Spatial Statistics 43 (2021) 100517



Contents lists available at ScienceDirect

Spatial Statistics

journal homepage: www.elsevier.com/locate/spasta

Efficiency assessment of approximated spatial predictions for large datasets



Yiping Hong^a, Sameh Abdulah^b, Marc G. Genton^a, Ying Sun^{a,*}

^a Statistics Program, King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia

^b Extreme Computing Research Center (ECRC), King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia

ARTICLE INFO

Article history: Received 7 September 2020 Received in revised form 2 May 2021 Accepted 4 May 2021 Available online 14 May 2021

Keywords: Covariance approximation Gaussian predictive process Loss of efficiency Maximum likelihood estimation Spatial prediction Tile Low-Rank approximation

ABSTRACT

Due to the well-known computational showstopper of the exact Maximum Likelihood Estimation (MLE) for large geospatial observations, a variety of approximation methods have been proposed in the literature, which usually require tuning certain inputs. For example, the recently developed Tile Low-Rank approximation (TLR) method involves many tuning parameters, including numerical accuracy. To properly choose the tuning parameters, it is crucial to adopt a meaningful criterion for the assessment of the prediction efficiency with different inputs. Unfortunately, the most commonly-used Mean Square Prediction Error (MSPE) criterion cannot directly assess the loss of efficiency when the spatial covariance model is approximated. Though the Kullback-Leibler Divergence criterion can provide the information loss of the approximated model, it cannot give more detailed information that one may be interested in, e.g., the accuracy of the computed MSE. In this paper, we present three other criteria, the Mean Loss of Efficiency (MLOE), Mean Misspecification of the Mean Square Error (MMOM), and Root mean square MOM (RMOM), and show numerically that, in comparison with the common MSPE criterion and the Kullback-Leibler Divergence criterion, our criteria are more informative, and thus more adequate to assess the loss of the prediction efficiency by using the approximated or misspecified covariance models. Hence, our suggested criteria are more useful for the

* Corresponding author.

E-mail addresses: yiping.hong@kaust.edu.sa (Y. Hong), sameh.abdulah@kaust.edu.sa (S. Abdulah), marc.genton@kaust.edu.sa (M.G. Genton), ying.sun@kaust.edu.sa (Y. Sun).

https://doi.org/10.1016/j.spasta.2021.100517

2211-6753/© 2021 Elsevier B.V. All rights reserved.

determination of tuning parameters for sophisticated approximation methods of spatial model fitting. To illustrate this, we investigate the trade-off between the execution time, estimation accuracy, and prediction efficiency for the TLR method with extensive simulation studies and suggest proper settings of the TLR tuning parameters. We then apply the TLR method to a large spatial dataset of soil moisture in the area of the Mississippi River basin, and compare the TLR with the Gaussian predictive process and the composite likelihood method, showing that our suggested criteria can successfully be used to choose the tuning parameters that can keep the estimation or the prediction accuracy in applications.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Geostatistical applications include modeling the spatial distribution of a set of observations (e.g., temperature, humidity, soil moisture, wind speed) taken at *n* locations regularly or irregularly spaced over a given geographical area. In geostatistics, the spatial datasets are often considered as a realization of a Gaussian process, defined by a mean function and a spatial covariance model. More specifically, we suppose that the data are observed from a stationary, isotropic Gaussian random field $\{Z(\mathbf{s}) : \mathbf{s} \in D \subset \mathbb{R}^d\}$, with mean zero and covariance function $C(h; \theta) := \operatorname{Cov}_{\theta}\{Z(\mathbf{s}_1), Z(\mathbf{s}_2)\}$ for any $\mathbf{s}_1, \mathbf{s}_2 \in D$ and $\|\mathbf{s}_1 - \mathbf{s}_2\| = h$, where θ is the unknown parameter vector. In recent years, the Matérn family has been a popular choice for the covariance function, since it represents a general form of many possible covariance models in the literature, due to its flexibility. The Matérn covariance function is defined as

$$C(h;\theta) = \frac{\sigma^2}{\Gamma(\nu)2^{\nu-1}} \left(\frac{h}{\alpha}\right)^{\nu} \mathcal{K}_{\nu}\left(\frac{h}{\alpha}\right),\tag{1}$$

where $\theta = (\sigma^2, \alpha, \nu)^{\top}$, $\sigma^2 > 0$, $\alpha > 0$, and $\nu > 0$ are the variance, range parameter, and smoothness parameter, respectively, and \mathcal{K}_{ν} is the modified Bessel function of the second kind of order ν .

The Maximum Likelihood Estimation (MLE) method has been widely used for estimating the parameter vector $\boldsymbol{\theta}$ of the spatial model. Denoting the spatial dataset by $\boldsymbol{Z} = \{Z(\boldsymbol{s}_1), \ldots, Z(\boldsymbol{s}_n)\}^\top$, where $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n$ are the observation locations, the MLE of the unknown parameter $\boldsymbol{\theta}$ can then be obtained by maximizing the following log-likelihood function:

$$l(\boldsymbol{\theta}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det\{\boldsymbol{\Sigma}(\boldsymbol{\theta})\} - \frac{1}{2}\boldsymbol{Z}^{\top}\boldsymbol{\Sigma}(\boldsymbol{\theta})^{-1}\boldsymbol{Z},$$
(2)

where $\Sigma(\theta)$ is the covariance matrix, with entries $[\Sigma(\theta)]_{i,j} = C(||\mathbf{s}_i - \mathbf{s}_j||; \theta)$ for i, j = 1, ..., n. Finding the exact MLE requires $O(n^3)$ computations and $O(n^2)$ memory, since evaluating the loglikelihood function involves the inverse and the determinant of the covariance matrix. Thus, the exact MLE is not feasible for large spatial datasets in applications, e.g. meteorological data, where n is often of an order of 10^5 or 10^6 .

To overcome this computational problem, finding approximation methods to compute the MLE has drawn considerable attention. The approximation can be applied to the spatial model, log-likelihood function, and covariance matrix. First, the spatial model can be approximated by a low-rank model, which is easier to compute. For instance, Cressie and Johannesson (2006, 2008) proposed the fixed rank kriging (FRK) method, which approximates the spatial dependence model by a linear combination of proper basis functions. Banerjee et al. (2008) introduced the Gaussian predictive process (GPP), where the spatial model is approximated by the kriging prediction using the observations on some pre-determined knots plus a nugget effect. Finley et al. (2009) modified this method by introducing the fine-scale process and fixed the problem that the marginal

variance is underestimated. Second, for the approximation of log-likelihood function, Vecchia (1988) and Curriero and Lele (1999) introduced the composite likelihood approach by ignoring the correlation of the observations at distant locations in the function. Stein et al. (2004) showed that this approximation could also be adapted to the restricted likelihood.

Third, the covariance matrix can be approximated by a sparse matrix. In the covariance tapering method (Furrer et al., 2006; Kaufman et al., 2008; Du et al., 2009), the covariance matrix is multiplied element-wise by a sparse covariance matrix, so the dependency between distant locations is neglected. Stein (2014) showed that one could approximate the covariance matrix by dividing the covariance matrix by several tiles and replacing the off-diagonal tiles by zero matrices. This approximation can provide a more accurate prediction compared with the lowrank model-based method. Naturally, one can introduce a more delicate sparse structure for the covariance matrix approximation. The H-matrix (Hackbusch, 1999) defines a hierarchical block structure for the matrix, which allows a coarse approximation for the block distant from the diagonal and a delicate approximation for the block near the diagonal. There are different kinds of H-matrix approximation, such as the HODLR (Aminfar et al., 2016), HSS (Ghysels et al., 2016), H2matrices (Borm and Christophersen, 2016; Sushnikova and Oseledets, 2016), and BLR/TLR (Pichon et al., 2017; Akbudak et al., 2017; Abdulah et al., 2018). The recently proposed Tile Low-Rank (TLR) approximation method (Akbudak et al., 2017; Abdulah et al., 2018) divides the covariance matrix into several tiles and performs low-rank approximations on the off-diagonal tiles. Abdulah et al. (2018) showed that it could improve the computation of the likelihood function on parallel architectures such as shared-memory, GPUs, and distributed-memory systems. Abdulah et al. (2019) also considered using different precisions for the diagonal and off-diagonal tiles in the Cholesky decomposition of covariance matrices, which can also improve the computational performance. One can also approximate the inverse of the covariance matrix, or the precision matrix, instead (Lindgren et al., 2011; Nychka et al., 2015). Sun and Stein (2016) introduced a sparse inverse Cholesky decomposition in the score equation and obtained the score equation approximation method. Besides the categories stated above, the MLE can also be approximated by algorithmic approaches, such as the metakriging (Minsker et al., 2014), the gapfill method (Gerber et al., 2018), and the local approximate Gaussian process (Gramacy and Apley, 2015). For a detailed review of the MLE approximation approaches in the literature, refer to Sun et al. (2012) and Heaton et al. (2019).

All the above approximation methods require certain types of tuning, to some extent. We can call them 'tuning parameters' to distinguish them from model parameters that need to be estimated from the data. For instance, for the covariance tapering method (Kaufman et al., 2008), the taper range is a tuning parameter. The composite likelihood method (Vecchia, 1988) approximates the conditional density $p(\mathbf{s}_i | \mathbf{s}_1, \dots, \mathbf{s}_{i-1})$ conditioning on a subset of $\mathbf{s}_1, \dots, \mathbf{s}_{i-1}$, such as *m* nearest neighbors of s_i , for which m is a tuning parameter. In the Gaussian predictive process model (Banerjee et al., 2008), the predetermined knots are tuning parameters. The TLR approximation (Abdulah et al., 2018) involves many tuning parameters, such as the matrix tile size, TLR maximum rank, TLR numerical accuracy, and optimization tolerance, which are introduced in Section 2. The tuning parameters should balance the computational burden and the estimation or prediction accuracy, so it is crucial to understand the impact of the tuning parameters on the statistical properties of the approximation methods. Finding a suggestion for these parameters serves as a motivation of our research, i.e., we would like to tune the TLR method input parameters, which can cut the computational time without losing too much estimation or prediction performance. The estimation performance is often evaluated via summary statistics and the plot of the estimations, such as the estimation variance (Kaufman et al., 2008) or the boxplots (Abdulah et al., 2018), but the prediction performance is not so straightforward to assess.

