbol{\beta}^R)'$ ,  $y_{i,r,j} = y_i - \sum_{(k_1, k_2) \neq (j,r)} K(\mathbf{s}_i, \mathbf{s}_{k_1}^{k_2}, \phi_{k_2}) \beta_{k_1}^{k_2}$ ,  $\mathbf{y}_{r,j} = (y_{1,r,j}, \dots, y_{n,r,j})'$ ,  $\mathbf{K}_{r,j} = (K(\mathbf{s}_1, \mathbf{s}_j^r, \phi_r), \dots, K(\mathbf{s}_n, \mathbf{s}_j^r, \phi_r))'$ . The full conditional distributions of  $\gamma$ ,  $\sigma^2$ ,  $\beta_j^r$  and  $\delta_{j,r}$  are readily available in closed form and are given by

- $\gamma | - \sim N((\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'(\mathbf{y} - \mathbf{K}\boldsymbol{\beta}), \sigma^2(\mathbf{X}'\mathbf{X})^{-1})$
- $\sigma^2 | - \sim IG\left(\frac{n}{2} + c, d + \frac{1}{2}\|\mathbf{y} - \mathbf{X}\gamma - \mathbf{K}\boldsymbol{\beta}\|^2\right)$
- $\beta_j^r | - \sim N\left(\frac{\frac{\mathbf{K}'_{r,j}\mathbf{y}_{r,j}}{\sigma^2}}{\frac{1}{\alpha_j^r} + \frac{\mathbf{K}'_{r,j}\mathbf{K}_{r,j}}{\sigma^2}}, \frac{1}{\frac{1}{\alpha_j^r} + \frac{\mathbf{K}'_{r,j}\mathbf{K}_{r,j}}{\sigma^2}}\right)$ .
- Recall the definition of father node in Section 2.2. Additionally define  $father^2(\mathbf{s}_j^r)$  as the father node of the father node of  $\mathbf{s}_j^r$ . Similarly,  $father^3, \dots, father^R$  node are defined. Let  $\alpha_{j,r,-1} = \prod_{l=r}^R \delta_{father^{r+1-l}(\mathbf{s}_j^l), l-1}$  and  $\alpha_{j,1,-1} = 1$ . Then  $\delta_1 | - \sim Gamma\left(1 + \frac{J(1) + \dots + J(R)}{2}, 1 + \frac{1}{2} \sum_{r=1}^R \sum_{j=1}^{J(r)} [(\beta_j^r)^2 / \alpha_{j,r,-1}]\right)$
- Let  $\alpha_{k,l,-r,-j} = \delta_1 \prod_{h=l}^{r+2} \delta_{father^{l+1-h}(\mathbf{s}_k^h), h-1} \prod_{h=r}^2 \delta_{father^{r+1-h}(\mathbf{s}_j^h), h-1}$ ,  $\alpha_{j,r,-r,-j} = 1$ . Then  $\delta_j^r | - \sim Gamma\left(c + \frac{\#\beta_{j,r}^{Subtree}}{2}, 1 + \frac{1}{2} \sum_{l \geq r, \mathbf{s}_k^l \in Subtree(\mathbf{s}_j^r)} (\beta_k^l)^2 / \alpha_{k,l,-r,-j}\right)$  for  $r > 1$ .
- Finally at each iteration, joint posterior distribution is maximized over a discrete grid of  $\eta$  values fixing all other parameters at the current iterate,  $\eta \in \{1, \dots, h_\eta\}$ ,  $h_\eta$  is an integer. In all simulation studies and in the real data analysis, we never found the

maximization of the posterior over  $\eta$  to occur for  $\eta$  values more than 5. Thus, we fix  $h_\eta = 5$  for all empirical investigations. We must mention that the posterior is maximized at  $\eta = 1, 2$  in most of the iterations.

### 3.2 Distributed computation and computational complexities

An important advantage of the multiscale DCT is that it facilitates distributed computation with little communication overhead at a large number of nodes, each only dealing with a small subset of the data. This section describes the distributed computing algorithm as well as the associated computation complexity.

To begin with, one must acknowledge that posterior updating of  $\boldsymbol{\gamma}, \sigma^2, \delta_{j,r}, \delta_1$  can be carried out rapidly without having to store the entire data in one processor. The main computational difficulty comes from updating of  $\boldsymbol{\beta}$ . Single updating of  $\beta_j^r$  introduces too much autocorrelation, while joint updating of  $\boldsymbol{\beta}$  requires inverting  $(\sum_{r=1}^R J(r)) \times (\sum_{r=1}^R J(r))$  matrix which is infeasible. The use of compactly supported basis functions offers an excellent solution by carefully exploiting conditional independence between blocks of  $\boldsymbol{\beta}$ . For  $m = 1, \dots, J(1)$ , define the *neighborhood function*  $\mathcal{N}(m)$  of  $m$  by

$$\mathcal{N}(m) = \{j : \|\mathbf{s}_j^1 - \mathbf{s}_m^1\| < 2\eta\}$$

Similarly, the *neighborhood data function* is defined as

$$\mathcal{N}_D(m) = \{j : \|\mathbf{s}_j^1 - \mathbf{s}_m^1\| < \eta\}$$

Let  $\boldsymbol{\beta}_{j,r}^{Subtree}$  be a vector composed of all elements in  $\mathcal{B}_{j,r}^{Subtree}$ . Clearly,  $\boldsymbol{\beta} = (\boldsymbol{\beta}_{1,1}^{Subtree}, \dots, \boldsymbol{\beta}_{1,J(1)}^{Subtree})'$ . Exploiting properties of the compactly supported basis functions, one obtains

$$\boldsymbol{\beta}_{s,1}^{Subtree} |_{-} \stackrel{\mathcal{L}}{=} \boldsymbol{\beta}_{s,1}^{Subtree} |_{\mathbf{y}_{\mathcal{N}_D(s)}, \boldsymbol{\beta}_{\mathcal{N}(s),1}^{Subtree}}, s = 1, \dots, J(1).$$

Here  $\mathcal{N}_D(s) = \{i_1, \dots, i_l\}$  will imply that  $\mathbf{y}_{\{i_1, \dots, i_l\}}$  and  $\mathbf{X}_{\{i_1, \dots, i_l\}}$  are the data responses and corresponding predictors in the domain  $\mathcal{D}_{i_1} \cup \dots \cup \mathcal{D}_{i_l}$ .

Algorithm 1 describes details of the computation strategy we adopt. As per Algorithm 1, the computation involves  $J(1)$  nodes with  $m$ th node storing  $\{\mathbf{y}_{\mathcal{N}_D(m)}, \mathbf{X}_{\mathcal{N}_D(m)}\}$  and executing posterior updates of  $\boldsymbol{\beta}_{m,1}^{Subtree}$ . The main computation cost involved in the  $m$ th node is in computing Cholesky decomposition of a  $dim(\boldsymbol{\beta}_{\mathcal{N}(m),1}^{Subtree}) \times dim(\boldsymbol{\beta}_{\mathcal{N}(m),1}^{Subtree})$  and multiplying a  $dim(\mathcal{N}_D(m)) \times (\sum_{r=1}^R J(r))$  matrix with a vector of dimension  $(\sum_{r=1}^R J(r))$ . They incur computation complexities of  $O(dim(\mathcal{N}(m))^3)$  and  $O(dim(\mathcal{N}_D(m)) \sum_{r=1}^R J(r))$  respectively. Since  $dim(\boldsymbol{\beta}_{1,\mathcal{N}(m)}^{Subtree}) = ((2d)^R - 1)/(2d - 1)$ , the computation time for the former is low. Choosing  $J(1)$  large enough one can reduce the computation time for the latter as well. The storage complexity is also dominated by  $dim(\mathcal{N}_D(m))$ .

### 3.3 Prediction and residual surface interpolation

Let  $\mathbf{s}_0$  be any location in the domain, where we seek to predict  $y(\mathbf{s}_0)$ , based on a given vector of predictors  $\mathbf{x}(\mathbf{s}_0)'$ . The spatial prediction at  $\mathbf{s}_0$  proceeds from the posterior predictive

---

**Algorithm 1** Distributed computing of the posterior distribution of  $\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma^2, \delta_{j,r}$ 


