



Published in final edited form as:

Technometrics. 2018 ; 60(4): 415–429. doi:10.1080/00401706.2018.1437476.

Permutation and Grouping Methods for Sharpening Gaussian Process Approximations

Joseph Guinness

North Carolina State University, Department of Statistics

Abstract

Vecchia's approximate likelihood for Gaussian process parameters depends on how the observations are ordered, which has been cited as a deficiency. This article takes the alternative standpoint that the ordering can be tuned to sharpen the approximations. Indeed, the first part of the paper includes a systematic study of how ordering affects the accuracy of Vecchia's approximation. We demonstrate the surprising result that random orderings can give dramatically sharper approximations than default coordinate-based orderings. Additional ordering schemes are described and analyzed numerically, including orderings capable of improving on random orderings. The second contribution of this paper is a new automatic method for grouping calculations of components of the approximation. The grouping methods simultaneously improve approximation accuracy and reduce computational burden. In common settings, reordering combined with grouping reduces Kullback-Leibler divergence from the target model by more than a factor of 60 compared to ungrouped approximations with default ordering. The claims are supported by theory and numerical results with comparisons to other approximations, including tapered covariances and stochastic partial differential equations. Computational details are provided, including the use of the approximations for prediction and conditional simulation. An application to space-time satellite data is presented.

1 Introduction

The Gaussian process model has become very popular for the analysis of time series, functional data, spatial data, spatial-temporal data, and the output from computer experiments. Likelihood-based methods for estimating Gaussian process covariance parameters were popularized by the work of Mardia and Marshall (1984), but it was quickly realized that the $O(n^2)$ memory and $O(n^3)$ flop burden made these methods infeasible for large datasets, necessitating the use of computationally efficient approximations. This article studies one of the earliest such approximations, due to Vecchia (1988), which has a number of advantages. First, the approximation is embarrassingly parallel, making it a good candidate for implementation on high performance computing systems. Second, the approximation corresponds to a valid multivariate normal distribution, meaning that the approximation can be used to generate ensembles of conditional simulations to characterize joint uncertainties in predictions. Third, the approximation enjoys the desirable statistical property that maximizing it corresponds to solving a set of unbiased estimating equations (Stein et al., 2004). Lastly, the approximation is very accurate; we demonstrate that with the improvements described in this article, it far outperforms state-of-the-art methods such as

stochastic partial differential equation (SPDE) approximations Lindgren et al. (2011) and covariance tapering (Furrer et al., 2006; Kaufman et al., 2008).

Suppose that the n observation locations are labeled 1 through n , as in $x_1, \dots, x_n \in \mathbb{R}^d$. Define $y_j := y(x_j) \in \mathbb{R}$ to be the observation associated with location x_j and the vector $y = (y_1, \dots, y_n)$. In a Gaussian process, the data y are modeled as $Y \sim \mathcal{N}(\mu, \Sigma_\theta)$, where Σ_θ is a covariance matrix with (i, j) entry determined by covariance function $K_\theta(x_i, x_j)$ depending on a vector of covariance parameters θ . Let $\tau : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ be a permutation of the integers 1 through n , and define the permuted vector y^τ , where $y_i^\tau = y_{\tau(i)}$. For any permutation τ , the joint density for the observations can be written as a product of conditional densities

$$p_\theta(y_1, \dots, y_n) = p_\theta(y_1^\tau) \prod_{i=2}^n p_\theta(y_i^\tau | y_1^\tau, \dots, y_{i-1}^\tau),$$

reflecting the invariance of the joint density to arbitrary relabeling of the observations. Vecchia’s approximation replaces the complete conditioning vectors $(y_1^\tau, \dots, y_{i-1}^\tau)$, with a subvector. Specifically, let $\{j_{i1}, \dots, j_{im_i}\}$ be a set of integers between 1 and $i - 1$, and define the approximation

$$p_{\theta, \tau, J}(y_1, \dots, y_n) = p_\theta(y_1^\tau) \prod_{i=2}^n p_\theta(y_i^\tau | y_{j_{i1}}^\tau, \dots, y_{j_{im_i}}^\tau), \quad (1)$$

where $J = \{J_1, \dots, J_n\}$ is collection of sets $J_i := \{j_{i1}, \dots, j_{im_i}, i\}$, which we refer to as the neighbors of observation i (by convention every observation neighbors itself). Since Vecchia’s approximation is defined as an ordered sequence of valid conditional distributions, $p_{\theta, \tau, J}$ forms a valid joint distribution. Also, its definition as a product of conditional distributions allows the approximation to be easily parallelized.

Clearly, the quality of the approximation depends on the choice of J since the approximate likelihood differs from the exact likelihood if any $m_i < i - 1$. The quality of the approximation also depends on the permutation τ , a fact acknowledged in the literature, but not yet carefully explored until now. Banerjee et al. (2014) outline some of the common criticisms of Vecchia’s approximation,

“However, the approach suffers many problems. First, it is not formally defined. Second, it will typically be sequence dependent, though there is no natural ordering of the spatial locations. Most troubling is the arbitrariness in the number of and choice of ‘neighbors.’ Moreover, perhaps counter-intuitively, we can not merely select locations close to s_j (as we would have with the full data likelihood) in order to learn about the spatial decay in dependence for the process. So, altogether, we do not see such approximations as useful approach.”

This paper takes the stance that viewing order dependence as an aspect of the approximation that can be tuned—rather than as an inherent deficiency—is a fruitful avenue for improving the approximation. Datta et al. (2016a) conducted a simulation study to compare three coordinate-based orderings and concluded that the inferences are “extremely robust to the ordering of the locations.” On the contrary, default orderings based on sorting the locations on a coordinate are often badly suboptimal compared to even a completely random ordering of the points. On two-dimensional domains, the more carefully constructed maximum minimum distance ordering is a further improvement and can achieve greater than 99% relative efficiency for estimating covariance parameters with as few as 30 neighbors, chosen to be simply the 30 nearest previous points. These results address most of the concerns raised by Banerjee et al. (2014). Order dependence can be exploited to improve the approximation. For maximum minimum distance ordering, a simple choice of nearest neighbors is effective for this ordering. We prove that the quality of the model approximation is nondecreasing as the number of neighbors increases, so that the choice of the number of neighbors is governed by a natural tradeoff between computational efficiency and model approximation.

In addition to the results on orderings, this paper describes how the approximation can be computed more efficiently, both in terms of memory burden and floating point operations, by grouping the observations into blocks and evaluating each group’s contribution to the likelihood simultaneously. The grouped version of the approximation is particularly interesting because not only does it reduce the computational burden, but it is provably guaranteed to improve the model approximation. A computationally efficient algorithm for grouping the observations is described and implemented, and numerical studies supporting the theoretical results are presented.

The consideration of arbitrary orderings presents some computational issues that are not problematic for coordinate-based orderings. First, obviously, is the issue of how to find the orderings. The maximum minimum distance ordering requires $O(n^3)$ floating point operations. To address this issue, we introduce an $O(n \log n)$ algorithm for finding an ordering that mimics the salient features of the maximum minimum distance ordering. Second, a naive search for ordered nearest neighbors requires $O(n^2 \log n)$ flops. There are well-known computationally efficient methods for finding nearest neighbors, and this paper describes an adaptation of those methods to the case when the neighbors must come from earlier in the ordering. We also describe methods for profiling out linear mean parameters and using the approximation for spatial prediction and for efficiently drawing from the conditional distribution at a set of unobserved locations given the data, which is a useful way to quantify joint uncertainty in predicted values.

Several articles have investigated properties and extensions of Vecchia’s approximation. Pardo- Iguzquiza and Dowd (1997) describe Fortran 77 software for computing the approximate likelihood and compare a random ordering versus a sorted coordinate ordering on a small dataset of size 41. Stein et al. (2004) describe how Vecchia’s approximation can be extended to residual/restricted maximum likelihood estimation, how blocking can be used to speed the computations, and how conditioning on nearest neighbors can be suboptimal when points are ordered according to the coordinates. Sun and Stein (2016) use Vecchia’s

approximation to define several unbiased estimating equations for covariance parameters. Datta et al. (2016a) use Vecchia's approximation as part of a hierarchical Bayesian specification of spatial processes, and Datta et al. (2016b) discuss interpretation of the approximation as an approximation to the inverse Cholesky factor of the covariance matrix and apply it to multivariate spatial data. Vecchia's approximation has been used in various applications, including for seismic data (Eide et al., 2002) and space-time SPDE models (Jones and Zhang, 1997).

The paper is organized as follows. Section 2 outlines formal definitions for the orderings. Section 3 presents the computational and theoretical results related to grouping the observations, establishing that grouping observations simultaneously reduces computational effort and improves the model approximations. Section 4 addresses additional computational issues described above. Section 5 contains numerical and timing experiments studying the relative efficiencies for various orderings, Kullback-Leibler divergence for various orderings and for other proposed Gaussian process approximations, and timing results. Section 6 includes an application of the methods to space-time satellite data, and the paper concludes with a discussion in Section 7.

2 Definitions of Orderings

In the numerical analysis literature on sparse matrix factorizations, it is widely recognized that row- column reordering schemes for large sparse symmetric positive definite matrices are essential for increasing the sparsity of the Cholesky factor (Saad, 2003). Finding an optimal such ordering is an NP-complete problem (Yannakakis, 1981), and so the algorithms in use are necessarily heuristic, but heuristic algorithms, such as the approximate minimum degree algorithm (Amestoy et al., 1996), have proven to be effective. The goals in Gaussian process approximation are related in that we search for an ordering that produces an approximately sparse inverse Cholesky factor. However, the problem at hand here is more difficult and perhaps less well-defined than sparse matrix factorization; our task is to find the best reordering of observations that produces accurate approximations to the joint density or an approximate likelihood function that delivers efficient parameter estimates, a criterion that depends on the derivative of the approximate and exact likelihood functions with respect to the covariance parameters. Thus, it is extremely unlikely that we will be able to find "optimal" orderings for large datasets. Nevertheless, as in the sparse matrix case, this paper shows that heuristically motivated orderings can offer significant improvements in statistical efficiency and model approximation over default orderings.