In the literature, the prediction performance is often evaluated by cross-validation. This method randomly leaves out *p* locations $\mathbf{s}_1, \ldots, \mathbf{s}_p$ from observation locations, and predicts the $Z(\mathbf{s}_1), \ldots, Z(\mathbf{s}_p)$ using the rest of the data at all other locations. Denote these predictions by $\hat{Z}(\mathbf{s}_1), \ldots, \hat{Z}(\mathbf{s}_p)$. The prediction performance is assessed by the deviation between the true and the predicted values, such as the Mean Square Prediction/Kriging Error (MSPE) (Abdulah et al., 2018)

$$MSPE = \frac{1}{p} \sum_{i=1}^{p} \{\hat{Z}(\boldsymbol{s}_i) - Z(\boldsymbol{s}_i)\}^2,$$
(3)

or the Mean Square Relative Prediction Error (Yan and Genton, 2018)

$$MSRPE = \frac{1}{p} \sum_{i=1}^{p} \left\{ \frac{\hat{Z}(\boldsymbol{s}_i) - Z(\boldsymbol{s}_i)}{Z(\boldsymbol{s}_i)} \right\}^2.$$

This performance can also be assessed by the deviation between the true observations and the corresponding predicted distributions, such as various kinds of proper scoring rules defined by Gneiting and Raftery (2007). For prediction intervals, the performance can be assessed by the empirical coverage of 95% prediction intervals on the left-out locations (Banerjee et al., 2008) or the interval score (Gneiting and Raftery, 2007). For more cross-validation based criteria, see Dai et al. (2007), Hengl et al. (2004), and Heaton et al. (2019). These criteria provide a straightforward measure of the performance of the prediction. However, they do not directly assess the loss of statistical efficiency when the approximated model is adopted instead of the true model, such as the extra Mean Square Errors (MSEs) caused by using the approximated model and the accuracy of estimated MSEs.

In the context of covariance model misspecification, Stein (1999) proposed the Loss of Efficiency (LOE) and the Misspecification of the MSE (MOM) criteria, based on the comparison of the MSEs between the true and the misspecified models. Using these criteria, Stein (1999) deduced that the simple kriging prediction is asymptotically optimal when the misspecified covariance model is equivalent to the true model. Stein (1999) also performed some simulations to assess the prediction performance of the kriging prediction under different settings of observation locations. However, all the results presented in Stein (1999) are for the case of a single prediction location.

In this article, we aim to give more appropriate criteria for the assessment of the loss of prediction efficiency when the true covariance model is approximated. Our suggested criteria can be used to assess the prediction efficiency of the approximation methods, e.g., the TLR method, with different tuning parameters, and help to choose the best value of these parameters. We suggest using the Mean Loss of Efficiency (MLOE) and the Mean Misspecification of the Mean Square Error (MMOM) criteria for multiple prediction locations as a generalization of the criteria proposed by Stein (1999). Here the MLOE and MMOM are relative errors. MLOE is strictly positive, while MMOM can be positive or negative at different locations. To avoid the possible issue of the cancellation of error over multiple locations in the MMOM criterion, we also introduce the Root mean square MOM (RMOM) criterion to evaluate the deviance of MOM from zero. Since the approximated covariance model can be viewed as a type of model misspecification, to show the MLOE, MMOM, and RMOM criteria are appropriate to assess the loss of prediction efficiency, we perform a similar simulation study from Stein (1999), where the exponential covariance model is misspecified as a Whittle covariance model plus a nugget effect, implying that the approximated covariance is smoother than the truth. Numerical results show that our criteria are better in assessing the prediction efficiency than the commonly used MSPE criterion and the Kullback-Leibler divergence criterion, which can be deduced from the logarithm score in Gneiting and Raftery (2007). As an application of our suggested criteria and a response to our research motivation, we use them to give a practical suggestion for selecting the tuning parameters in the TLR method, for which we investigate the performance of prediction and computation, using different tuning parameters from extensive simulation studies. For illustration of the validity of our suggested TLR tuning parameters, we fit a Gaussian-process model with a Matérn covariance function to a large spatial dataset of soil moisture in the area of the Mississippi basin; we then apply the TLR approximation method to obtain the MLEs and perform predictions with the suggested tuning parameters. Results show that our criteria are capable of selecting the tuning parameters of the TLR approximation since the TLR works well with our suggested parameters for the soil dataset. We also compare the TLR with the composite likelihood (Vecchia, 1988) and the Gaussian predictive process (Banerjee et al., 2008) for reference, suggesting that our criteria are successfully applied to the TLR tuning parameter selection in application. To the best of our knowledge, there is no previous work that considered choosing tuning parameters of an approximation method using a simulation-based criterion, rather than a data-based criterion such as the MSPE, for prediction performances.

The original LOE and MOM criteria proposed by Stein (1999) are for the assessment of model misspecification. We obtain its mean version for the understanding of the tuning parameters in

approximation methods for large spatial datasets, which can be used to help determine the tuning parameters in different approximation methods for real applications. For instance, to determine the tapering range in the covariance tapering method, one can first run a simulation for moderate datasets, using our criteria to compare the prediction efficiency with different tapering ranges. The best tapering range may be dependent on the range parameter α . Once the simulation results are obtained, one can choose the tapering range in the approximation problem for different real datasets, based on the simulation and a rough estimate of α . However, more advanced and accurate methods often involve more tuning parameters that require intensive simulation studies to understand the impact of each parameter and determine the tuning parameters that can provide the best trade-off between statistical properties and computational cost. Therefore we use a more delicate method (TLR) to show the effectiveness of our criteria.

The remainder of this article is organized as follows: Section 2 gives a brief background on the TLR approximation method and its tuning parameters. Section 3 introduces our suggested MLOE, MMOM, and RMOM criteria. In Section 4, we perform a simulation of the validity and sensitivity of the suggested criteria. In Section 5, we explain the simulation design to assess the TLR method, using different tuning parameters settings for which, our suggested criteria are used to measure the prediction accuracy and select the best specification of those tuning parameters. Section 6 shows the effectiveness of our suggested tuning parameters for the TLR method, using the soil moisture dataset. For this dataset, we also compare the estimation and prediction efficiency for the TLR method with our suggested parameters, with the composite likelihood and the Gaussian predictive process method in this section. Conclusions and discussions are provided in Section 7. More detailed numerical results about the specification of tuning parameters for the TLR method can be found in the Supplementary Material.

2. Tile Low-Rank (TLR) approximation

In this section, we give a brief background on the TLR approximation method, together with the tuning parameters associated with it in parallel hardware environments.

Tile-based algorithms have been developed on parallel architectures to speedup matrix-linear solver algorithms, for instance, PLASMA (Agullo et al., 2009) and Chameleon (cha, 2017) libraries. The given matrix is split into a set of tiles to allow the use of parallel execution, to a maximum degree, by weakening the synchronization points and bringing the parallelism in multithreaded BLAS (Blackford et al., 2002) to maximize the hardware utilization.

Since maximizing the log-likelihood in (2) and obtaining the MLE involves applying a set of linear-solver operations to the geospatial covariance matrix Σ , Abdulah et al. (2018a) have developed ExaGeoStat,¹ a framework that uses tile-based linear algebra algorithms to parallelize the MLE operations on leading-edge parallel hardware architecture. This framework has also been extended in Abdulah et al. (2018) to apply a TLR approximation to the covariance matrix. The new approximation technique aims at exploiting the data sparsity of the dense covariance matrix by compressing the off-diagonal tiles up to a user-defined accuracy threshold. The TLR method differs from existing low-rank approximation techniques, e.g., Banerjee et al. (2008), as the low-rank approximation is applied separately on each tile, instead of the whole matrix.

Fig. 1 gives an illustrative example of the TLR approximation method to a 8 × 8 covariance matrix, e.g., $\Sigma(\theta)$, where θ represents the parameter vector (i.e., variance, range, and smoothness parameters in the Matérn covariance function). Assuming a square positive-definite covariance matrix, the spatial covariance matrix with size $n \times n$ is divided into several tiles $D_{i,j}(\theta)$, where the size of each tile is $nb \times nb$. The Singular Value Decomposition (SVD) is used to approximate the off-diagonal tiles to a user-defined accuracy (i.e., tlr_acc , the tuning parameter argument used in ExaGeoStat as indicated in Table 1). In this case, the approximated tiles are the multiplication of two low-rank matrices, e.g., $D_{i,j}(\theta)$ is approximated by $\tilde{D}_{i,j}(\theta) = U_{i,j}(\theta)V_{i,j}(\theta)$, which can be deduced from the *k* most significant singular values and their associated left and right singular vectors.

¹ https://github.com/ecrc/exageostat.



Fig. 1. An illustrative example of a 8×8 covariance matrix TLR structure.

| Table 1 | | | | | | | | | | |
|-----------|-------|------|--------|------------|----|-----|-----|--------|----|-----|
| Arguments | for | the | tuning | parameters | of | the | TLR | method | in | the |
| ExaGeoSta | t fra | imew | ork. | | | | | | | |

| Name | Symbol |
|------------------------|--------------|
| Matrix tile size | nb |
| TLR maximum rank | tlr_max_rank |
| TLR numerical accuracy | tlr_acc |
| Optimization tolerance | opt_tol |

This approximation gives a data compression format that requires less memory and offers a faster computational speed of the matrix algebra. In the ExaGeoStat software (Abdulah et al., 2018), the TLR approximation is performed by the Hierarchical Computations on Manycore Architectures (HiCMA) numerical library (Abdulah et al., 2019a), which allows running the approximation on parallel systems with the help of StarPU (Augonnet et al., 2011).