---

- a. **No. of nodes used:** Use  $J(1)$  nodes for computation.
  - b. **MCMC initialization:** Initialize all parameters.
  - c. At the  $t$ th iteration, MCMC iterates are given by  $(\boldsymbol{\beta}_{m,1}^{Subtree})^{(t)}, m = 1, \dots, J(1), \sigma^{2(t)}, \boldsymbol{\gamma}^{(t)}, \delta_{j,r}^{(t)}, j = 2, \dots, J(r); r = 1, \dots, R$  and  $\delta_1^{(t)}$ .
  - d. Maximize posterior likelihood w.r.t.  $\eta \in \{1, \dots, h_\eta\}$ . Compute  $(\mathbf{y}_{\mathcal{N}_D(m)}, \mathbf{X}_{\mathcal{N}_D(m)})$  according to the maximized  $\eta$ . At the  $t$ th iteration store  $(\mathbf{y}_{\mathcal{N}_D(m)}, \mathbf{X}_{\mathcal{N}_D(m)})$  in the  $m$ th node.
  - e. For  $m = 1 : J(1)$  in parallel in  $J(1)$  different nodes
    - i.  $(t + 1)$  iterate of  $(\boldsymbol{\beta}_{m,1}^{Subtree})^{(t+1)}$  is obtained by drawing from  $\boldsymbol{\beta}_{1,m}^{Subtree} | (\boldsymbol{\beta}_{\mathcal{N}(m),1}^{Subtree})^{(t)}$ .
  - f. For  $m = 1 : J(1)$  in parallel in  $J(1)$  different nodes
    - i. Calculate  $\mathbf{X}'_m \mathbf{X}_m, \mathbf{y}_m - \mathbf{K}_m \boldsymbol{\beta}$ , where  $\mathbf{K}_m = (K(\mathbf{s}, \mathbf{s}_1^1, \phi_1), \dots, K(\mathbf{s}, \mathbf{s}_{J(R)}^R, \phi_R)), \mathbf{s} \in \mathcal{D}_m$ .
  - g. Use the fact that  $\sum_{m=1}^{J(1)} \mathbf{X}'_m \mathbf{X}_m = \mathbf{X}' \mathbf{X}$  and  $\mathbf{y} - \mathbf{K} \boldsymbol{\beta} = (\mathbf{y}_1 - \mathbf{K}_1 \boldsymbol{\beta}, \dots, \mathbf{y}_{J(1)} - \mathbf{K}_{J(1)} \boldsymbol{\beta})'$  to update from the full condition of  $\boldsymbol{\gamma}$ .
  - h. Update  $\delta_{j,r}^{(t+1)}$  and  $\delta_1^{(t+1)}$  at the  $(t + 1)$ th iteration.
- 

Time(multiple processor)	Time (single processor)
$\left[ (2\eta - 1) \frac{n}{J(1)} \right] \left[ \sum_{r=1}^R J(r) \right] + \frac{(2d)^{r-1}}{2d-1}$	$J(1) \left[ (2\eta - 1) \frac{n}{J(1)} \right] \left[ \sum_{r=1}^R J(r) \right]$

---

Table 1: Dominant terms in calculating time complexities for multiscale DCT with single and multiple processors.

distribution

$$p(\mathbf{y}(\mathbf{s}_0) | \mathbf{y}) = \int p(\mathbf{y}(\mathbf{s}_0) | \mathbf{y}, \boldsymbol{\Theta}) p(\boldsymbol{\Theta} | \mathbf{y}) d\boldsymbol{\Theta}, \quad (12)$$

using composition sampling, where  $\boldsymbol{\Theta} = (\sigma^2, \boldsymbol{\gamma}, (\beta_j^r)_{j,r=1}^{J(r),R}, (\delta_{j,r})_{j,r=1}^{J(r),R})$ . For each  $\{\boldsymbol{\Theta}^{(l)}\}$ ,  $l = 1, 2, \dots, L$ , obtained from the posterior distribution  $p(\boldsymbol{\Theta} | \mathbf{y})$ , draw  $\mathbf{y}(\mathbf{s}_0)^{(l)}$  from  $p(\mathbf{y}(\mathbf{s}_0) | \boldsymbol{\Theta}^{(l)})$ . The resulting  $\mathbf{y}(\mathbf{s}_0)^{(l)}$ ,  $l = 1, 2, \dots, L$  are samples from (12). This is especially simple for multiscale DCT as  $p(\mathbf{y}(\mathbf{s}_0) | \boldsymbol{\Theta})$  turns out to be a normal distribution.

For multiscale DCT, full Bayesian inference on the residual spatial surface at any unobserved location  $\mathbf{s}_0$  is trivially obtained. For each posterior sample  $\{\boldsymbol{\Theta}^{(l)}\}$ ,  $l = 1, 2, \dots, L$ , compute  $w(\mathbf{s}_0)_l = \sum_{r=1}^R \sum_{j=1}^{J(r)} K(\mathbf{s}_0 - \mathbf{s}_j^r, \phi_r) (\beta_j^r)^{(l)}$ .  $w(\mathbf{s}_0)_l$  are samples from the posterior distribution of the residual process. Surface interpolation is straightforward hereafter.

## 4 Theoretical properties

We establish convergence results for multiscale DCT regression model (10) under the simplifying assumptions<sup>1</sup> that the predictor coefficient  $\boldsymbol{\gamma} = (0, \dots, 0)$ .

---

<sup>1</sup>Simplifying assumption is merely to ease notation and calculations; all results generalize in a straightforward manner.

Define two metrics in the function space given by

$$\begin{aligned} \|w\|_\infty &= \sup_{\mathbf{s} \in \mathcal{D}} |w(\mathbf{s})|, \\ \|w\|_\zeta &= \max_{k \leq \lfloor \zeta \rfloor} \sup_{\mathbf{s} \in \mathcal{D}} |D^k w(\mathbf{s})| + \max_{\tilde{k} \leq \lfloor \zeta \rfloor} \sup_{\mathbf{s}, \mathbf{s}' \in \mathcal{D}} \frac{|D^{\tilde{k}} w(\mathbf{s}) - D^{\tilde{k}} w(\mathbf{s}')|}{\|\mathbf{s} - \mathbf{s}'\|_\zeta^{\lfloor \zeta \rfloor}}, \end{aligned}$$

where  $D^k = \frac{\delta^{k_1+k_2}}{\delta s_1^{k_1} \delta s_2^{k_2}}$ , for  $k_1, k_2 \in \mathbb{N}$  and  $\mathbf{s} = (s_1, s_2)'$ . Further define the sets

$$\begin{aligned} \Theta_\zeta &= \left\{ w(\mathbf{s}) : w(\mathbf{s}) = \sum_{r=1}^R \sum_{j=1}^{J(r)} K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) \beta_j^r, R \in \mathbb{N}, \mathbf{s}_j^r \in \mathcal{R}^2, \beta_j^r \in \mathcal{R}, \|w\|_\zeta < \infty, \right\} \\ \Theta_\zeta^n &= \{w \in \Theta_\zeta : \|w\|_\zeta < n^\alpha, \alpha \in (1/2, 1]\} \\ \Theta_{\zeta,c} &= \text{Closure under } \|\cdot\|_\infty \text{ of } \Theta_\zeta \\ \mathcal{B}_{\epsilon,n} &= \left\{ w \in \Theta_\zeta^n : \frac{1}{n} \sum_{i=1}^n |w(\mathbf{s}_i) - w_0(\mathbf{s}_i)| < \epsilon, \left| \frac{\sigma}{\sigma_0} - 1 \right| \right\}. \end{aligned}$$

**Theorem 4.1** *Let  $\mathcal{P}_{w_0, \sigma_0^2}$  denotes the true data generating joint distribution of  $\{y_i\}$ . Assume*

- (a)  $\mathcal{D}$  is compact.
- (b)  $K(\cdot, \cdot, \phi_r)$  is continuous.
- (c)  $w_0 \in \Theta_{\zeta,c}$ ,  $\|w_0\|_\zeta < \infty$ , for some  $\zeta$ .

Then for any  $(w_0, \sigma_0^2) \in \Theta_{\zeta,c} \times \mathcal{R}^+$  and for any  $\epsilon > 0$ ,

$$\lim_{n \rightarrow \infty} \Pi((w, \sigma^2) \in \mathcal{B}_{\epsilon,n} | y_1, \dots, y_n) = 0$$

almost surely under  $\mathcal{P}_{w_0, \sigma_0^2}$ .

Theorem 4.1 establishes consistency of estimating the data generating surface  $w_0$  and the true error variance  $\sigma_0^2$ . The proof proceeds along the same line of arguments outlined in Choi and Schervish (2007), Pillai (2008) and is provided in the Appendix.

## 5 Simulation studies

In this section, we use synthetic datasets to assess model performance with regard to interpolating unobserved residual spatial surface and predicting at new locations. To begin with, we present a one dimensional simulation experiment on a large dataset. The one dimensional simulation experiment helps to build intuition on how different resolutions capture large and small scale variabilities, including the advantage of choosing the tree shrinkage prior. Once computational and inferential aspects of multiscale DCT are discussed in one dimension, we provide a more realistic two dimensional example where computation time and performance of multiscale DCT will be compared with state-of-the-art and popular spatial models for big data. A non-distributed implementation of the methods are carried out in R version 3.3.1 on a 16-core machine (Intel Xeon 2.90GHz) with 64GB RAM.

## 5.1 One dimensional Example

For the one dimensional example, we simulated a dataset of size  $n = 20,000$  from the model with an intercept, a predictor and a spatial function  $w_0(s)$  in  $[0, 10]$  given by

$$w_0(s) = \begin{cases} \sin(2\pi s)s, & \text{if } 0 \leq s < 2 \\ |\sin(s - 3)|^3, & \text{if } 2 \leq s < 4 \\ 5|s - 5|, & \text{if } 4 \leq s < 6 \\ \sin(2\pi s)s, & \text{if } 6 \leq s < 10. \end{cases} \quad (13)$$

A plot of the true spatial function  $w_0(s)$  is provided in Figure 2. The function is piecewise differentiable which makes the estimation challenging.

We fit multiscale DCT with  $J(1) = 30$  to this dataset. As a competitor to multiscale DCT we implement

**DCT-GDP:** DCT-GDP uses the same basis functions as multiscale DCT, but replaces *multiscale tree shrinkage prior* by Generalized Double Pareto (GDP, Armagan et al. (2013)) shrinkage prior on the basis coefficients. GDP prior does not allow any multiscale structure, thereby asserting equivalent apriori shrinkage on all the basis coefficients.

**DCT-Normal:** DCT-Normal also uses the same basis functions with the prior on basis coefficients given by the independent normal prior distributions.