It appears that the default choice for Vecchia's approximation is to order the points by sorting on one of the coordinates. This is the ordering used in Vecchia (1988), Sun and Stein (2016), and Datta et al. (2016a). Stein et al. (2004) use an ordering based on sorting the points on the sum of their coordinates, which is equivalent to ordering on one of the coordinates in a system rotated by $\pi/4$. We refer to such orderings as *sorted coordinate* orderings. Sorted coordinate orderings are based on a heuristic from one-dimensional examples that each location can be separated by previous locations in the ordering by its nearest neighbors.

The numerical studies in Section 5 indicate that the following ordering scheme is effective for Matern covariance models in two dimensions. This ordering selects a point in the center first—the center being the mean location or some other measure of centrality, generically denoted as \bar{x} —then sequentially picks the next point to have maximum minimum distance to all previously selected points, that is

$$\begin{aligned}\tau(1) &= \arg \min_{i \in 1, \dots, n} \|x_i - \bar{x}\|, \\ \tau(j) &= \arg \max_{i \notin \tau(1), \dots, \tau(j-1)} \min_{k \in 1, \dots, j-1} \|x_i - x_{\tau(k)}\|, \quad j > 1.\end{aligned}$$

The result is that for every $k = 1, \dots, n$, the first k points form a space-covering set, none of which are too near each other. We refer to this ordering as a *maximum minimum distance* (MMD) ordering, and any approximation to it as an approximate MMD (AMMD) ordering. The MMD ordering is based on a heuristic of making sure that each location is surrounded by previous locations in the ordering.

The numerical studies in more than two dimensions indicate that it can be beneficial to sort the points by their distance to some point in the domain. For example, the points can be sorted based on their distance to \bar{x} . We use *middle out* ordering to refer to sorting based on distance to the center. The heuristic for middle out is similar to the sorted coordinate heuristic, with the difference that previous points fall inside a sphere with radius smaller than the radius of the current point. Finally, a *completely random* ordering is a draw from the uniform distribution on all permutations. Random orderings do not have a heuristic but tend to give orderings with the same surrounding heuristic as MMD and outperform sorted coordinate orderings in many cases. Examples of the four orderings are given in Figure 1.

3 Automatic Grouping Methods

In this section, a grouped version of Vecchia’s approximation is described and proven to be no worse (and possibly better) than its ungrouped counterpart in terms of Kullback-Leibler (KL) divergence. Moreover, this section outlines how an entire group’s contribution to the likelihood can be computed simultaneously, and it is demonstrated that if the grouping is chosen well, the simultaneous computation is lighter, both in terms of floating point operations and memory, which is important for implementation on shared memory parallel systems. Essentially, grouping gives us better approximations for free.

Stein et al. (2004) considered a blockwise version of Vecchia’s likelihood and demonstrated that blocking can reduce the computational burden. The grouping methods developed here differ in that they operate on the ordering and neighbor sets after they have been defined. This gives us the freedom to first choose an advantageous ordering and neighbor sets, and then use grouping to reduce the computational effort and improve the model accuracy. The grouping scheme here is also more general in that observations within a group are not required to be ordered contiguously, nor are observations within one group required to condition on all observations in another group.

Let $B = \{B_1, \dots, B_K\}$ be a partition of $\{1, \dots, n\}$, representing a grouping of the observations into blocks, and define $U_k = \bigcup_{b \in B_k} J_b$, the union of all neighbors of all indices B_k , representing the set of all neighbors of all observations in block B_k . Suppose that $i \in B_k$, and define \bar{J}_i as follows:

$$\bar{J}_i = \{j \in U_k : j \leq i\},$$

so that \bar{J}_i is the union of all neighbors of the observations with index in B_k , subject to the requirement that all elements of \bar{J}_i must be less than or equal to i . Thus \bar{J}_i is a set of integers between 1 and i that includes i . Therefore \bar{J}_i conforms to our convention and $p_{\theta, \tau, \bar{J}}$ is an approximation in the form of (1). Note that \bar{J}_i depends on the original neighbor choices J_1, \dots, J_n and the partition B . I refer to $p_{\theta, \tau, \bar{J}}$ as the grouped version of the ungrouped approximation $p_{\theta, \tau, J}$. Figure 2 gives an example of how \bar{J}_i is constructed for a group containing two points.

3.1 Improved Model Approximations

To gain an intuition for why grouping observations improves the accuracy of the model approximation, I briefly review one interpretation of the inverse Cholesky factor of a covariance matrix (Pourahmadi, 1999). Suppose $Y = (Y_1, \dots, Y_n)^T$ is a mean-zero multivariate normal vector with $E(YY^T) = \Sigma$, which has Cholesky decomposition $\Sigma = LL^T$. Let $\Gamma = L^{-1}$. The transformation $Z = \Gamma Y$ is decorrelating and can be interpreted as follows; for any $i = 2, \dots, n$,

$$Z_i = \sum_{j=1}^i Y_j = \Gamma_{ii} \left[Y_i - \sum_{j=1}^{i-1} \left(\frac{-\Gamma_{ij}}{\Gamma_{ii}} \right) Y_j \right],$$

is the standardized residual of the projection of Y_i onto (Y_1, \dots, Y_{i-1}) , Γ_{ii} is the reciprocal of the residual standard deviation, and $-\Gamma_{ij}/\Gamma_{ii}$ are the coefficients of the projection. The coefficients can also be interpreted as those defining the best linear unbiased predictor (BLUP) of Y_j using (Y_1, \dots, Y_{i-1}) , an interpretation which can be extended to the case of a linear mean function (Stein et al., 2004). Defining $z = \Gamma y$, the conditional density for Y_i given Y_1, \dots, Y_{i-1} can be written as

$$\log p_{\theta}(y_i | y_1, \dots, y_{i-1}) = -\frac{1}{2} \log 2\pi + \log \Gamma_{ii} - \frac{1}{2} z_i^2,$$

and thus the i th row of Γ encodes the conditional density for Y_i given Y_1, \dots, Y_{i-1} .

Vecchia's approximation replaces Γ with a sparse approximation $\tilde{\Gamma}$, since Y_i is projected onto a subset of (Y_1, \dots, Y_{i-1}) , and thus $\tilde{\Gamma}_{ij}$ is zero if Y_j is not in the conditioning set for Y_i . Because of the BLUP interpretation, $\tilde{\Gamma}_{ii} \leq \Gamma_{ii}$ for every i . Taking this a step further, suppose

$\tilde{\Gamma}^1$ and $\tilde{\Gamma}^2$ represent two approximations that use the same ordering with different conditioning sets given by J^1 and J^2 . If J_i^1 is contained in J_i^2 for every i , then $\tilde{\Gamma}_{ii}^1 \leq \tilde{\Gamma}_{ii}^2$, because the variance of the BLUP cannot increase when additional variables are added to the conditioning set.

The Kullback-Leibler (KL) divergence from p_i to p_0 is defined as $KL(p_0 \parallel p_i) = E(\log(p_0/p_i))$, where the expectation is with respect to p_0 . The following theorem establishes that the KL divergence is nonincreasing as observations are added to the conditioning sets. This result is needed to show that the grouped approximation is no worse than its ungrouped counterpart (and could be better).

Theorem 1.—If $J_i^1 \subset J_i^2$ for every $i = 1, \dots, n$, then $KL(p_\theta \parallel p_{\theta, \tau, J^2}) \leq KL(p_\theta \parallel p_{\theta, \tau, J^1})$.

Proof: Let Σ^1 and Σ^2 be the covariance matrices under p_{θ, τ, J^1} and p_{θ, τ, J^2} . The difference between the KL-divergences is given by

$$KL(p_\theta \parallel p_{\theta, \tau, J^1}) - KL(p_\theta \parallel p_{\theta, \tau, J^2}) = \frac{1}{2}(E(Y^T(\Sigma^1)^{-1}Y) - E(Y^T(\Sigma^2)^{-1}Y) - \log \det(\Sigma^1)^{-1} + \log \det(\Sigma^2)^{-1}).$$

Let Γ^k be the inverse Cholesky factor of Σ^k . According to Vecchia’s approximation, Γ^k is sparse and has entries constructed as follows. Let A be the covariance matrix for the vector

$$(Y_{j_{i1}}, \dots, Y_{j_{im}}, Y_i)^T,$$

and let C be the inverse Cholesky factor of A . Row $m + 1$ (the last row) of C contains the coefficients for the projection of Y_i onto $(Y_{j_{i1}}, \dots, Y_{j_{im}})$. The i th row of Γ^k has entries

$$\Gamma_{ij_{i\ell}}^k = C_{m+1, \ell}, \Gamma_{ii}^k = C_{m+1, m+1}, \Gamma_{ij}^k = 0 \text{ otherwise.}$$

Define $Z^k = \Gamma^k Y$. Then Z_i^k is

$$Z_i^k = \sum_{j=1}^i \Gamma_{ij}^k Y_j = \sum_{\ell=1}^m C_{m+1, \ell} Y_{j_{i\ell}} + C_{m+1, m+1} Y_i = C_{m+1, m+1} \left[Y_i - \sum_{\ell=1}^m \left(\frac{-C_{m+1, \ell}}{C_{m+1, m+1}} \right) Y_{j_{i\ell}} \right].$$

From the last expression it is evident that Z_i^k is the residual from the BLUP of Y_i given $(Y_{j_{i1}}, \dots, Y_{j_{im}})$, scaled by the standard deviation of that residual, and so Z_i^k has mean zero and variance 1. This ensures that $E(Y^T(\Sigma^k)^{-1}Y) = E(\sum_i (Z_i^k)^2) = n$, and thus both quadratic form terms are equal to n and cancel each other.

The determinant of Σ_k^{-1} can be computed by taking the square of the product of the entries of T^k . Since $J_i^1 \subset J_i^2$ for every i , we have that $0 \leq \Gamma_{ii}^1 \leq \Gamma_{ii}^2$ for every i , so $\log \det(\Sigma^1)^{-1} < \log \det(\Sigma^2)^{-1}$, and thus $KL(p_\theta \| p_{\theta, \tau, J^2}) \leq KL(p_\theta \| p_{\theta, \tau, J^1})$, as desired.