Applying the TLR approximation to the log-likelihood function requires tuning several inputs to control the performance and accuracy of the approximation, namely, nb, tlr_max_rank , tlr_acc , and opt_tol as shown in Table 1; nb controls the size of each tile $D_{i,j}(\theta)$, and tlr_max_rank determines the maximum possible rank of the approximated tiles, which affects the memory allocation process of the approximating low-rank matrices $U_{i,j}(\theta)$ and $V_{i,j}(\theta)$ in the HiCMA library. By adopting the suggested criteria when assessing the prediction efficiency, we herein determine the best combination of the four TLR inputs by tuning these inputs, and by evaluating the performance and the accuracy of the approximated MLE compared with the exact MLE solution.

The effectiveness of the TLR approximation method can be improved by well tuning these four inputs. For instance, the current implementation of TLR in HiCMA uses a fixed-rank method to allocate and process all the given matrix tiles, although different approximated tiles have different ranks. A value of *tlr_max_rank* that is too large causes unnecessary memory usage and more data movements in the case of distributed memory architectures, whereas a too small value may cause a failure in approximating the tile. Thus, the best value of *tlr_max_rank* should be the smallest possible value that makes the approximation feasible for all the off-diagonal tiles. The accuracy threshold *tlr_acc* is also important to control the approximation accuracy, such that the approximation $D_{i,i}(\theta)$ of each tile satisfies $\|\hat{D}_{i,i}(\theta) - D_{i,i}(\theta)\|_2 \le tlr_{acc}$, where $\|\cdot\|_2$ is the L₂-norm of a matrix. A lower accuracy (larger *tlr_acc*) brings the arithmetic intensity of the approximation close to the memory-bound regime, whereas a higher accuracy makes the approximation run in the computebound regime (Abdulah et al., 2018). Thus, the accuracy threshold is an application-specific value. Furthermore, the optimization tolerance opt_tol is the minimum difference between two loglikelihood values at different iterations to control the optimization convergence condition. More specifically, the iteration process of computing the maximum point stops when $|l(\theta_{opt}) - l(\theta_{sub})| \le 1$ *opt_tol*, where $l(\theta_{opt})$ is the largest value of the log-likelihood function over all iterations and $l(\theta_{sub})$ is the second largest one.

We hope the tuning parameters in the TLR can save as much computational time as possible, without losing too much in estimating the accuracy of the prediction. Abdulah et al. (2018) investigated the impact of *tlr_acc* by showing the boxplots of the estimated parameters and the MSPEs. Here, our work uses the more informative MLOE, MMOM, and RMOM criteria for assessing the spatial prediction efficiency, which we describe in more detail in Section 3.

In this study, we use the ExaGeoStatR² package to perform the experiments relating to the TLR approximation. ExaGeoStatR (Abdulah et al., 2019b) is the R-wrapper interface of ExaGeoStat developed to facilitate the exploitation of large-scale capabilities in the R environment. The package provides parallel computation to evaluate the Gaussian maximum likelihood function using shared memory, GPUs, and distributed systems, by mitigating its memory space and computing restrictions. This package provides three ways of computing the MLE on a large scale: exact, Diagonal Super Tile (DST) approximation (i.e., covariance tapering), and TLR approximation. We are targeting the R functions related to the TLR approximation. The function tlr_mle() in the ExaGeoStatR package allows the computation of the TLR approximation of the MLE for the Matérn covariance model. This function computes the estimation by substituting the covariance matrix with its TLR approximation in the exact MLE framework.

3. Efficiency criteria for approximated spatial predictions

In this section, we construct three criteria for assessing the spatial prediction accuracy when the covariance matrix in the log-likelihood in (2) is approximated. Our first two criteria are of the averaged form of the criteria called the Loss of Efficiency (LOE) and the Misspecification of the MSE (MOM), proposed by Stein (1999), in the context of spatial prediction with a misspecified covariance model. Our last criterion is the mean square form of the MOM to measure the MOM variability on different prediction locations.

We consider a zero-mean Gaussian random field $Z(\mathbf{s})$, where the observations are $\mathbf{Z} = \{Z(\mathbf{s}_1), \ldots, Z(\mathbf{s}_n)\}^{\top}$. When the covariance model is true, the kriging prediction of $Z(\mathbf{s}_0)$ at a point \mathbf{s}_0 is $\hat{Z}_t(\mathbf{s}_0) = \mathbf{k}_t^{\top} K_t^{-1} \mathbf{Z}$, with MSE given by MSE $(\mathbf{s}_0) = E_t \{e_t^2(\mathbf{s}_0)\} = k_{0t} - \mathbf{k}_t^{\top} K_t^{-1} \mathbf{k}_t$, where $e_t(\mathbf{s}_0) = \hat{Z}_t(\mathbf{s}_0) - Z(\mathbf{s}_0)$ is the error of the kriging predictor, $K_t = \operatorname{Cov}_t \{\mathbf{Z}, \mathbf{Z}^{\top}\}$, $\mathbf{k}_t = \operatorname{Cov}_t \{\mathbf{Z}, Z(\mathbf{s}_0)\}$, $k_{0t} = \operatorname{Var}_t \{Z(\mathbf{s}_0)\}$, E_t , Var_t , and Cov_t mean the expectation, variance, and covariance with respect to the true covariance model. However, when the covariance is approximated, the kriging predictor is $\hat{Z}_a(\mathbf{s}_0) = \mathbf{k}_a^{\top} K_a^{-1} \mathbf{Z}$ instead, where $K_a = \operatorname{Cov}_a \{\mathbf{Z}, \mathbf{Z}^{\top}\}$, $\mathbf{k}_a = \operatorname{Cov}_a \{\mathbf{Z}, Z(\mathbf{s}_0)\}$, and Cov_a means the covariance is computed under the approximated covariance model. Denoting the error of this predictor by $e_a(\mathbf{s}_0) = \hat{Z}_a(\mathbf{s}_0) - Z(\mathbf{s}_0)$, then the MSE of this prediction is actually $E_t \{e_a^2(\mathbf{s}_0)\} = k_{0t} - 2\mathbf{k}_t^{\top} K_a^{-1} \mathbf{k}_a + \mathbf{k}_a^{\top} K_a^{-1} \mathbf{k}_a$, and the calculated result of MSE is $E_a \{e_a^2(\mathbf{s}_0)\} = k_{0a} - \mathbf{k}_a^{\top} K_a^{-1} \mathbf{k}_a$, where $k_{0a} = \operatorname{Var}_a \{Z(\mathbf{s}_0)\}$, E_a and Var_a mean that the expectation and variance are computed using the approximated covariance model. Thus, following Stein (1999), the Loss of Efficiency of the prediction is defined as

$$\text{LOE}(\mathbf{s}_0) = \text{E}_t \{ e_a^2(\mathbf{s}_0) \} / \text{E}_t \{ e_t^2(\mathbf{s}_0) \} - 1,$$
(4)

and the Misspecification of the MSE is defined as

$$MOM(\mathbf{s}_0) = E_a\{e_a^2(\mathbf{s}_0)\} / E_t\{e_a^2(\mathbf{s}_0)\} - 1.$$
(5)

Our first two criteria are defined as the mean value of the Loss of Efficiency (4) and Misspecification of the MSE (5) over multiple prediction locations. More specifically, when the prediction locations are s_{01}, \ldots, s_{0m} , the Mean Loss of Efficiency is defined as

$$MLOE = \frac{1}{m} \sum_{i=1}^{m} LOE(\boldsymbol{s}_{0i}),$$
(6)

and the Mean Misspecification of the MSE is defined as

$$MMOM = \frac{1}{m} \sum_{i=1}^{m} MOM(\boldsymbol{s}_{0i}).$$
(7)

² https://github.com/ecrc/exageostatR.

Y. Hong, S. Abdulah, M.G. Genton et al.

For Gaussian random fields, the kriging predictor $\hat{Z}_t(\mathbf{s}_0)$ is the best predictor in terms of minimizing MSE (Stein, 1999), so $E_t\{e_t^2(\mathbf{s}_0)\} \leq E_t\{e_a^2(\mathbf{s}_0)\}$ and $LOE(\mathbf{s}_0) \geq 0$ for any \mathbf{s}_0 . However, $MOM(\mathbf{s}_{0i})$ may be positive or negative. If there are two MOM values which have opposite sign and large absolute values, they will eliminate each other in (7), causing an over optimistic MMOM result, although we believe that two MOM values with opposite signs can be considered better than with the same sign. To avoid this problem, which we call the canceling of error problem, we also define the following Root mean square MOM (RMOM) criterion:

$$\text{RMOM} = \sqrt{\frac{1}{m} \sum_{i=1}^{m} \{\text{MOM}(\boldsymbol{s}_{0i})\}^2}.$$
(8)

We choose the prediction locations of a regular grid in the observation region, so the value of MLOE, MMOM, and RMOM can describe the average prediction performance over the whole observation region. For instance, when the observation region is $[0, 1]^2$, the prediction locations can be (i/5, j/5) for i, j = 1, 2, 3, 4. The MLOE describes the average efficiency loss of the prediction when the approximated covariance model is used instead of the true one, whereas the MMOM describes the average misspecification between the computed and true MSEs. The RMOM describes the degree of deviance of the misspecification between the computed and true MSEs from zero.