The two competitors are mainly aimed at comparing the inferential advantage of the tree shrinkage prior over the ordinary shrinkage prior and normal prior distributions. Additionally, we fit multi-scale DCT with one and two resolutions to assess how the choice of  $R = 3$  improves inference. Multiscale DCT with one and two resolutions are referred to as MDCT(1) and MDCT(2) respectively.

Figure 2 reveals the role played by the three resolutions in estimating  $w_0(s)$ . Clearly, resolutions 1 and 2 mostly capture global trends. While resolution 1 mostly captures positive side of the sinusoidal curve, negative extremities of the sinusoidal curve is mostly reconstructed by resolution 2. Resolutions 3 captures the local variability found in the interval  $[4, 10]$ .

The inferential performance of MDCT is evaluated in estimating the spatial surface using mean squared error. To be more precise, let  $\hat{w}(s_i)$  be the posterior median of  $w(s_i)$ . Define mean squared error (MSE) by  $MSE = \frac{1}{n} \sum_{i=1}^n (\hat{w}(s_i) - w_0(s_i))^2$ . Average MSE along with the associated standard errors over multiple simulations for the three competing models are presented in Figure 2. It is evident from these figures that MDCT, with the same number of knots and same basis functions, provide improved inference, the reason being implementation of a structured prior distribution on the basis coefficients. The computation times to implement the three competitors are about the same, with one MCMC iteration in MDCT taking  $\approx 0.33$  seconds to run the full scale inference. Additionally, there seems to be a substantial improvement in terms of MSE with increasing resolutions, though it stabilizes after  $R = 3$ .

The one dimensional exploration of MDCTs presented above, points at a few important advantages. First of all, multi-scaling is able to capture local features succinctly, yielding superior inference with similar number of knots and same basis functions over one scale DCT. Secondly, the computational advantage of multiscale DCT is enormous given that

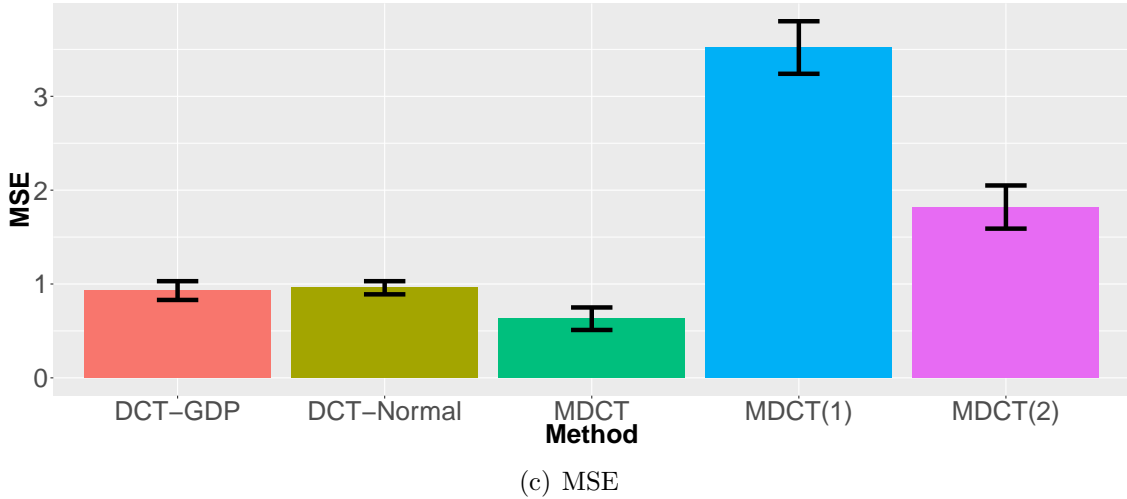
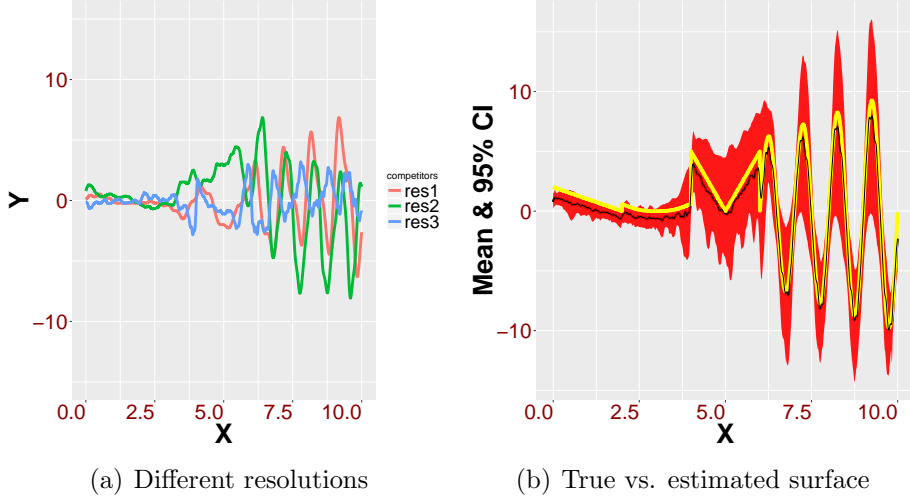


Figure 2: (a) Estimated mean function at different resolutions; (b) shows the true vs. the estimated function in  $R = 3$  resolutions. The true function is in yellow and the estimated function is in black. 95% confidence bands for the estimated function are displayed in red. (c) shows the MSE with associated standard errors for all competitors.

full Bayesian inference can be performed with a series of local computations. Given the architecture of MDCT, it is possible to implement MDCT by storing subsets of data in different processors. Moreover, it does not require storing large dimensional covariance matrices that incurs onerous storage burden, as is the case for Gaussian process based spatial models. In the next section a more involved comparative analysis of MDCT with popular competitors is presented in the context of two dimensional spatial examples.

## 5.2 Two dimensional example

In this section, we use two dimensional synthetic datasets to assess the performance of MDCT in comparison to popular models for large spatial data. For the sake of our expo-

sition, MDCT is implemented with 3 resolutions having a total of 2100 basis functions. As competitors to MDCT we implement:

(1) **Modified predictive process (MPP)**: a popular low rank model fitted to the entire data with full Bayesian implementation (Finley et al., 2009);

(2) **LatticeKrig**: LatticeKrig (Nychka et al., 2015) is a recently proposed multiresolution model for big data that employs kernel convolution with radial basis functions and Gaussian Markov Random Field (GMRF) distribution on the basis coefficients. `LatticeKrig` package in R is employed for non-Bayesian implementation of LatticeKrig with 3 resolutions having a total 12678 basis functions. LatticeKrig is a closely related competitor to MDCT with the major difference stems from using a GMRF prior on basis coefficients.

(3) **LaGP**: Local approximate Gaussian process (Gramacy and Apley, 2015). LaGP has emerged from the computer experiment literature and is devised to perform fast local neighborhood kriging with Gaussian processes. LaGP is not designed to provide full scale Bayesian inference and is only employed to compare predictive inference with other competitors.

We acknowledge **NNGP** (Datta et al., 2015) and **multiresolution predictive process** (Katzfuss, 2016) as the two important competitors of MDCT. However, both these methods are complicated in terms and implementation and till date there is no readily available open source software to implement these methods. To avoid the risk of implementing them incorrectly, we refrain from showing inferential results from these two methods. Moreover, it is argued in Gramacy and Apley (2015) that LaGP performs better than nearest neighbor methods in many applications.

Bayesian implementation of MPP is performed using the package `spBayes` in R. It is well known that MPP is not a suitable model when the sample size is very large. Therefore to accommodate MPP as a competitor in the analysis, we design simulation studies for datasets not larger than  $\sim 10500$  locations. Note that as the number of knots increases, performance of MPP improves, though adding much burden to computation. Therefore, while comparing MPP, we focus both on computation time and accuracy. Finally, the `laGP` package facilitates frequentist implementation of LaGP in R. All the interpolated spatial surfaces are obtained using the R package `MBA`.

To illustrate the performance of the competitors, 10,500 observations within  $[0, 1] \times [0, 1]$  domain are generated from the classical geostatistical model with likelihood  $\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\gamma} + \mathbf{w}_0, \sigma^2\mathbf{I})$ . The model includes an intercept  $\gamma_0$  and a predictor  $\mathbf{x}(s)$  drawn i.i.d from  $N(0, 1)$  with the corresponding coefficient  $\gamma_1$ ,  $\boldsymbol{\gamma} = (\gamma_0, \gamma_1)$ .  $\mathbf{w}_0 = (w_0(\mathbf{s}_1), \dots, w_0(\mathbf{s}_n))'$  is an  $n$  dimensional vector that follows a multivariate normal distribution with mean  $\mathbf{0}_n$  and the covariance matrix of the order  $n \times n$  with  $(i, j)$  entry given by  $v(\mathbf{s}_i, \mathbf{s}_j, \theta_1, \theta_2, \nu)$ .  $v$  is chosen from the popular Matern class of correlation functions given by

$$v(s, s', \theta_1, \theta_2, \nu) = \frac{\theta_1}{2^{\nu-1}\Gamma(\nu)} (\|s - s'\|_{\theta_2})^\nu \mathcal{K}_\nu(\|s - s'\|_{\theta_2}); \theta_2 > 0, \nu > 0, \quad (14)$$

with  $\theta_2, \nu$  controlling spatial decay and process smoothness respectively,  $\Gamma$  is the usual Gamma function and  $\mathcal{K}_\nu$  is a modified Bessel function of the second kind with order  $\nu$  (Stein, 2012). We fixed  $\nu = 0.5$  which reduces to the exponential covariance kernel and generates continuous but non-differentiable sample paths. For simulations, we use  $\theta_2 = 3$  and the ratio of spatial to noise variability is kept at 20. Among 10,500 observations, 10000



are randomly selected for model fitting and the rest are kept as a test dataset to facilitate predictive inference.