Corollary 1 follows from Theorem 1 when we note that, by definition, \bar{J}_i contains J_i , since $i \in B_k$ and all elements of J_i are less than or equal to i .

Corollary 1.—The grouped approximation has smaller KL divergence than its ungrouped counterpart, i.e.

$$KL(p_\theta \| p_{\theta, \tau, \bar{J}}) \leq KL(p_\theta \| p_{\theta, \tau, J})$$

3.2 Simultaneous Computation

Recall that U_k is the set of all neighbors of observations in block B_k . Let A^k be the covariance matrix for observations with indices in U_k , with observations ordered increasing in their indices. If we define C^k to be the inverse Cholesky factor of A^k , each row of C^k encodes the conditional distribution of an observation with index in U_k given those ordered before in U_k . Since \bar{J}_i contains precisely those indices ordered before i in U_k , we can calculate the contribution of observations with indices $i \in B_k$ to $p_{\theta, \tau, \bar{J}}$ simultaneously, by forming and factoring a single covariance matrix A^k . The total computational cost for the grouped version is on the order

$$\text{memory: } \sum_{k=1}^K (\#U_k)^2, \quad \text{flops: } \sum_{k=1}^K (\#U_k)^3,$$

where $(\#U_k)$ is the size of the set U_k . If a group of observations share a substantial number of neighbors, it is possible that these quantities are less than the memory and flop burden of the ungrouped approximation— $n(m+1)^2$ and $n(m+1)^3$ —and so computational savings can be gained by computing each group's contribution to the likelihood in this simultaneous way.

To illustrate this point, consider the canonical ordering and two observations y_i and y_{i+2} with $J_i = (i-m, \dots, i)$ and $J_{i+2} = (i-m+2, \dots, i+2)$, so that each observation conditions on the m previous ones. The computational cost of factoring the covariance matrices for the observations and their respective neighbors is $1/3(m+1)^3$ floating point operations each, for a total of $2/3(m+1)^3$ operations, and $(m+1)^2$ memory units each, for a total of $2(m+1)^2$ memory units. Now, suppose that $B_1 = \{i, i+2\}$, and thus $U_1 = \{i-m, \dots, i+2\}$, $\bar{J}_i = J_i, \bar{J}_{i+2} = \{i-m, \dots, i+2\}$, and A^1 is the $(m+3) \times (m+3)$ covariance matrix for $(y_{i-m}, \dots, y_{i+2})$. Row $m+1$ of C^1 encodes the conditional distribution of y_i given $(y_{i-m}, \dots, y_{i-1})$, and row $m+3$ of C^1 encodes the conditional distribution of y_{i+2} given $(y_{i-m}, \dots, y_{i+1})$, which is what is required for calculating the contribution from y_i and y_{i+2} to the grouped likelihood approximation. The computational cost of storing and factoring A^1 is $(m+3)^2$ memory units and $1/3(m+3)^3$ flops, so in addition to the increased accuracy of the

approximation (Corollary 1), we also achieve computational savings in flops if $m > 6$ and in memory if $m > 3$.

3.3 Grouping Algorithm

Since the groupings are a partition of the set of observations, it is of interest to find good partitions that improve the model and lessen computational cost. Here, we describe a fast greedy algorithm that depends only on the ordering and the neighbor sets. The algorithm starts with $B_k = \{k\}$, each index in its own block and proceeds to propose joining pairs of blocks.

Grouping algorithm:

```

Set  $B = \{\{1\}, \dots, \{n\}\}$ 
for  $\ell = 1, \dots, m$  do
  for  $i = 1, \dots, n$  do
    Identify the  $k$  for which  $i \in B_k$ 
    Identify the  $k'$  for which  $j_{i\ell} \in B_{k'}$ 
    if  $(\#(U_k \cup U_{k'}))^2 \leq (\#U_k)^2 + (\#U_{k'})^2$  then
      Set  $B_k = B_k \cup B_{k'}$ 
      Set  $B_{k'} = \{\}$ 
    end if
  end for
end for

```

The joining of two blocks is accepted if the square of the number of neighbors of the joined block is less than the sum of the squares of the number of neighbors in the two blocks. This ensures that the memory burden is made no worse when combining blocks, since the covariance matrices we need to form are governed by $(\#U_k)^2$. In practice, this rule also does not increase the computational burden since, when working with small covariance matrices, the computational demand is often dominated by filling in the entries of the covariance matrix rather than factoring the matrix. Only neighbors are considered as candidates for joining, ensuring that at most nm comparisons are made, and that we make comparisons between blocks that contain points that are near each other, which is advisable because distant points are unlikely to share many neighbors unless they appear early in the ordering.

4 Further computational considerations

This section describes computational considerations for the problems of finding orderings and nearest neighbors, approximately drawing from unconditional and conditional Gaussian process distributions, and profiling out linear mean parameters. Code for all methods is provided in online supplementary material.

4.1 Finding Orderings and Nearest Neighbors

Considering arbitrary permutations of the observations presents two computational issues that are not encountered in coordinate-based orderings. The first is finding the orderings. For example, finding the MMD ordering is $O(n^3)$ in computing time. The second problem is how to find the ordered nearest neighbors, that is, the set of m nearest neighbors to location $x_{(i)}$ among $x_{\tau(1)}, \dots, x_{\tau(i-1)}$. A naïve algorithm for finding neighbors would require $O(n \log n)$ operations to compute and sort distances to any point, and thus a total of $O(n^2 \log n)$ operations.

Searching for nearest neighbors in metric spaces is a well-studied problem. Vaidya (1989) describes an $O(n \log n)$ algorithm for finding the nearest neighbor to all n points. The FNN R package (Beygelzimer et al., 2013) has an implementation of a kd-tree algorithm. The algorithm uses a tree-based partitioning of the points to quickly narrow down the set of points that are candidates for being nearest each point. However, the FNN software cannot be used directly because the functions return nearest neighbors from the set of all points, rather than nearest previous neighbors, which is what we need for Vecchia's approximation. To bridge this gap, we first find the nearest $2m$ neighbor sets for each point. For the neighbor sets that contain at least m previous points, we have found the nearest m previous points. Those points are set aside, and for the remaining points, we find the nearest $4m$ points, again setting aside the points that have at least m previous neighbors. This is repeated until all points previous neighbors have been found.

It is possible to construct approximations to the MMD ordering that run in $O(n \log n)$ time. First divide the domain up into grid boxes, with the number of grid boxes proportional to the number of observations. Assigning each point to its grid box can be done in linear time with a simple rounding method. Using the center of each grid box as its location, then use MMD to order the grid boxes. Once the grid boxes are ordered, loop over the grid boxes according to the grid box ordering, each time picking the point from the current grid box with largest minimum distance to any selected points in the current or neighboring grid boxes, stopping when all points have been ordered. If the number of grid boxes g is large this algorithm can be used recursively to order the grid boxes, thus giving an $O(n \log n)$ algorithm.

4.2 Unconditional and Conditional Draws

Whereas Kriging interpolation requires only matrix multiplication and linear solves with the covariance matrix, which can usually be done in $O(n^2)$ operations with a good iterative solver, unconditional and conditional draws (simulations) from the Gaussian process model generally require us to factor the covariance matrix, which is $O(n^3)$ operations. One advantage that Vecchia's approximation has is that it defines an approximation to the inverse

Cholesky factor, which can be used to perform approximate draws from the Gaussian process model. This approximation can also be used to perform conditional draws of unobserved values given the observations, which, when done many times in an ensemble, is a useful way of quantifying the joint uncertainty associated with interpolated maps.

To simplify the following equations, write Γ as the inverse Cholesky factor of the $n \times n$ covariance matrix Σ , and $\tilde{\Gamma}$ as the approximation to Γ implied by Vecchia's likelihood. Let Z be an uncorrelated vector of standard normals of length n . Then $Y = \tilde{\Gamma}^{-1}Z$ is approximately $\mathcal{N}(0, \Sigma)$. We use Kullback-Leibler divergence in Section 5 to monitor the quality of the approximations. Since $\tilde{\Gamma}$ is sparse and triangular, solving for Y , and thus drawing approximately from $\mathcal{N}(0, \Sigma)$, can be done in $O(n)$ operations. Suppose Y is partitioned into two subvectors as $Y = (Y_1, Y_2)$ and write Σ as a 2×2 block matrix $\{\Sigma_{ij}\}_{i=1,2}$. To predict Y_2 from Y_1 , we compute $E(Y_2|Y_1) = \Sigma_{21}\Sigma_{11}^{-1}Y_1$. This can be rewritten in terms of the inverse Cholesky factor as $E(Y_2|Y_1) = -\Gamma_{22}^{-1}\Gamma_{21}Y_1$. Thus we can compute approximate conditional expectations as $-\tilde{\Gamma}_{22}^{-1}\tilde{\Gamma}_{21}Y_1$. Solving a linear system with $\tilde{\Gamma}_{22}$ can also be achieved in linear time because $\tilde{\Gamma}_{22}$ is sparse and lower triangular.

Conditional draws can be done at the cost of one unconditional draw and one conditional expectation. Suppose that the data are in the vector Y_1 , and that $Y^* = (Y_1^*, Y_2^*)$ is an unconditional draw from $\mathcal{N}(0, \Sigma)$. Then, conditionally on Y_1 , $-\tilde{\Gamma}_{22}^{-1}\tilde{\Gamma}_{21}(Y_1 - Y_1^*) + Y_2^*$ approximately has mean $\Sigma_{21}\Sigma_{11}^{-1}Y_1$ and covariance matrix $\Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}$ and thus is an approximate conditional draw of Y_2 given Y_1 . We use ensembles of conditional draws in Section 6 to quantify joint uncertainties in interpolations of satellite data.