In the case where the approximated model is an estimated model, Stein (1999) proposed alternative estimations for the MSE term $E_t\{e_a^2(\mathbf{s}_0)\}$ in (4) and (5). Let the random field $Z(\mathbf{s})$ follow a parametric model, where the unknown parameter is θ . Denote the true value of this parameter by θ_0 and the estimated value by $\hat{\theta} = \hat{\theta}(\mathbf{Z})$. Here the approximated model is related to a random variable $\hat{\theta}$, so the $E_t\{e_a^2(\mathbf{s}_0)\}$ and $E_a\{e_a^2(\mathbf{s}_0)\}$ terms in the LOE and MOM definitions are substituted by their corresponding estimations, such as the plug-in estimation. Denote by E_{θ} , $\operatorname{Var}_{\theta}$, and $\operatorname{Cov}_{\theta}$ the expectation, variance, and covariance computed using the parametric model with parameter θ , respectively; $\hat{Z}_{\theta}(\mathbf{s}_0)$ the kriging predictor under the parameter θ ; $e_{\theta}(\mathbf{s}_0) = \hat{Z}_{\theta}(\mathbf{s}_0) - Z(\mathbf{s}_0)$ the prediction error; $K_{\theta} = \operatorname{Cov}_{\theta}\{\mathbf{Z}, \mathbf{Z}^{\top}\}$; $\mathbf{k}_{\theta} = \operatorname{Cov}_{\theta}\{\mathbf{Z}, Z(\mathbf{s}_0)\}$; $k_{0,\theta} = \operatorname{Var}_{\theta}\{Z(\mathbf{s}_0)\}$. Thus, $E_t\{e_t^2(\mathbf{s}_0)\} = E_{\theta_0}\{e_{\theta_0}^2(\mathbf{s}_0)\} = k_{0,\theta_0} - \mathbf{k}_{\theta_0}^{\top}\mathbf{k}_{\theta_0}$. The $E_t\{e_a^2(\mathbf{s}_0)\}$ and $E_a\{e_a^2(\mathbf{s}_0)\}$ terms can be estimated by the following plug-in estimation:

$$\mathbf{E}_{t}\{\boldsymbol{e}_{a}^{2}(\boldsymbol{s}_{0})\} = \mathbf{E}_{\boldsymbol{\theta}_{0}}\{\boldsymbol{e}_{\hat{\boldsymbol{\theta}}}^{2}(\boldsymbol{s}_{0})\} \approx \mathbf{E}_{\boldsymbol{\theta}_{0}}\{\boldsymbol{e}_{\boldsymbol{\theta}}^{2}(\boldsymbol{s}_{0})\}|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = k_{0,\boldsymbol{\theta}_{0}} - 2\boldsymbol{k}_{\boldsymbol{\theta}_{0}}^{\top}\boldsymbol{K}_{\hat{\boldsymbol{\theta}}}^{-1}\boldsymbol{k}_{\hat{\boldsymbol{\theta}}} + \boldsymbol{k}_{\hat{\boldsymbol{\theta}}}^{\top}\boldsymbol{K}_{\hat{\boldsymbol{\theta}}}^{-1}\boldsymbol{K}_{\boldsymbol{\theta}_{0}}\boldsymbol{K}_{\hat{\boldsymbol{\theta}}}^{-1}\boldsymbol{k}_{\hat{\boldsymbol{\theta}}},$$
(9)

$$\mathbf{E}_{a}\{\boldsymbol{e}_{a}^{2}(\boldsymbol{s}_{0})\} \approx \mathbf{E}_{\theta}\{\boldsymbol{e}_{\theta}^{2}(\boldsymbol{s}_{0})\}|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = k_{0,\hat{\boldsymbol{\theta}}} - \boldsymbol{k}_{\hat{\boldsymbol{\theta}}}^{\top}\boldsymbol{K}_{\hat{\boldsymbol{\theta}}}^{-1}\boldsymbol{k}_{\hat{\boldsymbol{\theta}}}.$$
(10)

Stein (1999) noted that, when $\{\mathbf{Z}^{\top}, Z(\mathbf{s}_0)\}^{\top}$ is Gaussian, the conditional distribution of $e_{\hat{\theta}}(\mathbf{s}_0)$ given $\mathbf{Z} = \mathbf{z}$ is $N(e_{\hat{\theta}(\mathbf{z})}(\mathbf{s}_0) - e_{\theta_0}(\mathbf{s}_0), E_{\theta_0}\{e_{\theta_0}^2(\mathbf{s}_0)\})$, so

$$\mathsf{E}_{\theta_0}\{e^2_{\hat{\theta}(\mathbf{Z})}(\mathbf{s}_0)|\mathbf{Z}=\mathbf{Z}\}=\mathsf{E}_{\theta_0}\{e^2_{\theta_0}(\mathbf{s}_0)\}+\{e_{\hat{\theta}(\mathbf{Z})}(\mathbf{s}_0)-e_{\theta_0}(\mathbf{s}_0)\}^2.$$

Therefore, the $E_t \{e_a^2(\mathbf{s}_0)\}$ term can be estimated by

$$E_t\{e_a^2(\mathbf{s}_0)\} \approx E_{\theta_0}\{e_{\theta_0}^2(\mathbf{s}_0)\} + \{e_{\hat{\theta}}(\mathbf{s}_0) - e_{\theta_0}(\mathbf{s}_0)\}^2,\tag{11}$$

and $E_a\{e_a^2(\mathbf{s}_0)\}$ is still estimated by (10). When our suggested criteria are computed using (9) and (10), we say that the criteria are computed using the plug-in method. When our criteria are computed using (10) and (11), we say that the criteria are computed using Stein's method.

In the plug-in method, the computed $E_{\theta_0}\{e^2_{\theta}(\mathbf{s}_0)\}|_{\theta=\hat{\theta}}$ may be very slightly smaller than $E_{\theta_0}\{e^2_{\theta_0}(\mathbf{s}_0)\}$, possibly due to round-off error. In the subsequent simulations of this article, the smallest value of the computed $E_{\theta_0}\{e^2_{\theta}(\mathbf{s}_0)\}|_{\theta=\hat{\theta}}-E_{\theta_0}\{e^2_{\theta_0}(\mathbf{s}_0)\}$ is -6.3101×10^{-15} . In this case, we estimate the $E_t\{e^2_a(\mathbf{s}_0)\}$ term by $E_{\theta_0}\{e^2_{\theta_0}(\mathbf{s}_0)\}$ instead, to keep LOE (4) nonnegative on all prediction locations. In Stein's method, the computed LOE is always nonnegative, which is better than the plug-in method. However, Stein's method did not consider the model misspecification case, so we particularly recommend Stein's method for the case when the parametric model is correctly specified, and recommend the plug-in method for more general cases, e.g., when the parametric

model is misspecified. For simulations of this article, we will compute our suggested criteria using both of the methods.

Stein (1999) also introduced a resampling method to better estimate $E_t\{e_a^2(\mathbf{s}_0)\}$. This method first generates n_r independent samples of \mathbf{Z} and computes the estimate $\hat{\boldsymbol{\theta}}^{(1)}, \ldots, \hat{\boldsymbol{\theta}}^{(n_r)}$, then computes the kriging error terms $e_{\hat{\boldsymbol{\theta}}}(\mathbf{s}_0)$ and $e_{\theta_0}(\mathbf{s}_0)$ for each sample, which are denoted by $e_{\hat{\boldsymbol{\theta}}^{(j)}}^{(j)}(\mathbf{s}_0)$ and $e_{\theta_0}^{(j)}(\mathbf{s}_0)$, $j = 1, \ldots, n_r$, respectively. Since $E_{\theta_0}\{e_{\theta_0}^2(\mathbf{s}_0)\}$ remains unchanged for resampling, we have the following estimation:

$$\mathbb{E}_{t}\{e_{a}^{2}(\boldsymbol{s}_{0})\} \approx \mathbb{E}_{\boldsymbol{\theta}_{0}}\{e_{\boldsymbol{\theta}_{0}}^{2}(\boldsymbol{s}_{0})\} + \frac{1}{n_{r}}\sum_{j=1}^{n_{r}}\{e_{\hat{\boldsymbol{\theta}}^{(j)}}^{(j)}(\boldsymbol{s}_{0}) - e_{\boldsymbol{\theta}_{0}}^{(j)}(\boldsymbol{s}_{0})\}^{2}.$$

In the simulation framework, such resampling method is equivalent to estimate $E_t\{e_a^2(s_0)\}$ using Stein's method (11) with n_r replicates and report the mean value as the final result. When the number of replicates of the simulation is large enough, the increment of samples with a price of more computational burdens may not be necessary. Therefore, we will not perform this resampling in our simulation.

The computation of MLOE, MMOM, and RMOM criteria involves the inversion of the covariance matrix. In the simulation of Section 5, where the number of observations is n = 3,600, the computational times for these criteria are acceptable. If the direct computation of these criteria is not available due to the data size, one can adopt a matrix compression method, such as the TLR (Akbudak et al., 2017; Abdulah et al., 2018). This compression can provide an approximation of the covariance matrix and save computational time.

4. Simulation on the validity of the suggested criteria

We perform a numerical simulation to illustrate the validity and sensitivity of the suggested criteria, compared with the popular Mean Square Prediction Error (MSPE) criterion and the Kullback– Leibler Divergence criterion, which can be deduced by the logarithmic score introduced by Gneiting and Raftery (2007). Similar to the settings in Stein (1999), we focus on the case where the covariance model is misspecified.