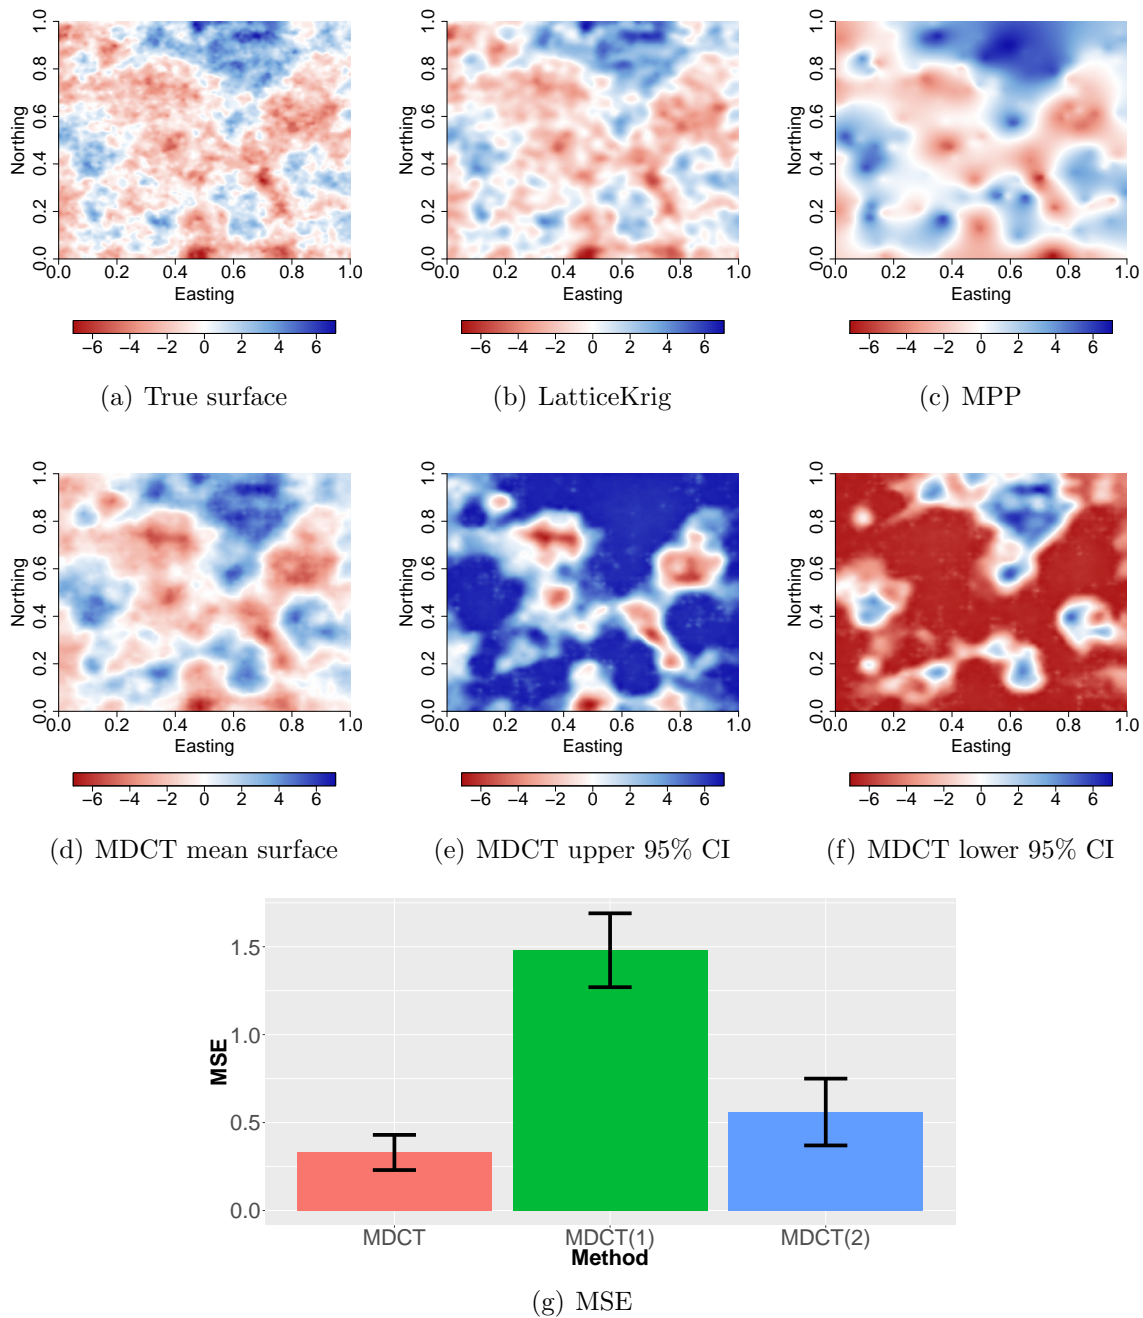


Figure 3: (a) True data generation surface along with posterior mean residual surface from (b) LatticeKrig (c) modified predictive process (d) multiscale DCT and; (e) and (f) present estimated 95% upper and lower quantile surfaces; (g) MSE along with standard errors for MDCT, MDCT(1) and MDCT(2)

Figure 3 presents true data generating surface and the estimated residual surfaces for

LatticeKrig (LK), MDCT and MPP. MPP shows excessive smoothing, while LatticeKrig and multiscale DCT yield essentially equivalent estimates of the spatial surface. Further MSE for multiscale DCT (MDCT) with one resolution (MDCT(1)) and two resolutions (MDCT(2)) are also plotted in Figure 3. As expected, MSE for MDCT decreases as the number of resolutions increases.

Next, we turn our attention to the predictive inference of the competitors. To this end, we compare all competitors based on their ability to produce accurate point prediction and predictive uncertainties. Point prediction of the competitors are judged based on the mean squared prediction error (MSPE) metric. For Bayesian competitors, predictive uncertainties are characterized by the length and coverage of 95% predictive intervals. The frequentist implementation of LK provides only predictive point estimates and standard errors, while LaGP yields posterior predictive mean and standard error (SE) at the unobserved locations. Thus, for these two competitors, approximate 95% predictive intervals are constructed by considering [predictive point estimate  $-1.96 SE$ , predictive point estimate  $+1.96 SE$ ].

It is clear from Figure 4 that MDCT yields a MSPE that is competitive with those of LK and LaGP, though the latter two slightly outperforms MDCT. Interestingly, MDCT with  $R = 3$  resolution shows significantly improved performance compared to MDCT with resolution one or two. Intuitively, one can explain such an upsurge in performance by noting that the true surface generated from the exponential correlation function has significant local behavior, thereby limiting the “borrowing of information” across space. Moreover, coverage and length of 95% predictive intervals of MDCT demonstrates accurate characterization of predictive uncertainty as opposed to MPP which shows some under-coverage. Likewise, the approximate 95% predictive interval of LaGP exhibits little under-coverage, while for LK we observe massive under-coverage.

To check the sensitivity with respect to the choice of  $R = 3$ , we run our analysis with  $R = 4, 5$  and compare to the MSPE obtained from  $R = 3$ . Table 2 clearly shows that beyond  $R = 3$  the improvement in MSPE performance is not commensurate with the increase in computation cost. We found this conclusion to hold across a number of simulation studies. Therefore  $R = 3$  is kept throughout this article.

<i>MDCT</i>	<i>R = 3</i>	<i>R = 4</i>	<i>R = 5</i>
MSPE	0.25 <sub>0.03</sub>	0.21 <sub>0.02</sub>	0.19 <sub>0.02</sub>

Table 2: Average mean squared prediction error for MDCT with  $R = 3, 4, 5$ . Associated standard errors over 5 repeated simulations are provided in the subscript.