4.3 Profile Likelihood with Linear Mean Parameters

The Gaussian process model often includes a mean function that is linear with respect to a set of spatial covariates. Stein et al. (2004) outlined methods for computing residual (also known as restricted) likelihoods to avoid having to numerically maximize the approximate likelihood over the mean parameters. If one prefers to obtain the maximum approximate likelihood estimates of all parameters, it is usually computationally advantageous to profile out the mean parameters. Writing $E(Y) = X\beta$, where X is the $n \times p$ design matrix, and β is the vector of linear mean parameters, Vecchia's approximate likelihood becomes

$$\log p_{\theta, \beta, \tau, J}(y) = -\frac{n}{2}\log(2\pi) + \log \det \tilde{\Gamma} - \frac{1}{2}\|\tilde{\Gamma}(y - X\beta)\|^2.$$

For fixed θ , the maximum approximate likelihood estimate for β is

$$\hat{\beta}(\theta) = (X^T \tilde{\Gamma}^{-1} X)^{-1} X^T \tilde{\Gamma}^{-1} y = [(\tilde{\Gamma} X)^T (\tilde{\Gamma} X)]^{-1} (\tilde{\Gamma} X)^T (\tilde{\Gamma} y).$$

Multiplying $\tilde{\Gamma}X$ can be done column-by-column, each of which takes the same computational effort as multiplying $\tilde{\Gamma}y$, which is the most computationally demanding task in the approximate likelihood evaluation. Thus profiling out the linear mean parameters can be done at the roughly the cost of p additional approximate likelihood evaluations per iteration when maximizing over θ .

5 Numerical and timing comparisons

This section contains numerical results studying differences in the relative quality Vecchia's approximation with respect to choices for the permutation and grouping, under different model settings and in two, three, and four dimensional domains. Vecchia's approximation is also compared to a simple block independent Gaussian process approximation, since this approximation is easily constructed and has been shown to be competitive with state-of-the-art approximations (Stein, 2014). Comparisons are also made with stochastic partial differential equation (SPDE) approximations (Lindgren et al., 2011) and tapered covariance approximations (Furrer et al., 2006; Kaufman et al., 2008). For all comparisons, the Matern covariance function is used,

$$M(r; \sigma^2, \alpha, \nu) = \frac{\sigma^2}{\Gamma(\nu)2^\nu - 1} \left(\frac{r}{\alpha}\right)^\nu \mathcal{K}_\nu\left(\frac{r}{\alpha}\right),$$

which has emerged as the model of choice for practitioners of spatial statistics (Guttorp and Gneiting, 2006). The covariance function has three positive parameters σ^2 , α , and ν , which are the variance, range, and smoothness parameters, respectively. The mean of the process is assumed to be known to be zero.

In the two-dimensional numerical studies, six different parameter choices are presented, with $\sigma^2 = 1$ in all six, while the range and smoothness cover all combinations of $\alpha \in \{0.1, 0.2\}$ and $\nu \in \{1/2, 1, 3/2\}$. The locations form an 80×80 regular grid on the unit square, giving 6400 locations. Example realizations from the six models are given in Figure 12 in Appendix A. Four orderings are considered: sorted coordinate, middle out, completely random, and AMMD. Ungrouped and grouped versions, and 30 and 60 nearest neighbors are considered. As mentioned in Section 3, calculations in the grouped versions effectively condition on more neighbors, so when discussing the results, "grouped version with 30 neighbors" refers to a grouped version of a likelihood approximation that initially conditioned on 30 neighbors. For each ordering and number of initial neighbors, Table 1 includes statistics on the number of blocks K , the sizes of the blocks ($\#U_k$), and the sizes of the conditioning sets ($\#J_k$).

The combinations above consist of (6 parameter settings) \times (4 orderings) \times (30 vs. 60 neighbors) \times (ungrouped vs. grouped) = 96 settings, not counting the block approximations or the SPDE approximations. Exploring such a large number of scenarios with simulations would not be feasible since many replicates would be required to control Monte Carlo error size. Instead, we use two deterministic criteria to evaluate the different approximations: (1) KL-divergence from the implied approximate model to the target model, (2) asymptotic

relative efficiency for estimating covariance parameters, computed using the the usual generalization of Fisher information for misspecified likelihoods (Heyde, 2008). If $\tilde{\ell}(\theta)$ is the approximate likelihood, the information criterion is

$$H(\theta) = E\left(\nabla^2 \tilde{\ell}(\theta)\right) E\left(\nabla \tilde{\ell}(\theta)\right) \left(\nabla \tilde{\ell}(\theta)\right)^T)^{-1} E\left(\nabla^2 \tilde{\ell}(\theta)\right),$$

where ∇ indicates gradient, and ∇^2 is the Hessian. The criterion $H(\theta)$ is of course equal to the Fisher information in the case that $\tilde{\ell}$ is equal to the true loglikelihood. The inverse of H gives the asymptotic covariance matrix for the three Vecchia likelihood parameter estimates, estimated simultaneously. Relative efficiencies are obtained by comparing the diagonal values of H^{-1} to the diagonal values of the inverse of the Fisher information.

In the three- and four-dimensional numerical studies, the grids on the unit square are of size 19^3 and 9^4 , giving totals of 6859 and 6561 locations. We use the exponential covariance function with range parameters 0.1, 0.2, and 0.4. Only the KL-divergence criterion and the ungrouped approximations are considered, but the same four orderings and up to 240 neighbors are used.

The section concludes with a timing study showing how the various computational tasks for Vecchia's approximation scale with the number of observations. All timing comparisons are done on a 2016 Macbook Pro with a 3.3GHz Intel Core i7 processor (two cores) with 16GB RAM. Code is mostly written in the R programming language (R Development Core Team, 2008), aside from the Vecchia likelihood functions, which have been implemented in C++. Computations are done while running R version 3.4.2, linked to Apple Accelerate multithreaded linear algebra libraries.

5.1 Kullback-Leibler Divergence Comparisons

For the $\nu = 1$ cases, an SPDE approximation (Lindgren et al., 2011) is also considered. The SPDE approximation provides computational savings by using a sparse inverse covariance matrix and is valid for integer values of the smoothness parameter in two spatial dimensions. To address edge effects in the SPDE approximation, we use the boundary extension described in Lindgren et al. (2011). Several choices of boundary parameters were considered, with only small differences among choices. For the $\nu = 1/2$ case, we also consider a tapered covariance approximation (Furrer et al., 2006; Kaufman et al., 2008) with three tapering ranges as a competitor. In contrast to Vecchia's approximation, the KL divergence from the SPDE or tapered approximations to the target model is not minimized at the target model's parameter values, so we perform an optimization to find the range and variance parameters to minimize the KL-divergence. SPDE computations were implemented using functions from the R-INLA package (Rue et al., 2009), tapering from the fields package (Nychka et al., 2016).

The KL divergence results are plotted in Figure 3, and the discussion of the results is organized according to the smoothness parameter setting. When $\nu = 1/2$, which corresponds to an exponential covariance function, AMMD ordering reduces KL divergence relative to sorted coordinate ordering by factors of 16 and 22 for the two range parameter settings when

the number of neighbors is 30. Both orderings run in roughly the same amount of time. When the observations are grouped, the approximations improve. When using the AMMD ordering, the improvement in KL divergence over the sorted coordinate ordering with no grouping increases to a factor of 64 when $\alpha = 0.1$ and a factor of 75 when $\alpha = 0.2$. For 60 neighbors the effect of ordering is more striking; AMMD ordering and grouping results reduces KL divergence by factors of 285 and 244. Covariance tapering is not competitive for any of the tapering ranges. The grouped approximation with AMMD ordering and 30 neighbors is more than 12,000 times more accurate than the tapered approximation that runs in roughly the same amount of time.

When $\nu=1$, the major story is that the grouped approximations with AMMD ordering run faster than the SPDE approximation and are far more accurate; for $\alpha = 0.1$, KL-divergence is reduced by a factor of 160, and for $\alpha = 0.2$ by a factor of 148. Computing times for the SPDE approximation include creating the mesh, forming the precision matrix (with boundary extension), and computing the sparse Cholesky factor. Reordering of the precision matrix is done by internal functions in the R Matrix package (Bates and Maechler, 2016). As before, the AMMD orderings are most accurate, especially with 60 neighbors.

When $\nu = 3/2$, which is the smoothest model considered, the AMMD ordering is again most accurate for 60 neighbors. For 30 ungrouped neighbors, the differences among orderings is less substantial. In all cases, grouped approximations are more accurate than their ungrouped counterparts. In all parameter settings, the block independent approximations are never competitive in accuracy. The 3×3 block approximation generally runs slower than the fastest grouped approximation and is orders of magnitude less accurate. Finer blocking in the block independent approximation will run faster but will decrease in accuracy.

KL-divergences for three and four dimensions are plotted in Figure 4, where we see that the middle out ordering is the best choice in all but two covariance settings. In three dimensions, the random and MMD orderings appear to have KL-divergences that decay faster as the number of neighbors increase, whereas the sorted coordinate orderings have a slower decay. However, the gains for random and MMD come only after a large number of neighbors, which are unlikely to be used in practice due to the cubic scaling in number of neighbors. In four dimensions, middle out performs best. KL divergences were also computed for covariance tapering with three different taper ranges. Covariance tapering was not competitive with Vecchia's approximation in any setting; KL divergences were larger than 10^2 in all settings, with some larger than 10^3 .

5.2 Relative Efficiency Comparisons

The results of the numerical study on relative efficiency are summarized in Figure 5. In every parameter setting, for all three parameters, and for both neighbor sizes, the AMMD ordering with grouping out-performs the default sorted coordinate ordering without grouping. In some cases, the difference is quite large. For example, when estimating the range parameter for the Matérn with smoothness 1 and range 0.2, using the default sorted coordinate without grouping has 93.2% relative efficiency, whereas AMMD with grouping has 99.7% relative efficiency. This gain in efficiency comes at no additional computational cost.