In this simulation, we consider a zero-mean stationary Gaussian random field $\{Z(\mathbf{s}), \mathbf{s} \in [0, 1]^2\}$ with Matérn covariance function (1). We set the true covariance model as the exponential model, with covariance function $C(h; \theta = (\sigma^2, \alpha, 0.5)^{\top})$, and consider two cases of model misspecification. In the first case, the covariance model is correctly specified, but the parameters σ^2 and α are misspecified as their maximum likelihood estimate. Under this kind of misspecification, the corresponding kriging prediction is called 'empirical best linear unbiased prediction' (EBLUP); the EBLUP does not significantly affect the prediction efficiency, according to the intuition and simulation results in the literature (Stein, 1999). In the second case, the covariance model is misspecified as a smoother (Whittle) covariance model plus a nugget effect term. In this case, the misspecified covariance function is $C(h; \theta = (\sigma^2, \alpha, 1.0)^{\top}) + \tau^2 I_{h=0}(h)$, where τ^2 is the nugget variance and $I_{h=0}$ is the indicator function. In the R package fields, the function for fitting a covariance model MLESpatialProcess() chooses the smoothness parameter $\nu = 1$ as the default value, which is smoother than the common setting $\nu = 0.5$. This motivates us to investigate the loss of prediction efficiency for this case.

The observation locations are set to be $\mathbf{s}_{r,l} = n^{-1/2}(r - 0.5 + \tilde{U}_{r,l}, l - 0.5 + \tilde{V}_{r,l})$, where n is the number of observations, $\tilde{U}_{r,l}$ and $\tilde{V}_{r,l}$ are *i.i.d.* samples from the uniform distribution U[-0.4, 0.4]. By ordering r, l lexicographically, these locations are also denoted by $\mathbf{s}_1, \ldots, \mathbf{s}_n$. We take $n = 12^2, 24^2$, or 48^2 . The true covariance function is set as $C(h; \boldsymbol{\theta} = (\sigma^2, \alpha, 0.5)^{\top})$, where $\sigma^2 = 1, \alpha = 0.2/(-\log(0.05))$, such that the true effective range of the model is 0.2. For each parameter setting, we generate 100 independent replications from the random field with the true covariance model at the same observation locations. First, we compute the MLE for the correctly specified and misspecified covariance models, then, we compute the plug-in kriging predictions for both covariance functions at the point (i/5, j/5), for i, j = 1, 2, 3, 4, denoted by $\mathbf{s}_{01}, \ldots, \mathbf{s}_{0p}$ for

p = 16, and compare the results with the kriging results from the exact covariance functions. Lastly, the performance of the prediction is comparatively assessed, using our suggested MLOE (6), MMOM (7), and RMOM (8) computed by the plug-in method and Stein's method, as well as MSPE in (3) and the Kullback–Leibler Divergence (K–L Divergence) criterion. As we have discussed in Section 3, the plug-in method is more suitable for computing our suggested criteria in this simulation of model misspecification.

The K–L divergence has been used to assess the estimation performance by comparing the approximated likelihood to the exact one (Huang and Sun, 2018). For predictions, we need to compare two predictive distributions. Let Q_t be the distribution of $\mathbf{Z}_p := \{Z(\mathbf{s}_{01}), \ldots, Z(\mathbf{s}_{0m})\}^{\top}$, conditional to the observations $\mathbf{Z} = \{Z(\mathbf{s}_1), \ldots, Z(\mathbf{s}_n)\}^{\top}$, computed using the true model, and Q_a be the computed distribution of \mathbf{Z}_p conditional to \mathbf{Z} using the approximated model. Denoting these two conditional distributions by $\{\mathbf{Z}_p | \mathbf{Z}\}_t$ and $\{\mathbf{Z}_p | \mathbf{Z}\}_a$, respectively, then the Kullback–Leibler Divergence is denoted by

$$D_{KL}(Q_t \parallel Q_a) = D_{KL}\left(\{\boldsymbol{Z}_p | \boldsymbol{Z}\}_t \parallel \{\boldsymbol{Z}_p | \boldsymbol{Z}\}_a\right) = \int \log\left\{\frac{q_t(\boldsymbol{Z}_p | \boldsymbol{Z})}{q_a(\boldsymbol{Z}_p | \boldsymbol{Z})}\right\} q_t(\boldsymbol{Z}_p | \boldsymbol{Z}) \mathrm{d}\boldsymbol{Z}_p,$$

where q_t and q_a are the conditional distributions corresponding to the true and the approximated model, respectively. When $Q_t \sim N(\mu_{Q_t}, \Sigma_{Q_t})$ and $Q_a \sim N(\mu_{Q_a}, \Sigma_{Q_a})$, the K–L divergence between these two multivariate Gaussian distribution satisfies

$$D_{KL}(Q_t \parallel Q_a) = \frac{1}{2} \left\{ \text{trace}(\boldsymbol{\varSigma}_{Q_a}^{-1} \boldsymbol{\varSigma}_{Q_t}) - \log \det(\boldsymbol{\varSigma}_{Q_a}^{-1} \boldsymbol{\varSigma}_{Q_t}) + (\boldsymbol{\mu}_{Q_a} - \boldsymbol{\mu}_{Q_t})^\top \boldsymbol{\varSigma}_{Q_a}^{-1} (\boldsymbol{\mu}_{Q_a} - \boldsymbol{\mu}_{Q_t}) - m \right\}, \quad (12)$$

where *m* is the dimension of Q_t or Q_a .

The K–L Divergence criterion comes from the logarithmic score criterion introduced by Gneiting and Raftery (2007). Let $\mathbf{x} \in \mathbb{R}^m$ be the *m*-dimensional observed value and \tilde{P} be the predicted distribution for this value, where \tilde{P} is assumed to be only related to its mean $\mu_{\tilde{P}}$ and covariance matrix $\Sigma_{\tilde{P}}$. Then the scoring rule

$$S(\tilde{P}, \mathbf{x}) = -\log \det \Sigma_{\tilde{P}} - (\mathbf{x} - \boldsymbol{\mu}_{\tilde{P}})^{\top} \Sigma_{\tilde{P}}^{-1} (\mathbf{x} - \boldsymbol{\mu}_{\tilde{P}})$$

is strictly proper relative to the class of Gaussian measures and is equivalent to the logarithmic score (Gneiting and Raftery, 2007). Therefore we call this scoring rule the logarithmic score. By Gneiting and Raftery (2007), the divergence function for this rule is

$$d(\tilde{P}, \tilde{Q}) = \operatorname{trace}(\boldsymbol{\Sigma}_{\tilde{P}}^{-1}\boldsymbol{\Sigma}_{\tilde{Q}}) - \log \operatorname{det}(\boldsymbol{\Sigma}_{\tilde{P}}^{-1}\boldsymbol{\Sigma}_{\tilde{Q}}) + (\boldsymbol{\mu}_{\tilde{P}} - \boldsymbol{\mu}_{\tilde{Q}})^{\top}\boldsymbol{\Sigma}_{\tilde{P}}^{-1}(\boldsymbol{\mu}_{\tilde{P}} - \boldsymbol{\mu}_{\tilde{Q}}) - m,$$

where \tilde{P} , \tilde{Q} are *m*-dimensional distributions with mean $\mu_{\tilde{P}}$, $\mu_{\tilde{Q}}$ and covariance matrix $\Sigma_{\tilde{P}}$, $\Sigma_{\tilde{Q}}$, respectively. Here the divergence function of a scoring rule is defined by $d(\tilde{P}, \tilde{Q}) := S(\tilde{Q}, \tilde{Q}) - S(\tilde{P}, \tilde{Q})$, where $S(\tilde{P}, \tilde{Q}) = \int S(\tilde{P}, \mathbf{x}) d\tilde{Q}(\mathbf{x})$. The $d(\tilde{P}, \tilde{Q})$ can be considered as a logarithm score divergence criterion, which equals two times of the K–L Divergence $D_{KL}(\tilde{Q} \parallel \tilde{P})$.

The simulation results are shown in Figs. 2–5. These figures show that, when the covariance model is correctly specified, the MLOE is very small in comparison with the exact model and has a decreasing trend when the number of observations *n* increases. The MMOM is larger, but concentrates near zero and shrinks when *n* increases. The RMOM also shrinks to zero when *n* increases. This shows that the plug-in kriging prediction does not lead to a significant loss of prediction efficiency, which is in agreement with the intuition and the simulation results introduced in Stein (1999). When the model is misspecified, the MLOE is clearly larger than that of the case where the model is correctly specified. The MMOM may likely have a mean value larger than zero, and the RMOM is larger than the correctly specified case. Thus, when a rougher covariance model is misspecified as a smoother model with a nugget effect, the plug-in prediction is suboptimal, and the MSE can be overestimated. In Fig. 4, the difference of boxplots between the case where the model is correctly specified or misspecified is not apparent, showing that our suggested MLOE, MMOM, and RMOM are more sensitive criteria for prediction accuracy. According to Fig. 5, the K–L Divergence is also a sensitive criterion for prediction accuracy, measuring the information loss when the predicted distribution is approximated. However, it cannot provide more detailed information



Fig. 2. Boxplots of the MLOE, MMOM, and RMOM computed by the plug-in method, with respect to the number of observations *n*, when the covariance model is correctly specified as the exponential model (True) or misspecified as the Whittle model (Presumed).



Fig. 3. Boxplots of the MLOE, MMOM, and RMOM computed by Stein's method, with respect to the number of observations *n*, when the covariance model is correctly specified as the exponential model (True) or misspecified as the Whittle model (Presumed).



Fig. 4. Boxplots of the MSPE for the predictions with respect to the exact model (Exact), plug-in prediction with correct covariance model (True), and the plug-in prediction with misspecified covariance model (Pres).

on the prediction for a spatial model that one may be interested in, such as the efficiency loss of MSE when the model is approximated and the accuracy of the computed MSE. In conclusion, our suggested criteria are valid, sensitive, and more informative tools to detect the loss of prediction efficiency caused by spatial model approximations in simulation studies.