Our analysis finds some interesting points about the competing methods. Recall that MDCT and LatticeKrig are similar in terms of their multiscale structure, the only difference being the distribution on the basis coefficients. Our investigation reveals that tree shrinkage priors are appropriately calibrated so as to yield similar point estimates with LatticeKrig with much less number of basis functions. The GMRF prior distribution on basis coefficients hinders efficient local computation in LatticeKrig. As a result LatticeKrig in its present form has less scope of being computationally efficient. In contrast, multiscale DCT is able draw full scale Bayesian inference with a series of parallelizable local computations. At the same time, tree shrinkage prior in multiscale DCT penalizes model overfitting with unnecessary resolutions. Similarly, LaGP has a notable disadvantage for not being model based. As a

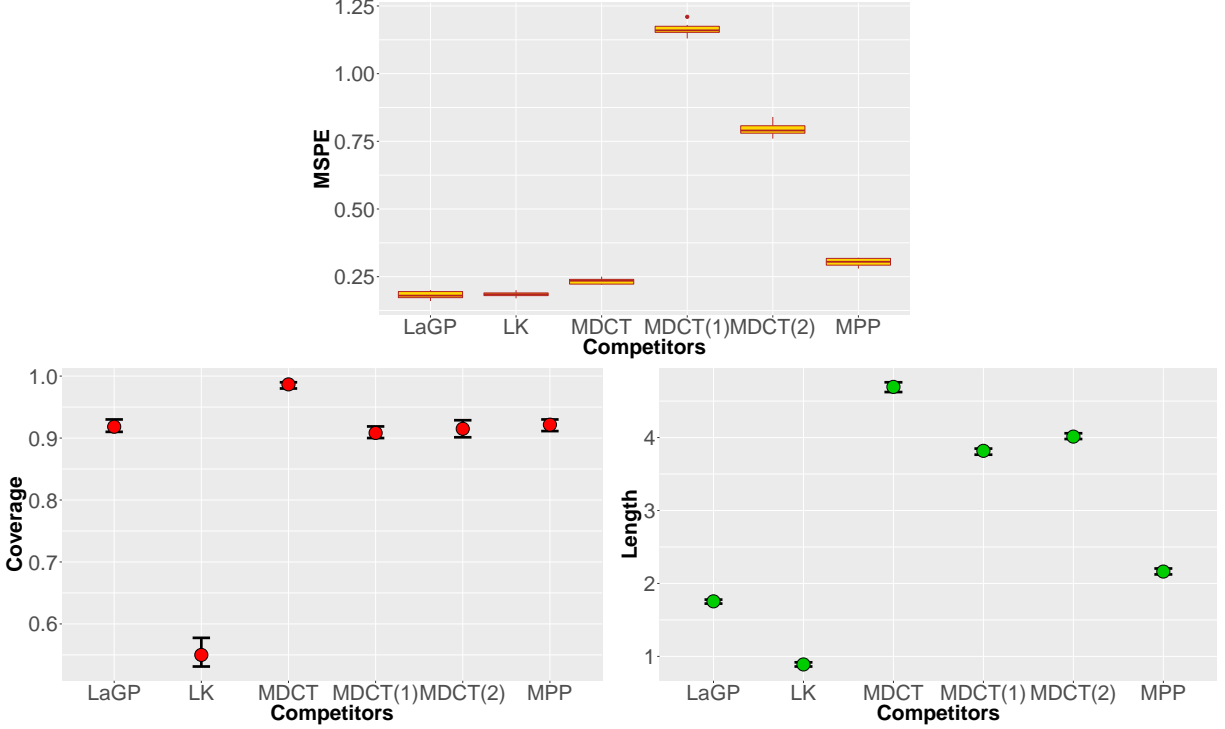


Figure 4: Plot at the top indicates boxplot of mean squared prediction error for all competitors over a few replications. Second and third plots show coverage and length of 95% predictive intervals for the competitors over the same replications. LatticeKrig shows extreme under-coverage compared to others.

result, it is not clear how to extend LaGP for non-Gaussian data, while MDCT structure can readily be embedded into a hierarchical structure to model non-Gaussian spatial data, as is described in the next section.

*Computation Time:* MDCT in this specific example takes approximately 3.07 seconds per iteration with non-optimized, non-parallel R implementation, while MPP implemented in C++ takes close to 7.2 seconds to run one MCMC iteration. We notice, though, that MPP performs the estimation of the basis functions, while MDCT assumes a fixed form with the empirical Bayes estimate of  $\eta$  at every iteration. However, even with fixed basis functions, MDCT is able to demonstrate superior inference to MPP. It is possible that, performing more elaborate inference on the kernel parameters of the MDCT would improve the predictive performance of the model. This would come at the cost of increased computational complexity, thus the benefits of such extension is unwarranted. Recall that, in this specific example MDCT is implemented with  $J(1) = 400$ . To understand how the computation time of MDCT varies vis a vis MPP with changing  $n$  and  $J(1)$ , we implement both MPP and MDCT with  $J(1) = 5^2, 10^2$  for different sample sizes. Figure 5 reports the computation time for the competitors using the R function `Sys.time`. It is be noted that MDCT can be implemented either by sequentially updating  $J(1)$  blocks of parameters or by parallelly updating these  $J(1)$  blocks independently in  $J(1)$  nodes. Thus the figure displays computation

time with both parallel and non-parallel implementation of the MDCT model. The blue line shows the computation time for MPP from  $n = 1000$  to  $n = 100,000$ , while the green and red lines show the same for MDCT with  $J(1) = 5^2$  and  $J(1) = 10^2$  respectively. Clearly, the computation time for MDCT increases linearly with  $n$  for both cases. It is also noticed that the computation time of MDCT with  $J(1) = 5^2$  is about 4-5 times faster than  $J(1) = 10^2$ . The increase in computation time is due to sequential updating of parameters in  $J(1)$  blocks. With a proper parallelized implementation of MDCT, arguably the increase in computation time from  $J(1) = 25$  to  $J(1) = 100$  will be minimal. We found that practical implementation of MPP becomes prohibitive due to both, memory allocation and exorbitant computation time, for  $n$  above 100,000. On the contrary, MDCT facilitates distributed storage of big data into multiple processors. It is worth noticing that frequentist implementation of LK and LaGP draw inference for a point estimate within a few minutes. In summary, 2D simulation examples comprehensively establishes MDCT as an effective tool for fast Bayesian implementation of large scale spatial data.




Figure 5: Computation time for MPP with 200 knots, MDCT with  $J(1) = 25$  and  $J(1) = 100$ . Computation times per MCMC iteration are presented for both MDCT and MPP.

### 5.2.1 Two dimensional illustration of MDCT with binary spatial data

To demonstrate the flexibility offered by MDCT as opposed to ad-hoc predictive methods (such as LaGP), performance of MDCT is investigated under non-Gaussian binary spatial data. For this purpose 10,500 observations within  $[0, 1] \times [0, 1]$  domain are generated from the probit spatial regression model. More precisely, with  $\mathbf{x}(\mathbf{s}_i)$  as the predictor vector at  $\mathbf{s}_i$ , the response  $y_i$  is simulated using

$$y_i \stackrel{ind}{\sim} Ber(p_i)$$

$$\Phi^{-1}(p_i) = \mathbf{x}(\mathbf{s}_i)' \boldsymbol{\gamma} + w_0(\mathbf{s}_i).$$

The model includes an intercept  $\gamma_0$  and a predictor  $\mathbf{x}(\mathbf{s})$  drawn i.i.d from from  $N(0, 1)$  with the corresponding coefficient  $\gamma_1$ ,  $\boldsymbol{\gamma} = (\gamma_0, \gamma_1)$ .  $\mathbf{w}_0 = (w_0(\mathbf{s}_1), \dots, w_0(\mathbf{s}_n))'$  is an  $n$  dimensional

vector that follows a multivariate normal distribution with mean  $\mathbf{0}_n$  and the covariance matrix of the order  $n \times n$  specified through the Matérn (14) class of correlation functions. A random subset of 10000 observations are selected for model fitting and the rest is used to judge performance of MDCT as a binary classifier.

To implement MDCT for binary spatial data, we employ

$$y_i \stackrel{ind}{\sim} Ber(p_i)$$

$$\Phi^{-1}(p_i) = \mathbf{x}(\mathbf{s}_i)' \boldsymbol{\gamma} + \sum_{r=1}^R \sum_{j=1}^{J(r)} K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) \beta_j^r.$$

The posterior computation of all model parameters proceeds with the standard data augmentation procedure (see Albert and Chib (1993)). Due to space constraint, we omit all the details and plan to elaborate it in a future article. For the sake of our exposition, MDCT is implemented with 3 resolutions having a total of 2100 basis functions.

Note that the binary regression precludes the possibility of employing LaGP as a competitor. On the other hand, `LatticeKrig` package implements LatticeKrig only for continuous response. Thus as a competitor, binary spatial regression with modified predictive process is implemented in R package `spBayes`.

Figure 6 shows the true surface and estimated surfaces from MDCT and MPP. Since the surface estimation from binary spatial regression is a notoriously challenging problem, it comes with no surprise that the performance of all competitors deteriorate when compared with Figure 3. However, among the two competitors MDCT outperforms MPP considerably. It becomes clear from Figure 6 that MPP undergoes massive oversmoothing and loses most of the local features in the spatial surface. MDCT also experiences smoothing, though to a much lesser degree than MPP. Referring to Figure 6(d), MDCT appears to be marginally better than MPP in terms of out of sample classification (Area under the ROC curve for MDCT is 0.74, while the same for MPP is 0.68). The binary spatial regression analysis further corroborates the flexibility and accuracy of MDCT. Next section discusses performance of MDCT along with its competitors on a sea surface temperature dataset.