In every case, using 30 neighbors with the AMMD ordering and grouping is superior to using 60 neighbors with the default sorted coordinate ordering without grouping. Since the computational complexity scales with the cube of the number of neighbors, this means our proposed improvements to Vecchia's likelihood can achieve increases in relative efficiency while simultaneously reducing the computational cost by a factor of 8. Finally, perhaps the most surprising result of this numerical study is that the default orderings are almost always outperformed by a completely random ordering of the points. This is also true in the KL divergence study. In both cases, the AMMD ordering offers further improvement.

5.3 Timing

In the context of analyzing data, it is useful to note which calculations must be carried out just once versus multiple times, and which calculations can be parallelized, although no explicit parallelization has been used in these studies. The ordering, grouping, and nearest neighbor calculations are computed just once. The likelihood calculations generally need to be repeated many times in the process of either numerically maximizing it with respect to covariance parameters, or sampling from posterior distributions in MCMC algorithms, but likelihood evaluations are embarrassingly parallel due to the separability of Vecchia's specification. The nearest neighbor searches can be done in parallel.

Figure 6 presents the results of the timing study for an increasing number of observations with the exponential covariance function and AMMD ordering. The study is carried out for 30 and 60 neighbors and up to 10^5 observations. For 30 neighbors, the slowest operation is grouping, followed by ordering, then finding neighbors, then the likelihood evaluations. The grouped likelihood for 10^5 observations was computed in 1.5 seconds. For 60 neighbors, the grouping algorithm was slowest, followed by ordering. The grouped likelihood with 60 neighbors required 7.4 seconds for 10^5 observations.

6 Jason-3 Satellite Wind Speed Observations

Launched in January 2016, the Jason-3 satellite is the latest in a series of satellites equipped with radar altimeters for measuring ocean surface height, wave height, and ocean wind speed (NASA-JPL, 2016). Jason-3 orbits the earth every 112 minutes along a path that repeats every 9.9 days. The data we consider are ocean surface wind speed values reported roughly once per second between August 4 and August 9, 2016 and are available at <http://www.nodc.noaa.gov/sog/jason/>. The goal of this analysis is to create interpolated maps of wind speed and quantify the joint uncertainty in the interpolations.

Data Preprocessing:

The Jason-3 data files include rain and ice flags to signify the presence of liquid water and ice along the radar signal's path. These disruptions degrade the quality of the signal, and the Jason-3 products handbook (CNES et al., 2016) states that flagged measurements should be ignored. As a conservative measure, we discard any measurement taken within 30 seconds of passing over rain or ice. In order to ensure that we can perform extensive comparisons among statistical analyses for data spanning a reasonably long time period, we average the one second measurements within 10 second intervals, discarding any intervals that have

missing values. The resulting data vector has 18,973 values. A map of the data is plotted in Figure 7.

Since the satellite can measure wind speeds at only a single location at any given time, it is expected that uncertainties in interpolated maps at specific times and locations will depend on whether there are nearby observations in space and time. This feature can be captured by modeling the data with a space-time Gaussian process. Specifically, for wind speeds at location $x \in \mathbb{S}^2$ and $t \in \mathbb{R}$, we consider ageostatistical model of the form

$$Y(x, t) = \mu + Z(x, t) + \varepsilon(x, t),$$

where μ is considered nonrandom, $\varepsilon(x, t)$ is uncorrelated $N(0, \tau^2)$ (nugget term), and $Z(x, t)$ is a spacetime Gaussian process with a Matern covariance function

$$\text{Cov}(Z(x_1, t_1), Z(x_2, t_2)) = \frac{\sigma^2}{2^\nu - 1 \Gamma(\nu)} (d_{12})^\nu \mathcal{K}_\nu(d_{12}),$$

Where \mathcal{K}_ν is a modified Bessel function of the second kind, and

$$d_{12} = \sqrt{\frac{\|x_1 - x_2\|^2}{\alpha_1^2} + \frac{|t_1 - t_2|^2}{\alpha_2^2}},$$

so that the covariance function is isotropic in space, stationary in time, but with different range parameters for the space and time dimensions. This method for constructing covariance functions on the sphere-time domain was originally used in Jun and Stein (2007). The Matérn covariance is not generally valid with great circle distance metric on the sphere (Gneiting, 2013), so Euclidean distance is used. Porcu et al. (2016) constructed some alternative space-time covariance functions, but Guinness and Fuentes (2016) argued that there is no reason to expect the use of covariance functions constructed on Euclidean spaces and restricted to the sphere to cause any distortions. The numerical results in Porcu et al. (2016) support this idea. The model we consider has five unknown covariance parameters, $(\sigma^2, \alpha_1, \alpha_2, \nu, \tau^2)$.

We consider four orderings for the observations in Vecchia's likelihood. The first is ordered in time, which I consider to be a default choice for this application. The second is completely random, the third is MMD in time, and the fourth is MMD in space. Since the orbital pattern of the satellite follows a regular path over time, the MMD in time ordering provides good spatial coverage early in the ordering, as can be seen in Figure 8, which shows the locations of the first 1000 observations from the MMD in time ordering.

For neighbor selection, Datta et al. (2016c) defined "nearest" neighbors in space and time based on the correlation function. Here, we define distance based on the spatial dimension only. The reason for this choice is to encourage the neighbor sets to contain observations

from different time points so that the conditional likelihoods contain information about the temporal range parameter.

Since the quality of Vecchia's approximation increases as neighbors are added, and maximum likelihood estimates are obtained in an iterative numerical search procedure, it is natural to consider a sequence of maximum approximate likelihood estimates, where each step in the sequence uses a larger number of neighbors than in the previous step. The estimates from an iterative procedure such as this can be monitored for convergence and stopped when successive parameter estimates do not change beyond some tolerance level. For each ordering, we start with 10 neighbors, find the maximum approximate likelihood estimate, and then we consider 20 neighbors, starting the optimization at the 10-neighbor parameter estimate, and so on up to 100 neighbors. As long as the optimization procedures are finding the legitimate maximum of the approximate likelihoods, the sequence of estimates from all orderings will all converge to the same maximum likelihood estimate, since all of the likelihood approximations converge to the exact likelihood when we condition on all possible past observations. Orderings can be compared to each other based on how many neighbors are required for convergence.

The results for the four orderings are plotted in Figure 9. Only the variance, spatial range, temporal range, and smoothness parameters are plotted because the nugget variance is essentially zero for every ordering and number of neighbors. It is clear that the parameter estimates from the random ordering and the two MMD orderings are converging more rapidly than are the estimates from the time ordering.

The MMD orderings appear to be converging slightly more quickly than the random ordering. For the random and MMD orderings, the estimates for all parameters are within 2% of the 100 neighbor estimates with just 50 neighbors, whereas the time ordered estimates do not appear to have settled down, even with 100 neighbors. The exception is the smoothness parameter, whose estimates for 10 neighbors are within 2% of the 100 neighbor estimates for all orderings.

Spatial interpolations are computed on an evenly spaced grid of size 120×240 in latitude and longitude at two time periods. The first is at the average of all the observation times, to represent a hindcast—a prediction of past values from past data. The second is at the time of the last observation, to represent real-time interpolation of the data. Using ordering 3 for the observations, a random ordering for the prediction points, and 60 nearest neighbors with grouping, we compute 4000 conditional draws and use the sample variances of the conditional draws to estimate the conditional variances at each of the prediction points. The predictions and their standard deviations are plotted in Figure 10. As expected, the standard deviations are generally smaller for the hindcast predictions, and one can clearly see the effect of the orbital path of the satellite; locations along recently visited paths have smaller standard deviations. Further, the standard deviations are small at high and low latitudes, reflecting the fact that the satellite passes near the poles on every orbit. Finally, Figure 11 contains two individual hindcast conditional draws and maps of conditional correlations with two points in the southern Atlantic Ocean. One can see that the conditional correlation structure is quite complex and inhomogeneous among the two points. A possible explanation

is that this region has nearby observations in space and time, as can be seen from bottom left panel of Figure 10.

7 Discussion

This paper demonstrates that reordering and grouping operations can lead to substantial improvements in the quality of Vecchia's approximation, sometimes by more than two orders of magnitude compared to default ungrouped approximations. Grouping also reduces the computational effort over ungrouped versions, and the grouping algorithm introduced here is general in that it can be applied to any choice of ordering and neighbor sets. I have also provided the R code written for reordering, finding nearest neighbors, automatic grouping, and likelihood evaluations.

One perhaps surprising result of this study is that Vecchia's approximation with MMD ordering and grouping with 30 neighbors runs faster than an SPDE approximation, while being two orders of magnitude more accurate in terms of KL divergence. Certainly more numerical results are needed in more cases, but this result coupled with the fact that Vecchia's approximation is valid for general covariance functions rather than being restricted to Matern with integer smoothness, suggests that Vecchia's approximation should be considered as a candidate for approximating Matern models whenever the SPDE approximation is considered.

The paper introduces a grouped version of Vecchia's approximation, based on a partitioning of the observations into blocks. Each block's contribution to the likelihood can be computed simultaneously.

A theorem is provided to explain why grouping improves the model approximation, and an algorithm is described for finding a partition that is guaranteed control the memory burden. Exploring new partitioning algorithms for Vecchia's approximation could be a fruitful avenue for future work.

While developing some theory for grouping has been successful, a general theory for reordering remains elusive. One reason for this is the diversity of covariance functions and observation settings to be considered. Results on the screening effect (Stein, 2002) may be helpful for pushing the theory forward in some special cases, and perhaps theoretical results on estimates from subsamples could be useful as well (Hung and Zhao, 2016). However, I caution against drawing strong conclusions based on one-dimensional results. For example, for the exponential covariance in one-dimension, Vecchia's approximation is exact when observations are sorted along the coordinate and just a single nearest neighbor is used. Unsorted orderings require two neighbors for at least one of the points. I suspect that intuition gleaned from these one-dimensional examples has been incorrectly applied to the two-dimensional case, leading to the preference for sorted coordinate orderings in two dimensions. The numerical results in three and four dimensions are an attempt to study how ordering influences Vecchia's approximation in higher dimensions, which is relevant for computer experiment applications (Sacks et al., 1989; Santner et al., 2013). There is some evidence that KL divergence decays faster with the number of neighbors under random or

MMD ordering in three dimensions, but middle out ordering appears to be best in four dimensions, at least in the examples considered here. This is further evidence to be wary of applying intuition from lower dimensions to higher dimensions.