5. Simulation experiments on the tuning parameters

As an application of the suggested MLOE, MMOM, and RMOM criteria in Section 3, we aim at assessing the performance of the TLR approximation method based on these criteria. We define how to tune the TLR associated inputs based on the target data and the application requirements,



Fig. 5. Boxplots of the K–L Divergence for the predictions with respect to the number of observations *n*, when the covariance model is correctly specified as the exponential model (True) or misspecified as the Whittle model (Presumed).

using simulation experiments, which is the answer to the motivation of our study. All experiments being carried out in this section are conducted on a dual-socket 8-core Intel Sandy Bridge-based Xeon E5-2670 CPU running at 2.60 GHz.

5.1. Simulation settings

Here, we provide an outline of our simulation settings. Similar to the settings in Sun and Stein (2016), our simulation experiments are performed on a set of synthetic datasets generated using the built-in data generator tool in ExaGeoStatR at irregular locations in a 2D space (i.e., simulate_data_exact() function). The generation process assumes a zero-mean stationary Gaussian random field $\{Z(s), s \in [0, 1]^2\}$. The observation locations s_1, \ldots, s_n are generated by the same settings as those detailed in Section 4. Given the set of *n* locations, the covariance matrix Σ is constructed using the Matérn covariance function.

The simulation is to illustrate the effectiveness of using the TLR approximation method for the MLE estimation. The assessments include the total execution time, estimation accuracy, and prediction accuracy. Instead of the MSPE criterion in Abdulah et al. (2018), here the prediction accuracy is investigated by MLOE, MMOM, and RMOM, using the plug-in method and Stein's method stated in Section 3, whose effectiveness and sensitivity have been shown in Section 4. The assessment includes the kriging performance obtained by using the estimated parameters to predict unknown sets of values at various specific locations. Unless otherwise specified, our suggested criteria are computed by the plug-in method. Results of Stein's method are shown in Tables 13–16 in Supplementary Material, indicating that all conclusions drawn from the plug-in method and Stein's method are consistent. All the symbols in this section follow the abbreviations illustrated in Table 1.

In the simulation experiments conducted by Abdulah et al. (2018), both the estimation and the prediction accuracy of the TLR method were shown by a set of boxplots representing the estimation accuracy of different model parameters and the MSPE, and compared with the exact method. The simulation performed in Abdulah et al. (2018) also assessed the impact of using different TLR accuracy levels *tlr_acc* for both the accuracy and the execution time. Here, we use our suggested criteria and not only consider *tlr_acc*, but also consider the impact of other tuning parameters, i.e., tile size, maximum rank, and optimization tolerance to the overall execution time, estimation accuracy, and prediction accuracy. Moreover, we consider two different smoothness levels of the underlying random field, i.e., v = 0.5 and 1, whereas the simulations of Abdulah et al. (2018) only considered v = 0.5.

Y. Hong, S. Abdulah, M.G. Genton et al.

All the experiments in this section use spatial data where the number of locations is n = 3600. For the true values of the parameters in (1), we consider $\sigma^2 = 1$, $\nu = 0.5$ or 1, and α is chosen such that the effective range of the model can be $h_{\text{eff}} = 0.2, 0.4, 0.8$, or 1.6. First, a set of the following experiments aims at comparing the performance of the TLR approximation under different tile size *nb* with suitable value of maximum rank *tlr_max_rank*, while the other tuning parameters *tlr_acc* and *opt_tol* are fixed at a moderate value which does not affect the estimation accuracy of the approximations. Second, we compare the performance under different accuracy levels *tlr_acc* and optimization tolerance *opt_tol*, where the tile size *nb* and maximum rank *tlr_max_rank* are fixed at the suggested value obtained in the previous step. The reason for adopting these two steps is, according to Abdulah et al. (2018), that *nb* and *tlr_max_rank* mainly affect the computational time, whereas *tlr_acc* and *opt_tol* mainly affect the prediction efficiency. Recall that a larger *tlr_acc* corresponds to a coarser tile approximation, so the maximum rank necessary for the approximation is smaller. Thus the value of *tlr_max_rank* does not directly affect the prediction efficiency; it could affect the efficiency via different *tlr_acc*.

5.2. Performance using different tile sizes

The parallel TLR approximation computation depends on dividing the matrix into a set of tiles where the tile size is $nb \times nb$. Here, nb should be tuned in different hardware platforms to obtain the best performance that corresponds to the trade-off between the arithmetic intensity and the degree of parallelism. We illustrate the performance and accuracy using different values of nb, i.e., nb = 400, 450, 600, and 900. We fix $tlr_acc = 10^{-9}$ and $opt_tol = 10^{-6}$ since these values have little impact on the estimation performance. The tlr_max_rank is fixed to the smallest feasible value for TLR computation obtained before the simulation. The tlr_max_rank actually affects the memory allocation process and communication cost in case of distributed memory systems. A value of tlr_max_rank that is too large can slow down the computation due to the unnecessary allocations, whereas a too small value may cause the failure of the SVD approximation of each off-diagonal tile. Thus, for each value of nb, we try to compute TLR approximations for $tlr_max_rank = 10, 20, \ldots$, until the value of tlr_max_rank can make the approximation feasible for all replicates.

We generate 100 independent replications from the random field with the true covariance model for each parameter setting. The synthetic datasets are generated using the simulate_data_exact() function in the ExaGeoStatR package. The estimation performed uses both the exact and TLR methods, by the exact_mle() and the tlr_mle() functions in the same package, respectively, and estimates both the execution time and the estimation accuracy of each method, for different *nb* values. The last step is to compute the MLOE, MMOM, and RMOM on prediction locations (*i*/5, *j*/5), where *i*, *j* = 1, 2, 3, 4. The prediction performance is then evaluated by the mean and standard deviations for both the values of our criteria. In our estimation, the value of ν is fixed at its true value and the optimization bound for estimating σ^2 and α is [0.01, 5]. The optimization tolerance of the exact MLE is set as 10^{-9} in order to get more accurate estimation results for comparison.

Selecting the smallest tlr_max_rank value for each tile size is important to obtain the best performance. Thus, we perform a set of experiments to select the tlr_max_rank value corresponding to each nb when n = 3600 (Table 2). The reported values show that the feasible tlr_max_rank does not simply increase when the tile size nb increases. In fact, when the number of tiles is divisible by the number of underlying CPUs, i.e., nb = 450 or 900, the maximum rank of each tile tlr_max_rank is relatively small compared to the nb. The required tlr_max_rank is significantly smaller when the model is relatively smoother. Thus, when the number of locations is n = 3600, we recommend to choose the tlr_max_rank as the largest values shown in Table 2 for the corresponding values of ν and nb.

For the MLE and different TLR approximations, we only show typical results with v = 0.5, $h_{\text{eff}} = 0.2$ and v = 1, $h_{\text{eff}} = 1.6$, shown in Table 3. This table shows that the TLR approximation has a similar estimation and prediction performance for different tile sizes *nb*, whereas the fastest computational time is obtained when nb = 450.

Table 2

Smallest *tlr_max_rank* that makes the TLR approximation applicable to different values of *nb*, and the parameters of the Matérn covariance. The number of locations is n = 3600.

| ν | Eff.range | Tile size (nb) | | | | | | |
|-----|-----------|----------------|-----|-----|-----|--|--|--|
| | | 400 | 450 | 600 | 900 | | | |
| | 0.2 | 260 | 210 | 310 | 270 | | | |
| 0.5 | 0.4 | 250 | 210 | 310 | 270 | | | |
| 0.5 | 0.8 | 250 | 210 | 300 | 260 | | | |
| | 1.6 | 250 | 210 | 300 | 260 | | | |
| | 0.2 | 220 | 170 | 250 | 210 | | | |
| 10 | 0.4 | 220 | 170 | 250 | 210 | | | |
| 1.0 | 0.8 | 210 | 170 | 250 | 210 | | | |
| | 1.6 | 210 | 180 | 250 | 210 | | | |

Table 3

Estimation and prediction performances of MLE and TLR approximation estimates for different values of *nb*. Bias(·) means the estimate of the parameter minus its true value, whereas the estimation time means the computational time of the corresponding estimation. The value of MLOE for all cases ($\nu = 0.5$ and $\nu = 1.0$) is multiplied by 10⁶.

| Mean (sd) | v = 0.5, l | $h_{\rm eff}=0.2$ | | | | $v = 1.0, h_{\rm eff} = 1.6$ | | | | | |
|------------------------|------------|---|----------|----------|----------|------------------------------|----------|-------------------------|----------|----------|--|
| | MLE | $\frac{\text{TLR approximations }(nb)}{400 450 600 900}$ | | | | MLE | TLR appr | TLR approximations (nb) | | | |
| | | 400 | 450 | 600 | 900 | | 400 | 450 | 600 | 900 | |
| $Bias(\sigma^2)$ | -0.0080 | -0.0079 | -0.0079 | -0.0079 | -0.0079 | 0.0163 | 0.2173 | 0.2236 | 0.2210 | 0.2153 | |
| | (0.0908) | (0.0908) | (0.0908) | (0.0908) | (0.0908) | (0.6546) | (0.7339) | (0.7793) | (0.7624) | (0.7419) | |
| $Bias(\alpha)$ | -0.0006 | -0.0006 | -0.0006 | -0.0006 | -0.0006 | -0.0144 | 0.0577 | 0.0577 | 0.0579 | 0.0571 | |
| | (0.0063) | (0.0063) | (0.0063) | (0.0063) | (0.0063) | (0.1186) | (0.1348) | (0.1394) | (0.1370) | (0.1357) | |
| MLOE ($\times 10^6$) | 3.3945 | 3.3756 | 3.3756 | 3.3758 | 3.3756 | 0.0273 | 0.0109 | 0.0110 | 0.0109 | 0.0109 | |
| | (5.9930) | (5.9474) | (5.9474) | (5.9477) | (5.9475) | (0.0669) | (0.0242) | (0.0243) | (0.0242) | (0.0241) | |
| ММОМ | 0.0017 | 0.0011 | 0.0011 | 0.0011 | 0.0011 | 0.0014 | -0.1428 | -0.1428 | -0.1428 | -0.1428 | |
| | (0.0232) | (0.0232) | (0.0232) | (0.0232) | (0.0232) | (0.0227) | (0.0223) | (0.0222) | (0.0223) | (0.0223) | |
| RMOM | 0.0185 | 0.0185 | 0.0185 | 0.0185 | 0.0185 | 0.0182 | 0.1428 | 0.1428 | 0.1428 | 0.1428 | |
| | (0.0141) | (0.0140) | (0.0140) | (0.0140) | (0.0140) | (0.0135) | (0.0223) | (0.0222) | (0.0223) | (0.0223) | |
| Estimation | 146.5 | 110.2 | 90.3 | 146.4 | 146.6 | 277.5 | 122.3 | 108.1 | 143.5 | 186.5 | |
| Time (s) | (20.2) | (15.3) | (13.5) | (21.7) | (19.5) | (75.6) | (35.6) | (31.6) | (46.6) | (63.0) | |

Remark 1. In Table 3, the standard deviation of the MLOE is larger than the corresponding mean. Fig. 6 shows the typical case of boxplots for the MLOE in the simulation, indicating that the distribution of MLOE is skewed to the right, causing a larger standard deviation.