## 6 Analysis of sea surface temperature data

A description of the evolution and dynamics of the oceans' temperature is a key component of the study of the Earth's climate. Historical records of ocean data have been collected for the purpose of understanding the properties of water masses and their changes in time. They are also used to assess, initialize and constrain numerical models of the climate. Sea surface temperature data from ocean samples have been collected by voluntary observing ships, buoys, military and scientific cruises for decades. During the last 20 years or so, this wealth of data has been complemented by regular streams of remotely sensed observations from satellite orbiting the earth. Increasingly sophisticated climatological research requires, not only the description of the mean state and the relevant trends in ocean data, but also a careful quantification of the data variability at different spatial and temporal scales. A number of articles have appeared to address this issue in recent years, see e.g. Higdon (1998),

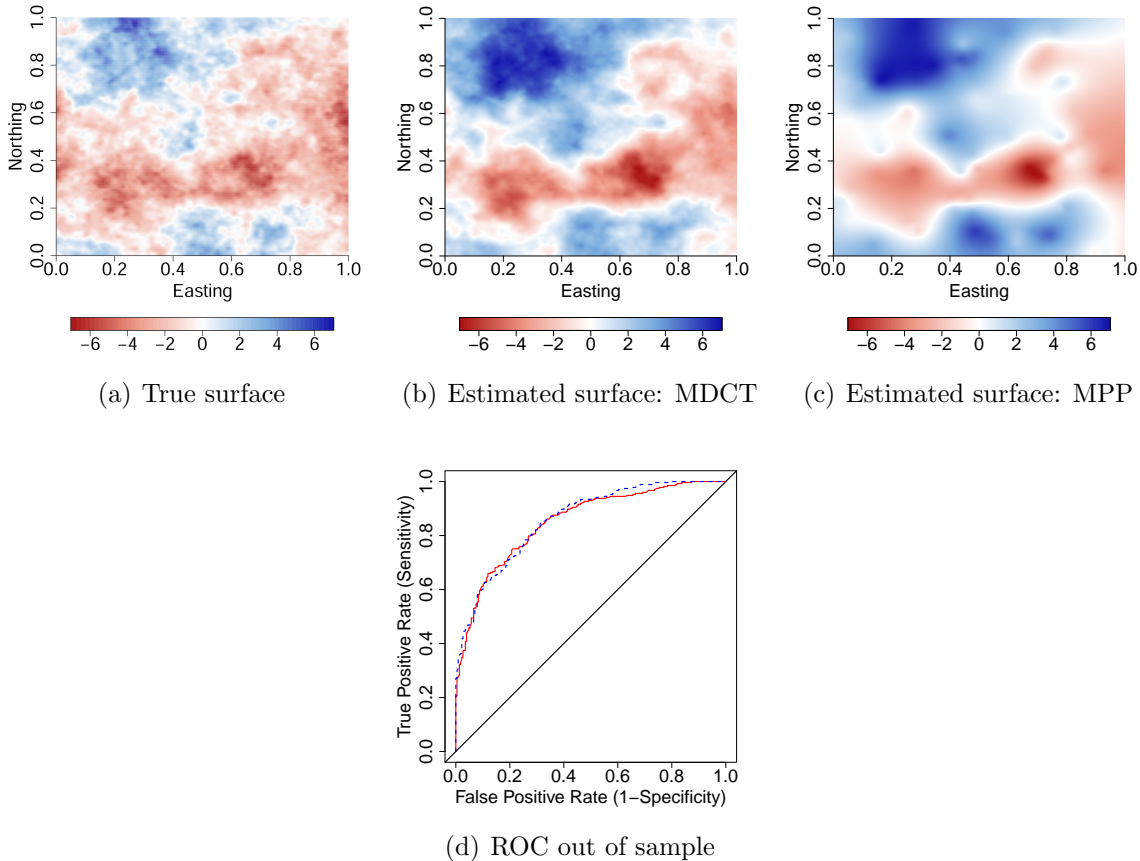


Figure 6: (a) True surface generating the data. Figures (b) and (c) present the posterior predictive mean of estimated spatial surfaces from MDCT and MPP. (d) shows out of sample ROC curves for MPP and MDCT. Dotted line presents ROC for MDCT, while solid line presents ROC for MPP.

Lemos and Sansó (2009), Lemos and Sansó (2006), Berliner et al. (2000).

This article considers the problem of capturing the spatial trend and characterizing the uncertainties in the sea surface temperature (SST) in the West coast of mainland USA, Canada and Alaska between  $30^{\circ} - 60^{\circ}$  N. latitude and  $122^{\circ} - 152^{\circ}$  W. longitude. The dataset is obtained from NODC World Ocean Database 2016 and we use the data collected in the month of October for all the spatial locations. Note that, for this example, we ignore the temporal component. We perform screening of the data to ensure quality control and then choose a random subset of 113,412 spatial observations over the domain of interest. Out of the total observations, about 90%, i.e 100,000 observations are used for model fitting and rest are used for prediction. We replicate this procedure 5 times to eliminate any chance factor in our analysis. The domain of interest is large enough to allow considerable spatial variation in SST from north to south and provides an important first step to extend these models for the analysis of global scale SST database.

The plot of the sea surface temperature along with coastal lines of Western United States and Canada is shown in Figure 7(a). The data show a clear decreasing SST trend with

increasing latitude. Consequently, we add latitude and longitude as linear predictors to explain the long-range directional variability in the SST. We fitted a non-spatial model with latitude and longitude as linear predictors using ordinary least square (OLS). The resulting residuals are shown in Figure 7(b). The residual plot reveals spatial dependence with no obvious pattern of anisotropy. Thus a multiscale DCT model with latitude and longitude as predictors seem to be a desirable model for this data.

The proposed MDCT model for the training data uses  $R = 3$  resolutions with the first resolution having  $J(1) = 100$  knots. To minimize edge effect, some knots are also kept inside land. Similar to the simulation studies, we proposed a  $IG(2,1)$  prior on  $\sigma^2$  and multiscale tree shrinkage prior on  $\beta_j^r$ 's with the same hyper-parameters described in (11). We implement Algorithm 1 and run it for 2000 iterations to discover that  $\eta = 1$  appears overwhelmingly among 2000 iterations. Thus to reduce unnecessary storage complexity and to speed up the computations for a dataset of this scale, we run the rest of the iterations with  $\eta = 1$ . The model thereafter is run for 5000 iterations, convergence diagnostics is performed with the `coda` package in R which indicates that 2000 iterations are sufficient as burn-in, to achieve practical convergence. As competitors to the MDCT, we fitted LatticeKrig and LaGP to the data. MPP is computationally prohibitive for the size of the dataset and is omitted from the comparison.

Table 3: Mean squared prediction error (MSPE), length and coverage of 95% predictive intervals of MDCT, MDCT(1), MDCT(2), laGP and LatticeKrig.

	MDCT	MDCT(1)	MDCT(2)	laGP	LatticeKrig
MSPE	0.18	0.52	0.36	0.11	0.10
Length of 95% PI	2.49	2.38	2.42	1.26	0.48
Coverage of 95% PI	0.98	0.95	0.97	0.93	0.63

The predictive power of the proposed model, along with that of its competitors, is assessed based on mean squared prediction error (MSPE), coverage and length of 95% predictive intervals. The non-spatial model and MDCT yield MSPE 1.34 and 0.18 respectively. The dramatic improvement in MSPE due to the inclusion of a spatial structure, that is evident from Table 3 corroborates the fact that there is a strong spatial dependence in the field that can not be explained by a linear effect of longitude and latitude. From the results in Table 3 we observe that LaGP has a little better predictive performance than MDCT. The smallest MSPE in the table corresponds to LatticeKrig. Nevertheless, we observe that LatticeKrig suffers heavily in characterizing predictive uncertainty. Overall, MDCT turns out to be a competitive performer in predictive inference. Moreover, greater flexibility offered by MDCT in terms of modeling non-Gaussian spatial data has already been discussed in Section 6.

Predictive surfaces in Figure 7 further corroborate this fact. Importantly, even with non-parallel implementation MDCT takes about 26 seconds to run one iteration. As shown in Figure 5, the computation time can be reduced by multiple folds through efficient parallel implementation. On the contrary, even the frequentist implementation of LatticeKrig takes about 2 hours. Fitting MDCT beyond  $R = 3$  unnecessarily exacerbates computational burden with minimal improvement of inferential and predictive performance.