This work considers the effect of ordering with the rule for choosing neighbor sets held constant.

Stein et al. (2004) presented examples where including some distant neighbors can help in estimating parameters that control the behavior of the covariance function away from the origin. I have tried this approach but was not able to improve on nearest neighbors. A possible explanation is that I consider the Matern model with a different parameterization than in Stein et al. (2004). I note here that the likelihood with MMD ordering automatically includes information about distant relationships since observations early in the ordering necessarily condition on distant observations since early observations are necessarily distant from one another.

The work on the local approximate Gaussian process (laGP) (Gramacy and Apley, 2015) explores fast automatic neighbor selection in a different context and could be relevant to this question. While neighbors figure prominently in both the present paper and in laGP, neighbors in laGP can be selected from anywhere in the ordering, whereas Vecchia's approximation requires neighbors to come from previous in the ordering. Ordered neighbors ensure that the approximation corresponds to a valid joint density, which is intended to approximate a specified joint density, whereas laGP constructs a framework for interpolation and does not target any particular global model.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Number 1613219.

Appendix

A: Realizations from Models Studied

Figure 12 contains example realizations from the models used in the KL-divergence and relative efficiency studies.

References

- Amestoy PR, Davis TA, and Duff IS (1996). An approximate minimum degree ordering algorithm. *SIAM Journal on Matrix Analysis and Applications*, 17(4):886–905.
- Banerjee S, Carlin BP, and Gelfand AE (2014). *Hierarchical modeling and analysis for spatial data*. CRC Press.
- Bates D and Maechler M (2016). *Matrix: Sparse and Dense Matrix Classes and Methods*. R package version 12–4.

- Beygelzimer A, Kakadet S, Langford J, Arya S, Mount D, and Li S (2013). FNN: Fast Nearest Neighbor Search Algorithms and Applications. R package version 1.1.
- CNES EUMETSAT, JPL, and NOAA/NESDIS (2016). Jason-3 products handbook.
- Datta A, Banerjee S, Finley AO, and Gelfand AE (2016a). Hierarchical nearest-neighbor gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, 111(514):800–812. [PubMed: 29720777]
- Datta A, Banerjee S, Finley AO, and Gelfand AE (2016b). On nearest-neighbor Gaussian process models for massive spatial data. *Wiley Interdisciplinary Reviews: Computational Statistics*, 8(5): 162–171. [PubMed: 29657666]
- Datta A, Banerjee S, Finley AO, Hamm NA, Schaap M, et al. (2016c). Nonseparable dynamic nearest neighbor gaussian process models for large spatio-temporal data with an application to particulate matter analysis. *The Annals of Applied Statistics*, 10(3): 1286–1316. [PubMed: 29657659]
- Eide AL, Omre H, and Ursin B (2002). Prediction of reservoir variables based on seismic data and well observations. *Journal of the American Statistical Association*, 97(457):18–28.
- Furrer R, Genton MG, and Nychka D (2006). Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3):502–523.
- Gneiting T (2013). Strictly and non-strictly positive definite functions on spheres. *Bernoulli*, 19(4): 1327–1349.
- Gramacy RB and Apley DW (2015). Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24(2):561–578.
- Guinness J and Fuentes M (2016). Isotropic covariance functions on spheres: Some properties and modeling considerations. *Journal of Multivariate Analysis*, 143:143–152.
- Guttorp P and Gneiting T (2006). Studies in the history of probability and statistics XLIX on the Matern correlation family. *Biometrika*, 93(4):989–995.
- Heyde CC (2008). Quasi-likelihood and its application: a general approach to optimal parameter estimation. Springer Science & Business Media.
- Hung Y and Zhao Y (2016). Variable selection for gaussian process models using experimental design-based subagging. *Statistica Sinica*, pages 1–10.
- Jones RH and Zhang Y (1997). Models for continuous stationary space-time processes In *Modelling longitudinal and spatially correlated data*, pages 289–298. Springer.
- Jun M and Stein ML (2007). An approach to producing space-time covariance functions on spheres. *Technometrics*, 49(4):468–479.
- Kaufman CG, Schervish MJ, and Nychka DW (2008). Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association*, 103(484): 1545–1555.
- Lindgren F, Rue H, and Lindstrom J (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(4):423–498.
- Mardia KV and Marshall R (1984). Maximum likelihood estimation of models for residual covariance in spatial regression. *Biometrika*, 71(1):135–146.
- NASA JPL (2016). Jason-3 mission website. <http://sealevel.jpl.nasa.gov/missions/jason3/>. Accessed: 2016–09–03.
- Nychka D, Furrer R, Paige J, and Sain S (2016). fields: Tools for Spatial Data. R package version 8.3–6.
- Pardo-Igúzquiza E and Dowd PA (1997). AMLE3D: A computer program for the inference of spatial covariance parameters by approximate maximum likelihood estimation. *Computers & Geosciences*, 23(7):793–805.
- Porcu E, Bevilacqua M, and Genton MG (2016). Spatio-temporal covariance and cross-covariance functions of the great circle distance on a sphere. *Journal of the American Statistical Association*, 111(514):888–898.
- Pourahmadi M (1999). Joint mean-covariance models with applications to longitudinal data: unconstrained parameterisation. *Biometrika*, 86(3):677–690.

- R Development Core Team (2008). R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
- Rue H, Martino S, and Chopin N (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.
- Saad Y (2003). Iterative methods for sparse linear systems. Society for Industrial and Applied Mathematics.
- Sacks J, Welch WJ, Mitchell TJ, and Wynn HP (1989). Design and analysis of computer experiments. *Statistical Science*, 4(4):409–423.
- Santner TJ, Williams BJ, and Notz WI (2013). The design and analysis of computer experiments. Springer Science & Business Media.
- Stein ML (2002). The screening effect in Kriging. *Annals of Statistics*, 30(1):298–323.
- Stein ML (2014). Limitations on low rank approximations for covariance matrices of spatial data. *Spatial Statistics*, 8:1–19.
- Stein ML, Chi Z, and Welty LJ (2004). Approximating likelihoods for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(2):275–296.
- Sun Y and Stein ML (2016). Statistically and computationally efficient estimating equations for large spatial datasets. *Journal of Computational and Graphical Statistics*, 25(1):187–208.
- Vaidya PM (1989). An $O(n \log n)$ algorithm for the all-nearest-neighbors problem. *Discrete & Computational Geometry*, 4(2):101–115.
- Vecchia AV (1988). Estimation and model identification for continuous spatial processes. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 297–312.
- Yannakakis M (1981). Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 2(1):77–79.

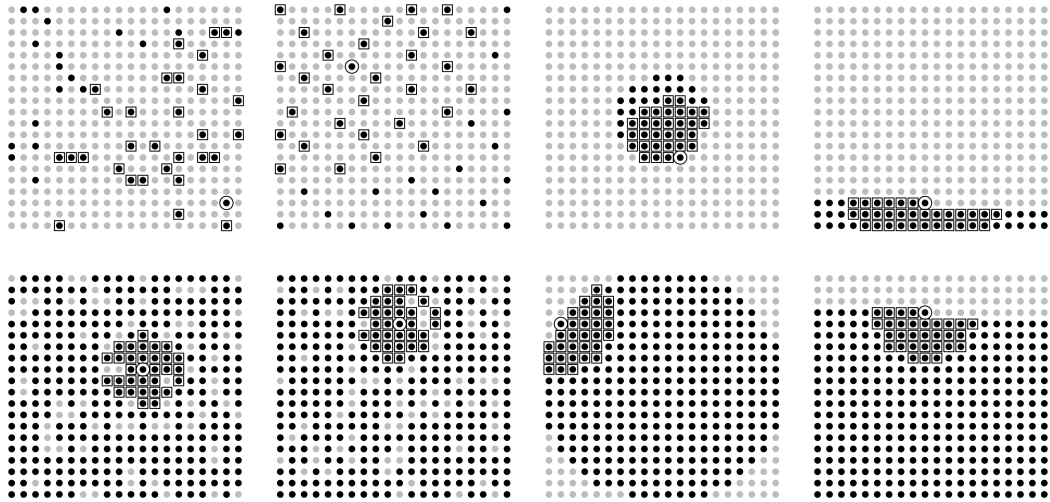


Figure 1:
 Examples of four orderings of 400 locations on a 20×20 grid. From left to right are completely random, maximum minimum distance, middle out, and sorted on vertical coordinate. On the top row, the black circles are ordered 1–50, circled point is number 50, and squared points are 30 nearest neighbors to point 50 among previous points. Bottom row is the same for point ordered number 330.

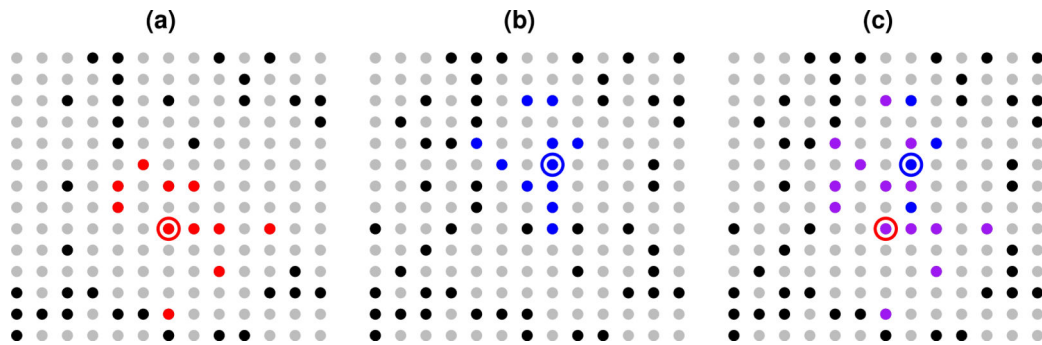


Figure 2:

Example with $B_i = \{45, 60\}$. (a) Point 45 (circled), points 1 through 45 (black or red), J_{45} (red). (b) Point 60 (circled), points 1 through 60 (black or blue), J_{60} (blue). (c) Points 1 through 60 (not gray), \bar{J}_{45} (purple), \bar{J}_{60} (purple + blue). \bar{J}_{45} has three more points than J_{45} , and \bar{J}_{60} has 7 more points than J_{60} .