The main reason for the larger standard deviation is, when a normal distributed fluctuation is introduced in the model parameters, the difference of kriging prediction results between the original model and the fluctuated model may have a heavy-tailed distribution. We have run a simple illustrative example to show this. Consider a stationary Gaussian random field $Z(\mathbf{s})$ with exponential covariance function $C(h; \theta) = \sigma^2 \exp(-h/\alpha)$, where the observation locations are $\mathbf{s}_1 = (0, 0), \mathbf{s}_2 =$ $(0, 1), \mathbf{s}_3 = (1, 0), \mathbf{s}_4 = (1, 1)$, and the prediction location is $\mathbf{s}_0 = (0.5, 0.5)$. We generate 10,000 replicates of the observations, where the parameter $\theta = (\sigma^2, \alpha)^{\top}$ has true value $\sigma_0^2 = 1, \alpha_0^2 = 0.1$. In each replicate, the presumed values of σ^2 and α are independently drawn from normal distributions $N(1, 0.01^2)$ and $N(0.1, 0.01^2)$. We computed the difference of kriging prediction $\hat{Z}_{\theta}(\mathbf{s}_0) - \hat{Z}_{\theta_0}(\mathbf{s}_0)$ and LOE(\mathbf{s}_0) for plug-in and Stein's methods, where θ is the presumed value of parameters. Results are shown in Fig. 7, indicating that the difference $\hat{Z}_{\theta}(\mathbf{s}_0) - \hat{Z}_{\theta_0}(\mathbf{s}_0)$ follows a heavy-tailed distribution and LOE(\mathbf{s}_0) is skewed to the right. The mean and standard deviation of LOE(\mathbf{s}_0) are 1.8323 × 10⁻⁶ and 4.2347 × 10⁻⁶ (Plug-in method) or 1.8647 × 10⁻⁶ and 8.1274 × 10⁻⁶ (Stein's method), respectively. Note that for Stein's method, LOE(\mathbf{s}_0) = $\{\hat{Z}_{\theta}(\mathbf{s}_0) - \hat{Z}_{\theta_0}(\mathbf{s}_0)\}^2/E_{\theta_0}\{e_{\theta_0}^2(\mathbf{s}_0)\}$, so the right-skewed distribution of the LOE comes from a square of a heavy-tailed distribution. The



Fig. 6. Boxplots of the MLOE corresponding to MLE and TLR approximation estimates for different values of *nb*, where $\nu = 0.5$ and $h_{eff} = 0.2$.



Fig. 7. Results of the difference of prediction results (Diff. of Pred.) and the LOE for an illustrative example of the prediction with misspecification. The left panel is the histogram of the difference of prediction; The middle panel is the boxplot of the difference of prediction; The right panel is the boxplot of the LOE computed by different methods.

heavy-tailed distribution for the difference of kriging prediction results also appears in our simulation on different TLR tuning parameters.

Although one can improve the accuracy of MLOE using a resampling method, we do not apply this in our simulation because the resampling method is equivalent to increasing the number of replications in the original simulation, as we have discussed in Section 3. For the simple illustrative example stated above, one can run a resampling of 10,000 cases and report the mean of the LOEs computed by Stein's method, e.g., 1.8647×10^{-6} , as the final LOE result. However, computing LOE without resampling for each replicate and reporting the mean and standard deviation, or the boxplot in Fig. 7, is more informative.

5.3. Performance using different TLR accuracy levels

We investigate the effect of *tlr_acc* and *opt_tol* for the TLR approximations, where nb = 450 and *tlr_max_rank* is chosen from Table 2. To compare the effect of different values of *tlr_acc*, we fix *opt_tol* = 10^{-6} and choose *tlr_acc* = 10^{-5} , 10^{-7} , 10^{-9} , or 10^{-11} . We also compare the effect of different *opt_tol* values; to do so, we fix *tlr_acc* = 10^{-9} and choose *opt_tol* = 10^{-3} , 10^{-6} , 10^{-9} , or 10^{-12} .

Table 4

Estimation and prediction performances of the exact MLE and TLR approximation estimates for different tlr_acc values. Bias(·) means the estimate of the parameter minus its true value, and the estimation time means the computational time of the corresponding estimation. The value of MLOE for all cases ($\nu = 0.5$ and $\nu = 1.0$) is multiplied by 10⁶. The missing part in the table (–) means that the result is not available, because the covariance matrix is numerically non positive-definite.

| Mean (sd) | v = 0.5, l | $h_{\rm eff} = 0.2$ | | $v = 1.0, h_{\rm eff} = 1.6$ | | | | | | |
|--------------------------|---------------------|---------------------------------|---------------------|------------------------------|---------------------|---------------------|---------------------------------|---------------------|---------------------|---------------------|
| | MLE | TLR accuracy (<i>tlr_acc</i>) | | | MLE | TLR a | TLR accuracy (<i>tlr_acc</i>) | | | |
| | | 10 ⁻⁵ | 10 ⁻⁷ | 10 ⁻⁹ | 10 ⁻¹¹ | | 10 ⁻⁵ | 10 ⁻⁷ | 10 ⁻⁹ | 10 ⁻¹¹ |
| $Bias(\sigma^2)$ | -0.0080 (0.0908) | -0.0023 (0.1095) | -0.0079 (0.0908) | -0.0079 (0.0908) | -0.0079 (0.0908) | 0.0163 (0.6546) | - | 0.3800 (0.3154) | 0.2236 (0.7793) | 0.2276 (0.7927) |
| $Bias(\alpha)$ | -0.0006 (0.0063) | -0.0002 (0.0077) | -0.0006 (0.0063) | -0.0006 (0.0063) | -0.0006 (0.0063) | -0.0144 (0.1186) | - | 0.1033 (0.0686) | 0.0577 (0.1394) | 0.0581 (0.1408) |
| MLOE (×10 ⁶) | 3.3945 (5.9931) | 3.5691 (6.4517) | 3.3756 (5.9476) | 3.3756 (5.9474) | 3.3757 (5.9474) | 0.0273 (0.0669) | - | 0.0079 (0.0167) | 0.0110 (0.0243) | 0.0109 (0.0242) |
| MMOM | 0.0017 (0.0232) | 0.0009 (0.0233) | 0.0011 (0.0232) | 0.0011 (0.0232) | 0.0011 (0.0232) | 0.0014 (0.0227) | - | -0.1436 (0.0222) | -0.1428 (0.0222) | -0.1428 (0.0222) |
| RMOM | 0.0185 (0.0141) | 0.0186 (0.0140) | 0.0185 (0.0140) | 0.0185 (0.0140) | 0.0185 (0.0140) | 0.0182 (0.0135) | _ | 0.1436 (0.0222) | 0.1428 (0.0222) | 0.1428 (0.0222) |
| Estimation Time (s) | 168.1 (22.3) | 69.8 (11.9) | 77.8 (9.5) | 90.4 (13.5) | 112.1 (15.5) | 274.3 (74.6) | - | 62.7 (26.7) | 106.7 (31.2) | 111.6 (33.7) |

Table 5

Estimation and prediction performances of the exact MLE and TLR approximation estimates for different *opt_tol* values. Bias(·) means the estimate of the parameter minus its true value, and the estimation time means the computational time of the corresponding estimation. The value of MLOE for all cases ($\nu = 0.5$ and $\nu = 1.0$) is multiplied by 10⁶.