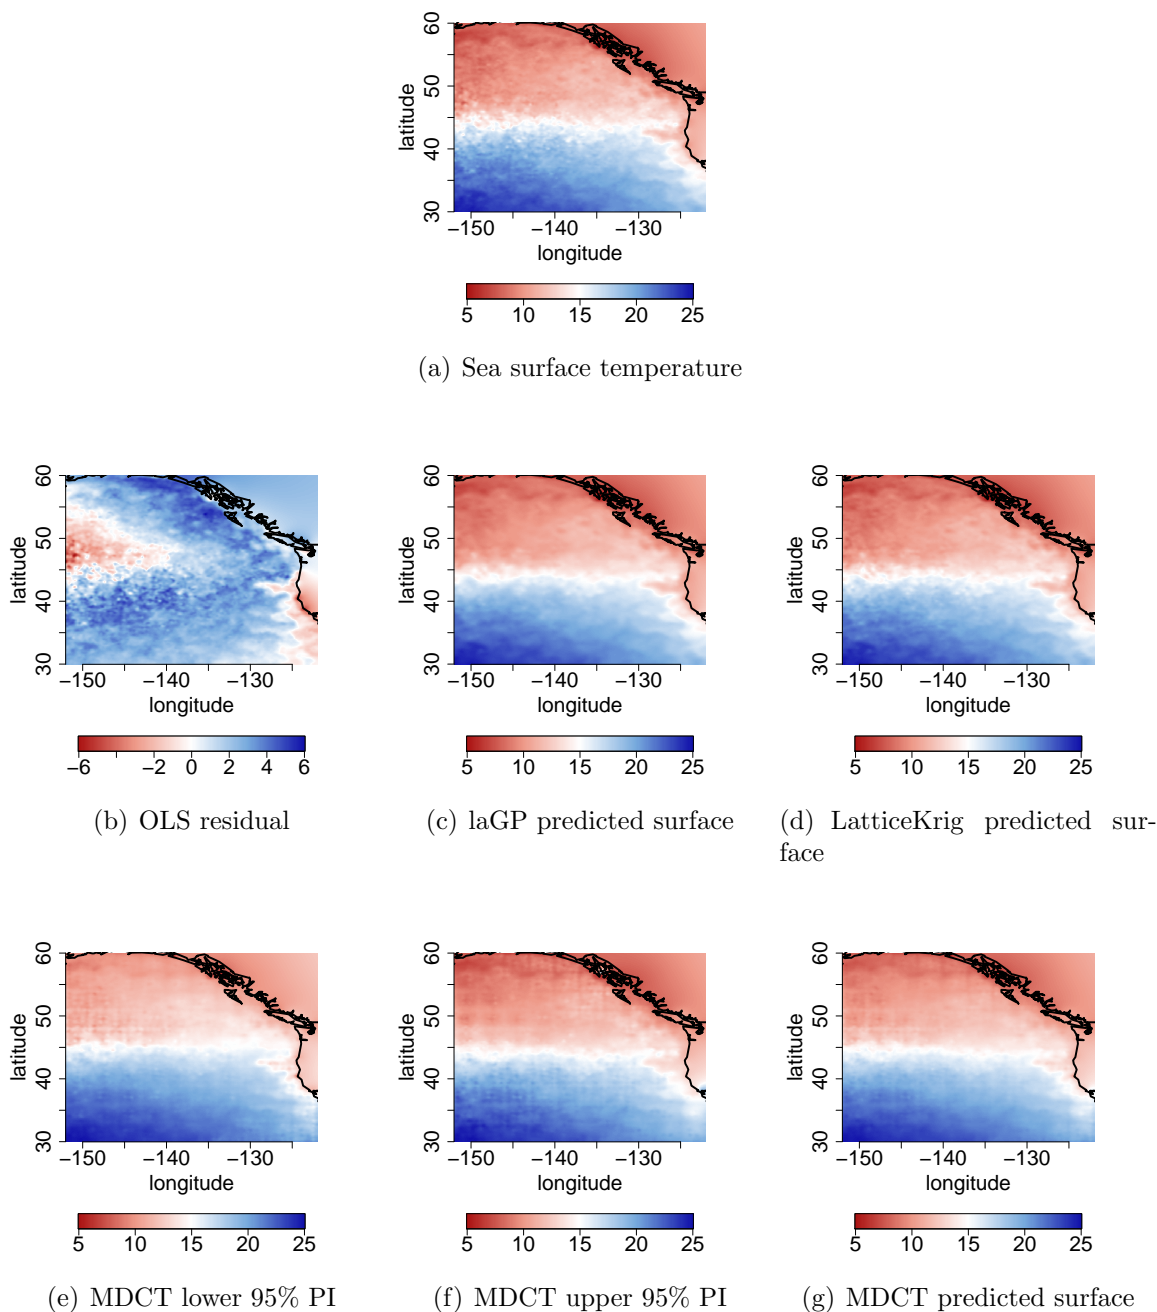


Figure 7: (a) Sea surface temperature in October 2016 for a portion of the North Pacific. (b) Estimated OLS residual from the non-spatial. Panels (c), (d) and (g) show the estimated mean predictive surfaces for three competing models. Figures (e) and (f) present the point-wise predictive bands for the MDCT. Temperatures are in degree centigrade.

## 7 Conclusion

This article proposes a novel multiscale kriging model for spatial datasets. The model writes the unknown spatial surface as an independent sum of processes at different scales and is



able to approximate a broad class of both 1D and 2D spatial processes of various degree of smoothness. One key ingredient of our mutiscale model is the kernel convolution with a compactly supported kernel of minimal degree and knots placed in a regular grid at every resolution. Theoretically, it allows us to completely characterize the space of functions generated from the mutiscale spatial model. Another important contribution of the current article is to propose a new class of *multiscale tree shrinkage prior* distribution for the basis coefficients. The construction of a tree shrinkage prior is governed by the consideration that as the model moves to the higher resolutions, more and more basis coefficients become irrelevant. The idea of multiscale tree shrinkage prior is novel in statistics and can find applications outside of spatial analysis.

Besides the important methodological and theoretical contributions that the proposed model entails, there is an equally important contribution in computational efficiency for massive datasets. The research on multiscale spatial models is largely motivated by the quest to build a complex and flexible spatial model that allows accurate spatial inference and prediction for massive datasets and yet allows rapid Bayesian computation. The compactly supported kernel together with the multiscale shrinkage priors allow easy MCMC of the model parameters. We develop a strategy for posterior computation within our modeling paradigm that ensures computation of the model parameters locally. More specifically, our strategy requires inverting a large number of  $[\frac{(2d)^R - 1}{2d - 1}] \times [\frac{(2d)^R - 1}{2d - 1}]$  matrices in parallel at every MCMC iteration, leading to efficient computation for Bayesian spatial models. We have demonstrated utility of MDCT both in spatial models with Gaussian error and in binary spatial regression. As mentioned in the article, the state-of-the-art competitors such as LatticeKrig and LaGP are implemented only for Gaussian data.

Several interesting new directions open up from this article. First of all, the current framework of mutiscale Bayesian modeling of spatial datasets can readily be extended to spatio-temporal datasets. Secondly, the recent idea of spatial meta kriging (Guhaniyogi and Banerjee, 2016) allows scalability by fitting a spatial model independently on partitions of a big data followed by combining the inferences. It is established in this article that the proposed multiscale framework can scale up to  $\approx 1$ -2 million spatial locations, but may struggle with tens of million. If we have resources to run on  $\approx 15$  different subsets, then SMK combined with our approach can yield full Bayesian inference on  $\approx 30$  million locations. Finally, this article proposes one specific rectangular partition of the domain. There is a scope of future research as to how adaptive partitioning of the domain is to be implemented using techniques such as the Voronoi tessellation. Adaptive partitioning with the appropriate placement of knots might significantly reduce the number of knots required to yield acceptably accurate inference. We will explore these approaches in future.

## Appendix

### Proof of Theorem 2.1:

Use the fact that  $\kappa$  is a compactly supported polynomial of minimal degree for two dimensions that possesses continuous derivatives upto second order. By Theorem 10.10 and 10.35 in

Wendland (2004), we obtain that the Fourier transform of  $\kappa$ , denoted by  $\hat{\kappa}$  satisfies

$$c_1(1 + \|\omega\|_2)^{-d-3} \leq \hat{\kappa}(\omega) \leq c_2(1 + \|\omega\|_2)^{-d-3},$$

for some  $c_1, c_2 > 0$ . The result now follows using Corollary 10.13 in Wendland (2004).

**Proof of Theorem 4.1:**

We begin by stating and proving a lemma that will be useful in the proof of the theorem.

**Lemma 7.1** *Consider a ball of radius  $\delta$  around  $(w_0, \sigma_0^2)$  given by*

$$B_\delta(w_0, \sigma_0^2) = \left\{ (w, \sigma^2) : \|w - w_0\|_\infty < \delta, \left| \frac{\sigma^2}{\sigma_0^2} - 1 \right| < \delta \right\}.$$

Then  $\pi(B_\delta(w_0, \sigma_0^2)) > 0$ , for all  $\delta > 0$ .

**Proof** Since  $w_0 \in \Theta_c$ ,  $\exists w_*(\mathbf{s}) = \sum_{r=1}^{R^*} \sum_{j=1}^{J(r)} K(\mathbf{s}, \mathbf{s}_j^{r*}, \phi_r) \beta_j^{r*}$ , s.t.  $\|w_* - w_0\|_\infty < \delta/2$ . Note that  $K(\cdot, \cdot, \phi_r)$  is a continuous function on a compact set  $\mathcal{D}$ , implying  $K(\cdot, \cdot, \phi_r)$  to be a uniformly continuous function. Thus,  $\exists M$ , s.t.  $M = \sup_{\mathbf{s} \in \mathcal{D}} \max_{r=1, \dots, R^*; j=1:J(r)} |K(\mathbf{s}, \mathbf{s}_j^{r*}, \phi_r)|$ .

Assume further that  $\eta = \sum_{r=1}^{R^*} \sum_{j=1}^{J(r)} |\beta_j^{r*}|$ . Since  $K$  is uniformly continuous, one can choose  $\mathbf{s}_j^r$ 's such that  $\sup_{\mathbf{s} \in \mathcal{D}} |K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) - K(\mathbf{s}, \mathbf{s}_j^{r*}, \phi_r)| < \frac{\delta}{4\eta \sum_{r=1}^{R^*} J(r)}$ . Define the set

$$\mathcal{I} = \left\{ \{\beta_j^r\} : |\beta_j^r - \beta_j^{r*}| < \frac{\delta}{4M \sum_{r=1}^{R^*} J(r)} \right\}.$$

Clearly, for the set of all  $w(\mathbf{s}) = \sum_{r=1}^{R^*} \sum_{j=1}^{J(r)} K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) \beta_j^r$ , with  $\mathbf{s}_j^r$  is chosen as above and  $\beta_j^r$  chosen from  $\mathcal{I}$ , we have

$$\begin{aligned} |w_0(\mathbf{s}) - w(\mathbf{s})| &\leq |w_0(\mathbf{s}) - w_*(\mathbf{s})| + |w_*(\mathbf{s}) - w(\mathbf{s})| \\ &\leq \frac{\delta}{2} + \sum_{r=1}^{R^*} \sum_{j=1}^{J(r)} |K(\mathbf{s}, \mathbf{s}_j^r, \phi_r)| |\beta_j^r - \beta_j^{r*}| + \sum_{r=1}^{R^*} \sum_{j=1}^{J(r)} |\beta_j^{r*}| |K(\mathbf{s}, \mathbf{s}_j^r, \phi_r) - K(\mathbf{s}, \mathbf{s}_j^{r*}, \phi_r)| \\ &\leq \frac{\delta}{2} + \frac{\delta M \sum_{r=1}^{R^*} J(r)}{4M \sum_{r=1}^{R^*} J(r)} + \frac{\delta \eta \sum_{r=1}^{R^*} J(r)}{4\eta \sum_{r=1}^{R^*} J(r)} = \delta. \end{aligned}$$

Thus  $\mathcal{I} \times \left\{ \sigma^2 : \left| \frac{\sigma^2}{\sigma_0^2} - 1 \right| < \delta \right\} \subseteq B_\delta(w_0, \sigma_0^2)$ . Since, the prior on all  $\beta_j^r$  are continuous on the entire real line and the prior on  $\sigma^2$  is also continuous on  $\mathcal{R}^+$ , it trivially holds that  $\pi(B_\delta(w_0, \sigma_0^2)) \geq \pi\left(\mathcal{I} \times \left\{ \sigma^2 : \left| \frac{\sigma^2}{\sigma_0^2} - 1 \right| < \delta \right\}\right) > 0$ . This concludes the proof of the lemma.