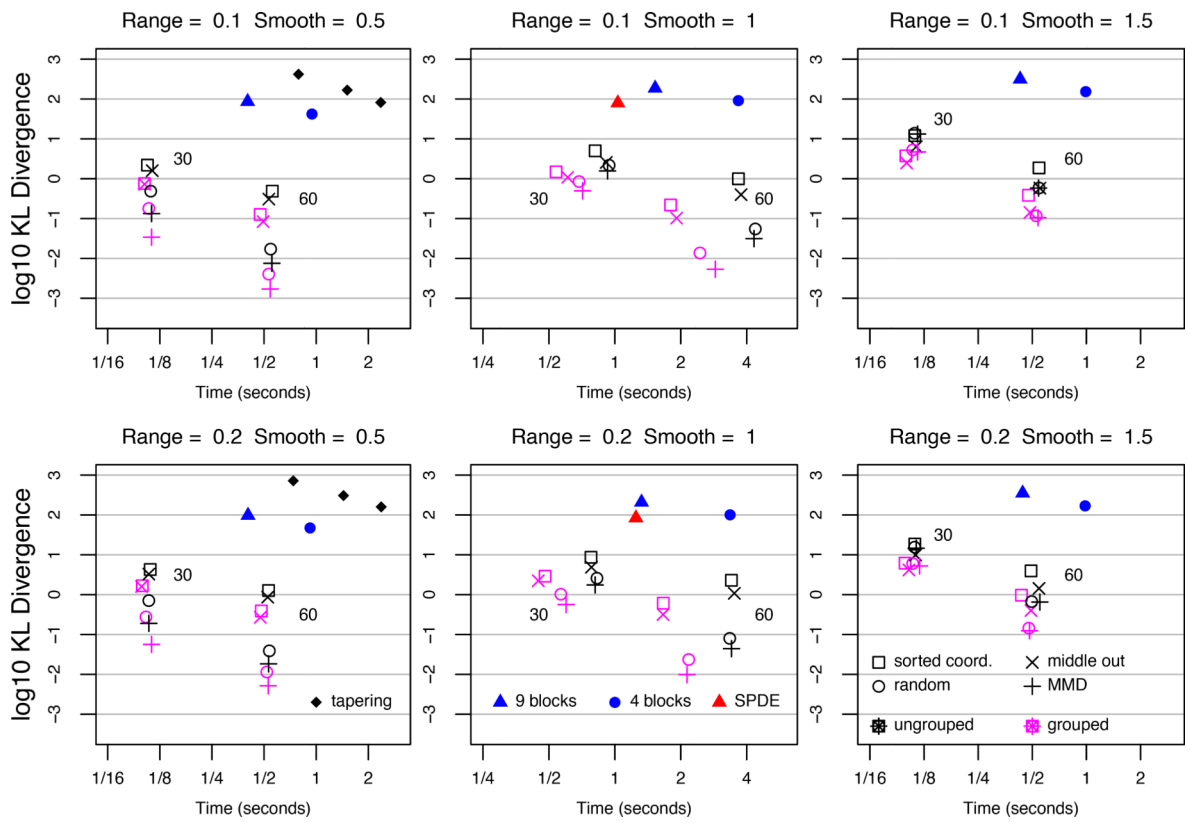


Figure 3: KL-divergences for Matérn covariances with range $\in \{0.1, 0.2\}$, and smoothness $\in \{1/2, 1, 3/2\}$. Locations form 80×80 grid on the unit square. Number of neighbors (30 and 60) indicated next to ungrouped symbols.

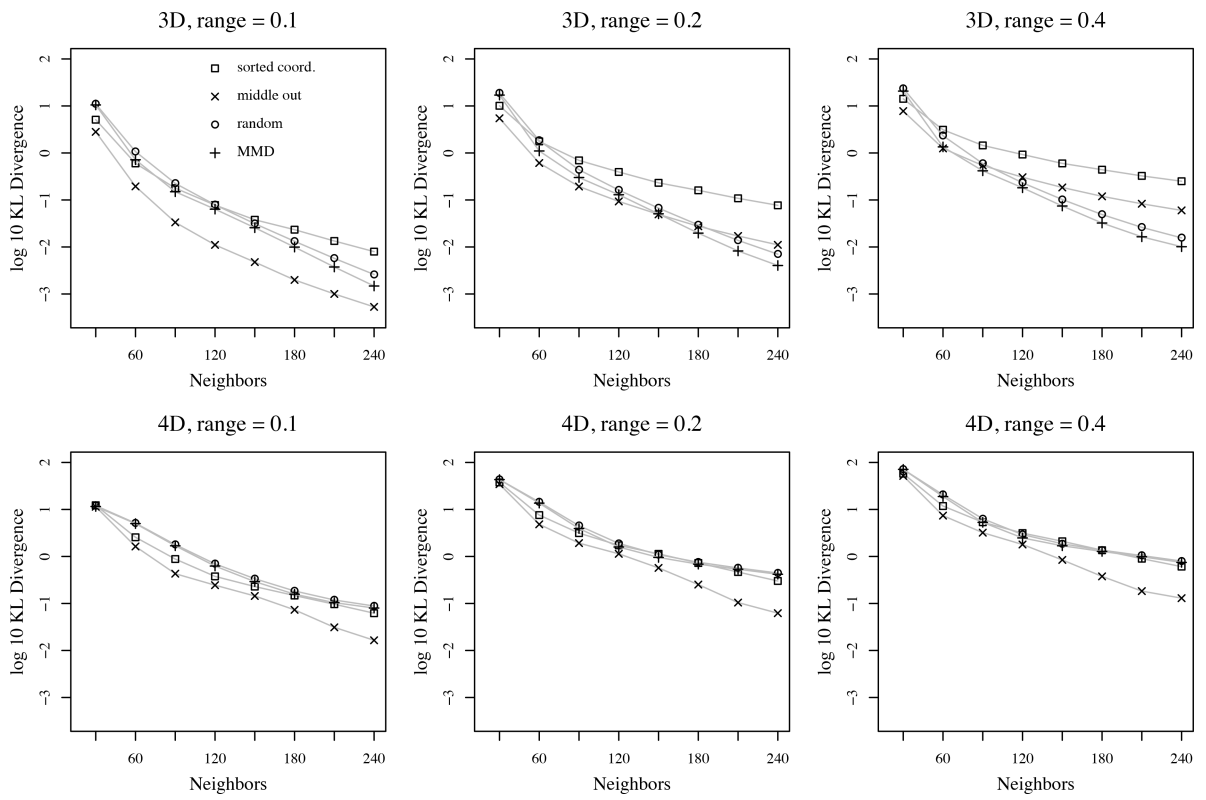


Figure 4: KL-divergences as a function of ordering and number of nearest neighbors for ungrouped version of Vecchia’s approximation with exponential covariance. Locations form a regular grid in 3 and 4 dimensions, totaling 6859 and 6561 locations.

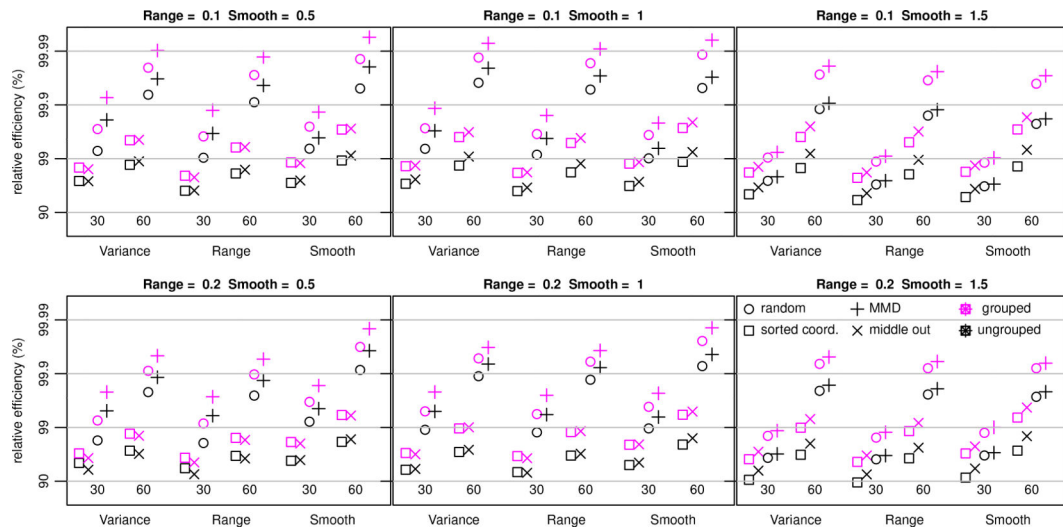


Figure 5: Relative efficiencies for estimating three Matérn covariance parameters, variance = 2, range $\in \{0.1, 0.2\}$, and smoothness $\in \{1/2, 1, 3/2\}$. Locations form 80×80 grid on the unit square, and the four orderings considered are sorted coordinate, middle out, completely random, and AMMD. Results for ungrouped (black) and not grouped (magenta) are presented.

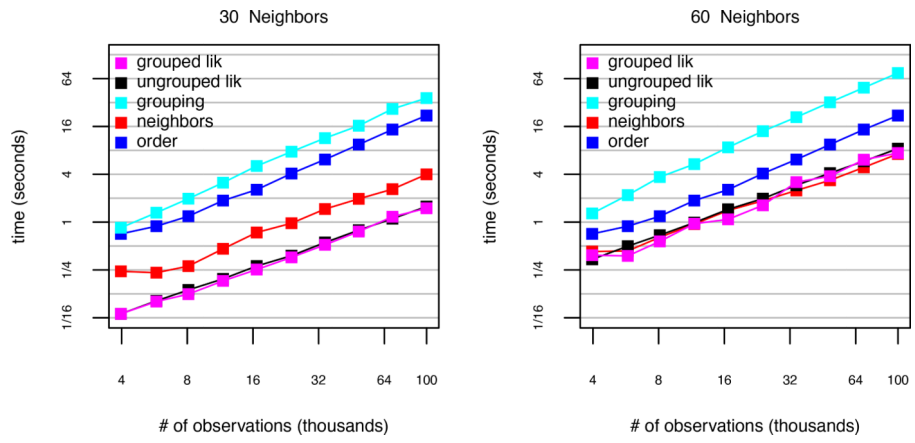


Figure 6: Timing Results for increasing number of observations regularly spaced on a square grid.

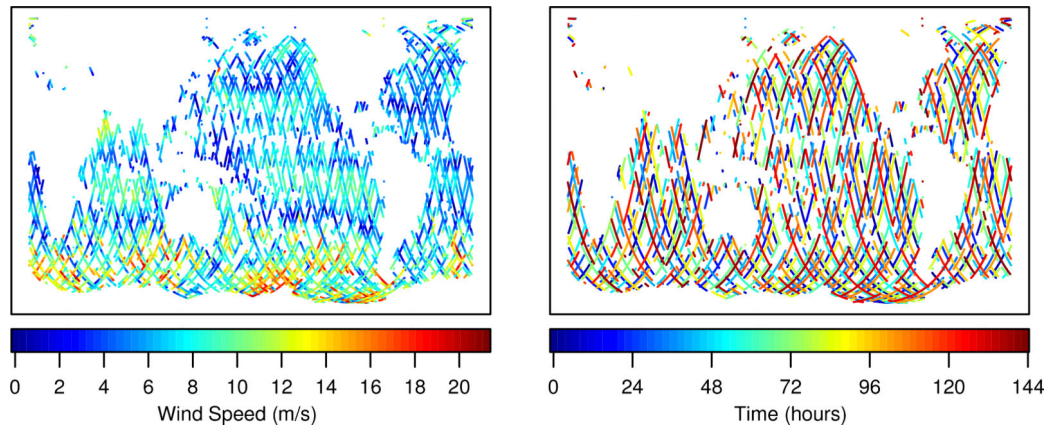


Figure 7:
Jason-3 wind speed values and observation times.



Figure 8:
For MMD time ordering, first 1000 locations (black) and all locations (gray).

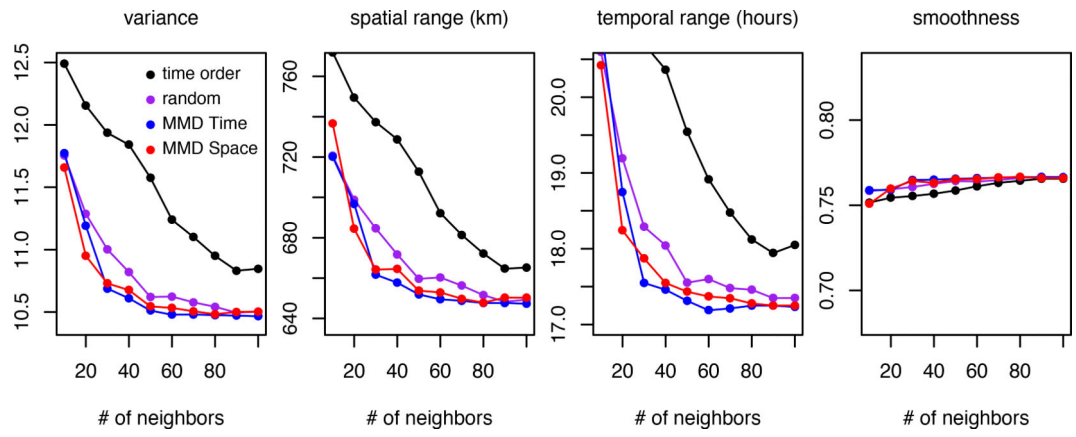


Figure 9: Estimated parameters for increasing number of neighbors. Vertical axis heights are 20% of the value of the last estimate from ordering 3.

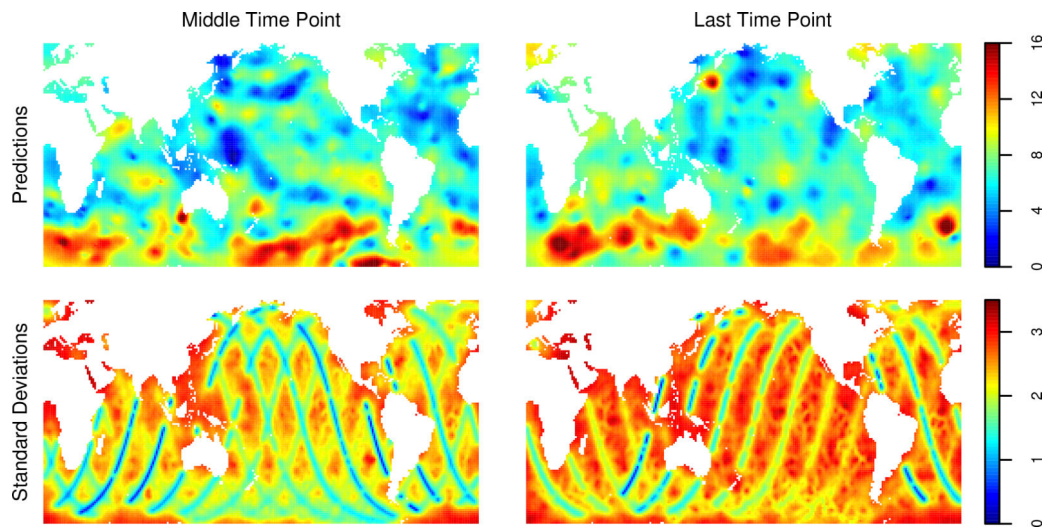


Figure 10: Predictions (top row) and simulated prediction standard deviations (bottom row) of windspeed at the mean observation time (left) and the last observation time (right).

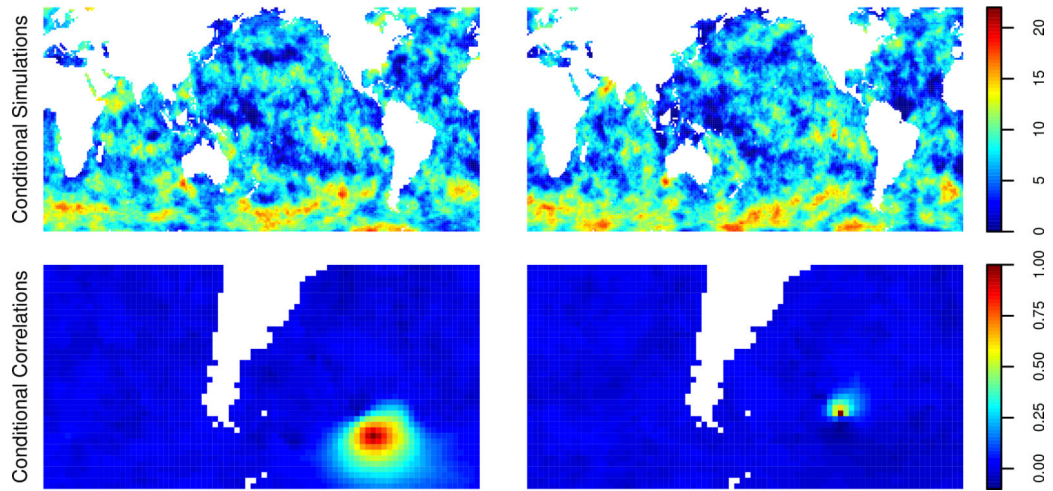


Figure 11: For hindcast, two individual conditional draws (top row) and empirical conditional correlations at two points in the southern Atlantic.

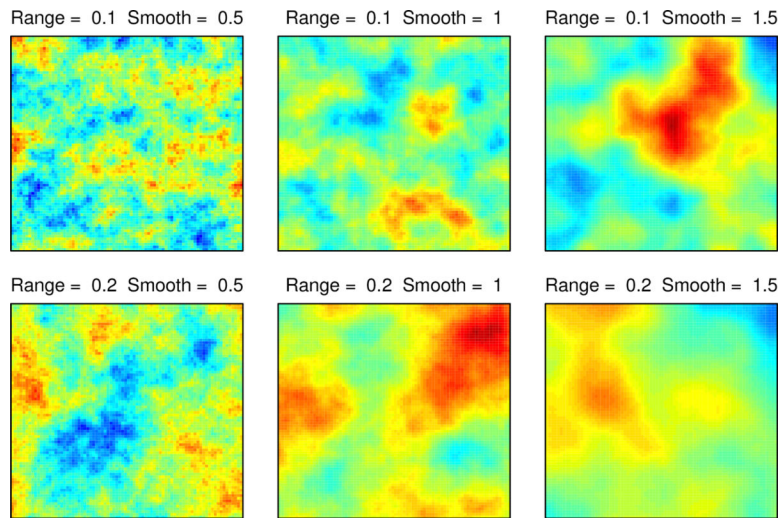


Figure 12: Gaussian process realizations at the six Matérn parameter settings.

Table 1:

For each ordering scheme and for both initial number of neighbor choices m , statistics on the number of blocks returned by the grouping algorithm K , the size of the neighbor sets in each group ($\#U_k$), and the actual number of neighbors ($\# \bar{J}_i$). Here $k = 1, \dots, K$, and $i = 1, \dots, n$.

Ordering	m	K	mean ($\#U_k$)	max ($\#U_k$)	mean ($\# \bar{J}_i$)	max ($\# \bar{J}_i$)
MMD	30	1406	51.35	124	55.83	123
random	30	1361	50.82	114	53.01	113
coordinate	30	788	58.95	171	66.02	170
middle out	30	773	60.33	154	60.35	153
MMD	60	754	116.46	291	133.36	290
random	60	714	118.46	308	129.74	307
coordinate	60	358	136.83	450	166.49	449
middle out	60	320	145.27	418	155.87	417