We will now proceed with the proof of Theorem 4.1. Our aim is to check that all conditions of Theorem in Choi and Schervish (2007) are satisfied. Let  $H_i = \frac{N(y_i|w_0(\mathbf{s}_i), \sigma_0^2)}{N(y_i|w(\mathbf{s}_i), \sigma^2)}$ , and  $K_i(w, w_0) =$

$E(H_i)$  and  $V_i(w, w_0) = \text{Var}(H_i)$ . It is easy to check that (Choi and Schervish (2007))

$$K_i(w, w_0) = \frac{1}{2} \log \frac{\sigma^2}{\sigma_0^2} - \frac{1}{2} \left( 1 - \frac{\sigma_0^2}{\sigma^2} \right) + \frac{1}{2} \frac{(w(\mathbf{s}) - w_0(\mathbf{s}))^2}{\sigma^2}$$

$$V_i(w, w_0) = \frac{1}{2} \left( \frac{\sigma_0^2}{\sigma^2} - 1 \right)^2 + \frac{\sigma_0^4}{\sigma^4} (w(\mathbf{s}) - w_0(\mathbf{s}))^2.$$

Thus for every  $\epsilon > 0$ , there exists a  $\delta > 0$  such that  $(w(\cdot), \sigma^2) \in B_\delta(w_0, \sigma_0^2)$  implies  $K_i(w, w_0) < \epsilon, \forall i$  and  $\sum_{i=1}^{\infty} \frac{V_i(w, w_0)}{i^2} < \infty$ . Thus condition (i) is satisfied. Condition (ii), i.e. the prior positivity has already been proved to be satisfied by Lemma 7.1.

Finally, the condition of having an exponentially consistent sequence of tests follows along the same line as the proof of Theorem 2 in Choi and Schervish (2007). This concludes the theorem.

## References

- Albert, J. H. and S. Chib (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association* 88(422), 669–679.
- Armagan, A., D. B. Dunson, and J. Lee (2013). Generalized double Pareto shrinkage. *Statistica Sinica* 23(1), 119.
- Banerjee, S., B. P. Carlin, and A. E. Gelfand (2014). *Hierarchical modeling and analysis for spatial data*. Crc Press.
- Banerjee, S., A. E. Gelfand, A. O. Finley, and H. Sang (2008). Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 70(4), 825–848.
- Berliner, L. M., C. K. Wikle, and N. Cressie (2000). Long-lead prediction of Pacific ssts via Bayesian dynamic modeling. *Journal of climate* 13(22), 3953–3968.
- Carvalho, C. M., N. G. Polson, and J. G. Scott (2009). Handling sparsity via the horseshoe. In *AISTATS*, Volume 5, pp. 73–80.
- Choi, T. and M. J. Schervish (2007). On posterior consistency in nonparametric regression problems. *Journal of Multivariate Analysis* 98(10), 1969–1987.
- Clyde, M., H. Desimone, and G. Parmigiani (1996). Prediction via orthogonalized model mixing. *Journal of the American Statistical Association* 91(435), 1197–1208.
- Crainiceanu, C. M., P. J. Diggle, and B. Rowlingson (2008). Bivariate binomial spatial modeling of loa loa prevalence in tropical Africa. *Journal of the American Statistical Association* 103(481), 21–37.
- Cressie, N. and G. Johannesson (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 70(1), 209–226.
- Cressie, N. and C. K. Wikle (2015). *Statistics for spatio-temporal data*. John Wiley & Sons.
- Datta, A., S. Banerjee, A. O. Finley, and A. E. Gelfand (2015). Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association* (just-accepted), 00–00.
- Du, J., H. Zhang, V. Mandrekar, et al. (2009). Fixed-domain asymptotic properties of tapered maximum likelihood estimators. *the Annals of Statistics* 37(6A), 3330–3361.
- Eidsvik, J., B. A. Shaby, B. J. Reich, M. Wheeler, and J. Niemi (2014). Estimation and prediction in spatial models with block composite likelihoods. *Journal of Computational and Graphical Statistics* 23(2), 295–315.
- Finley, A. O., S. Banerjee, P. Waldmann, and T. Ericsson (2009). Hierarchical spatial modeling of additive and dominance genetic variance for large spatial trial datasets. *Biometrics* 65(2), 441–451.

- Furrer, R., M. G. Genton, and D. Nychka (2012). Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*.
- Gelfand, A. E., P. Diggle, P. Guttorp, and M. Fuentes (2010). *Handbook of spatial statistics*. CRC press.
- George, E. I. and R. E. McCulloch (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association* 88(423), 881–889.
- Geweke, J. (2004). Getting it right: Joint distribution tests of posterior simulators. *Journal of the American Statistical Association* 99(467), 799–804.
- Gramacy, R. B. and D. W. Apley (2015). Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics* 24(2), 561–578.
- Gramacy, R. B. and H. K. Lee (2012). Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*.
- Guhaniyogi, R. and S. Banerjee (2017). Meta kriging: scalable Bayesian modeling and inference for large spatial datasets. <https://www.soe.ucsc.edu/research/technical-reports/UCSC-SOE-17-03>.
- Guhaniyogi, R., A. O. Finley, S. Banerjee, and A. E. Gelfand (2011). Adaptive Gaussian predictive process models for large spatial datasets. *Environmetrics* 22(8), 997–1007.
- Guinness, J. (2016). Permutation methods for sharpening gaussian process approximations. *arXiv preprint arXiv:1609.05372*.
- Hans, C. (2009). Bayesian lasso regression. *Biometrika* 96(4), 835–845.
- Higdon, D. (1998). A process-convolution approach to modelling temperatures in the north Atlantic ocean. *Environmental and Ecological Statistics* 5(2), 173–190.
- Higdon, D. (2002). Space and space-time modeling using process convolutions. In *Quantitative methods for current environmental issues*, pp. 37–56. Springer.
- Katzfuss, M. (2013). Bayesian nonstationary spatial modeling for very large datasets. *Environmetrics* 24(3), 189–200.
- Katzfuss, M. (2016). A multi-resolution approximation for massive spatial datasets. *Journal of the American Statistical Association* (just-accepted).
- Kaufman, C. G., M. J. Schervish, and D. W. Nychka (2008). Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association* 103(484), 1545–1555.
- Lemos, R. T. and B. Sansó (2006). Spatio-temporal variability of ocean temperature in the Portugal current system. *Journal of Geophysical Research: Oceans* 111(C4).

- Lemos, R. T. and B. Sansó (2009). A spatio-temporal model for mean, anomaly, and trend fields of north Atlantic sea surface temperature. *Journal of the American Statistical Association* 104(485), 5–18.
- Nychka, D., S. Bandyopadhyay, D. Hammerling, F. Lindgren, and S. Sain (2015). A multi-resolution Gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics* 24(2), 579–599.
- Park, C. and D. Apley (2017). Patchwork kriging for large-scale gaussian process regression. *arXiv preprint arXiv:1701.06655*.
- Park, T. and G. Casella (2008). The bayesian lasso. *Journal of the American Statistical Association* 103(482), 681–686.
- Pillai, N. S. (2008). *Lévy random measures: posterior consistency and applications*. Ph. D. thesis, Citeseer.
- Polson, N. G. and J. G. Scott (2012). Local shrinkage rules, Lévy processes and regularized regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 74(2), 287–311.
- 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(2), 319–392.
- Shaby, B. and D. Ruppert (2012). Tapered covariance: Bayesian estimation and asymptotics. *Journal of Computational and Graphical Statistics* 21(2), 433–452.
- Stein, M. L. (2007). Spatial variation of total column ozone on a global scale. *The Annals of Applied Statistics*, 191–210.
- Stein, M. L. (2012). *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media.
- Stein, M. L. (2014). Limitations on low rank approximations for covariance matrices of spatial data. *Spatial Statistics* 8, 1–19.
- Stein, M. L., Z. Chi, and L. J. Welty (2004). Approximating likelihoods for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 66(2), 275–296.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Vecchia, A. V. (1988). Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society. Series B (Methodological)*, 297–312.
- Wendland, H. (2004). *Scattered data approximation*, Volume 17. Cambridge university press.

- Zhang, R., C. D. Lin, and P. Ranjan (2016). Local Gaussian process model for large-scale dynamic computer experiments. *arXiv preprint arXiv:1611.09488*.
- Zou, H. and T. Hastie (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67(2), 301–320.