| Mean (sd) | v = 0.5, 1 | $h_{\rm eff} = 0.2$ $\nu = 1.0, h_{\rm eff} = 1.6$ | | | | | | | | |
|--------------------------|------------|--|-----------|------------------|-------------------|----------------------------------|------------------|-----------|------------------|------------|
| | MLE | Optimization tolerance (opt_tol) | | | MLE | Optimization tolerance (opt_tol) | | | | |
| | | 10 ⁻³ | 10^{-6} | 10 ⁻⁹ | 10 ⁻¹² | | 10 ⁻³ | 10^{-6} | 10 ⁻⁹ | 10^{-12} |
| $Bias(\sigma^2)$ | -0.0080 | 0.3654 | -0.0079 | -0.0079 | -0.0079 | 0.0163 | 0.3263 | 0.2236 | 0.2234 | 0.2234 |
| | (0.0908) | (0.3017) | (0.0908) | (0.0908) | (0.0908) | (0.6546) | (0.4421) | (0.7793) | (0.7787) | (0.7787) |
| $Bias(\alpha)$ | -0.0006 | 0.0253 | -0.0006 | -0.0006 | -0.0006 | -0.0144 | 0.0896 | 0.0577 | 0.0577 | 0.0577 |
| | (0.0063) | (0.0210) | (0.0063) | (0.0063) | (0.0063) | (0.1186) | (0.0904) | (0.1394) | (0.1393) | (0.1393) |
| MLOE ($\times 10^{6}$) | 3.3945 | 19.3107 | 3.3756 | 3.3756 | 3.3756 | 0.0273 | 0.0150 | 0.0110 | 0.0110 | 0.0110 |
| | (5.9931) | (13.9768) | (5.9474) | (5.9475) | (5.9475) | (0.0669) | (0.0708) | (0.0243) | (0.0243) | (0.0243) |
| ММОМ | 0.0017 | -0.0061 | 0.0011 | 0.0011 | 0.0011 | 0.0014 | -0.1432 | -0.1428 | -0.1428 | -0.1428 |
| | (0.0232) | (0.0234) | (0.0232) | (0.0232) | (0.0232) | (0.0227) | (0.0220) | (0.0222) | (0.0222) | (0.0222) |
| RMOM | 0.0185 | 0.0196 | 0.0185 | 0.0185 | 0.0185 | 0.0182 | 0.1432 | 0.1428 | 0.1428 | 0.1428 |
| | (0.0141) | (0.0140) | (0.0140) | (0.0140) | (0.0140) | (0.0135) | (0.022) | (0.0222) | (0.0222) | (0.0222) |
| Estimation | 168.1 | 33.8 | 90.4 | 102.2 | 113.0 | 274.3 | 29.4 | 106.7 | 124.1 | 136.2 |
| Time (s) | (22.3) | (13.6) | (13.5) | (13.2) | (13.0) | (74.6) | (21.7) | (31.2) | (34.8) | (36.0) |

The parameter settings and the simulation procedures are similar to those given in Section 5.2. When $tlr_acc = 10^{-11}$, the tlr_max_rank value in Table 2 is not large enough. Thus, we use increased values of tlr_max_rank in this case, namely, when v = 0.5, we set $tlr_max_rank = 270$ for $h_{\text{eff}} = 0.2$ and $tlr_max_rank = 260$ for other cases; when v = 1, we set $tlr_max_rank = 200$. We only provide the estimation and prediction performances for two typical cases, when v = 0.5, $h_{\text{eff}} = 0.2$, and when v = 1, $h_{\text{eff}} = 1.6$. Table 4 shows the results obtained with different values of tlr_acc , and Table 5 presents the results for different opt_tol values. For more details, please refer to the Supplementary Material.

Tables 4 and 5 indicate that the exact MLE and the TLR approximations can provide accurate prediction results since the small MLOE values suggest that the loss of prediction efficiency is very small. The MMOM results indicate that the computed MSEs are also accurate, except when

Table 6

Prediction performance and the computational time for TLR approximations with different combinations of *tlr_acc* and *opt_tol*. The estimation time means the computational time of the corresponding estimation. The value of MLOE for all cases ($\nu = 0.5$ and $\nu = 1.0$) is multiplied by 10⁶.

| Mean (sd) | (tlr_acc, opt | t_tol), $v = 0.5$ | $h_{\rm eff} = 0.2$ | | $(tlr_acc, opt_tol), v = 1.0, h_{eff} = 1.6$ | | | | |
|--------------------------|----------------------|---|---|---|--|---|---|---|--|
| | $(10^{-7}, 10^{-3})$ | (10 ⁻⁹ , 10 ⁻³) | (10 ⁻⁷ , 10 ⁻⁶) | (10 ⁻⁹ , 10 ⁻⁶) | $(10^{-7}, 10^{-3})$ | (10 ⁻⁹ , 10 ⁻³) | (10 ⁻⁷ , 10 ⁻⁶) | (10 ⁻⁹ , 10 ⁻⁶) | |
| MLOE (×10 ⁶) | 19.0060 | 19.3107 | 3.3756 | 3.3756 | 0.0138 | 0.0150 | 0.0079 | 0.0110 | |
| | (13.6484) | (13.9768) | (5.9476) | (5.9474) | (0.0688) | (0.0708) | (0.0167) | (0.0243) | |
| MMOM | -0.0062 | -0.0061 | 0.0011 | 0.0011 | -0.1436 | -0.1432 | -0.1436 | -0.1428 | |
| | (0.0232) | (0.0234) | (0.0232) | (0.0232) | (0.0219) | (0.0220) | (0.0222) | (0.0222) | |
| RMOM | 0.0193 | 0.0196 | 0.0185 | 0.0185 | 0.1436 | 0.1432 | 0.1436 | 0.1428 | |
| | (0.0142) | (0.0140) | (0.0140) | (0.0140) | (0.0219) | (0.0220) | (0.0222) | (0.0222) | |
| Estimation | 28.8 | 33.1 | 76.6 | 88.5 | 19.4 | 29.8 | 63.5 | 108.1 | |
| Time (s) | (11.4) | (13.3) | (9.4) | (13.2) | (8.8) | (22.1) | (27.1) | (31.3) | |

 $\nu = 1$ and $h_{\text{eff}} = 1.6$, which shows that the plug-in kriging based on TLR approximations may underestimate the prediction MSEs for a smoother random field with a larger effective range. The RMOM results exclude the canceling of error problem in the MMOM results. The plug-in kriging based on the exact MLE works well for all cases.

Table 4 shows that the TLR approximations give similar and relatively satisfactory performances of the estimation when $tlr_acc \le 10^{-9}$, and that the prediction performs well when $tlr_acc \le 10^{-7}$. The computational time increases when tlr_acc decreases, so we suggest $tlr_acc = 10^{-9}$ for maintaining estimation performance and $tlr_acc = 10^{-7}$ for maintaining prediction performance.

Table 5 shows that the estimation performs relatively well when $opt_tol \le 10^{-6}$. For prediction performances, the case of $opt_tol = 10^{-3}$ performs well enough, though the MLOE values are larger compared with other cases. So we suggest $opt_tol = 10^{-6}$ for keeping estimation performances and $opt_tol = 10^{-3}$ for keeping prediction performances because of the significantly faster computational speed in this case.

To further investigate the impact of different combinations of tlr_acc and opt_tol for prediction performances, we also try the cases where tlr_acc can be 10^{-7} , 10^{-9} and opt_tol can be 10^{-3} , 10^{-6} . Table 6 shows that choosing $tlr_acc = 10^{-7}$ and $opt_tol = 10^{-3}$ can provide a faster computation without losing too much prediction efficiency; we therefore suggest to select $tlr_acc = 10^{-7}$ and $opt_tol = 10^{-3}$ for keeping the prediction performances.

In conclusion, the TLR approximation method can significantly reduce the computational time and maintain the prediction efficiency. The only problematic aspect of the TLR method is that, when $\nu = 1$ and the effective range is large, the prediction MSE may be underestimated. For tuning the inputs in the TLR approximation, we recommend a moderate value of *nb* that makes the number of tiles divisible by the total number of CPUs, and a smallest feasible *tlr_max_rank*, which can be obtained by our simulations or by some simple trials. For instance, in Section 6, we choose *nb* according to the simulation in this section, such that the number of tiles remains the same for different *n*; *tlr_max_rank* is determined by trials similar to the simulation. We suggest *tlr_acc* = 10^{-9} , *opt_tol* = 10^{-6} for maintaining estimation performances; and suggest *tlr_acc* = 10^{-7} , *opt_tol* = 10^{-3} , when only the prediction performances are necessary to maintain.

Our suggested MLOE, MMOM, and RMOM criteria can successfully assess the loss of spatial prediction efficiency of the TLR method with different tuning parameters. We have also successfully detected the changes of the prediction efficiency for different *tlr_acc* and *opt_tol*. For different *nb*, the criteria's values are similar, indicating that the *nb* mainly affects the computational time, rather than the efficiency.

Remark 2. Table 4 shows that the TLR method with $tlr_acc = 10^{-5}$ can maintain the prediction performance for the exponential covariance model, but cannot for the Whittle covariance model, suggesting that one may need a lower tlr_acc value for a smoother process. Thus, if the process is



Fig. 8. Image plot of a subset of the soil moisture dataset residuals with N = 64,648 for the case study.

smoother than the Whittle covariance model and $tlr_{acc} = 10^{-7}$ is not applicable, then one can choose smaller tlr_{acc} such as 10^{-9} .

6. Application to soil moisture data

To show the effectiveness of our suggested TLR tuning parameter settings for real datasets, we compare the estimation and prediction performance of the TLR approximation to the exact MLE for the soil moisture dataset, with a 64-bit 20-core Intel Xeon Gold 6248 CPU running at 2.50 GHz, allowing the computation of the exact MLE by the ExaGeoStatR framework. We use four nodes, each node has 16 underlying CPUs, so the number of tiles is divisible by the total number of CPUs.

This dataset describes the daily soil moisture percentage at the top layer of the Mississippi basin, U.S., on January 1st, 2004, including the observation locations and the residuals of the fitted linear model in Huang and Sun (2018), and can be obtained from the website https://ecrc.github. io/exageostat/md_docs_examples.html, containing the example data of the ExaGeoStat package. The full dataset consists of about 2 million locations, however, we select a region of N = 64,648 locations that can be considered as representative regions for the whole area. For our computational experiment, we consider a subset of this dataset, where the latitude and longitude of the locations lie within [33.0, 35.2] × [-106.1, -103.9], as shown in Fig. 8. We use latitude and longitude as the coordinates of the observation locations in our computation.

In this numerical experiment, we randomly choose n = 3600, 14,400, 32,400, or 57,600 points for the estimation, and use the remaining points for assessing the prediction performance. For estimation, the smoothness parameter ν is either treated as unknown or fixed at $\nu = 0.5$. The searching intervals for optimizing the likelihood function are $\sigma^2 \in [0.01, 5]$, $\alpha \in [0.01, 5]$, and $\nu \in [0.01, 5]$ for the unknown case. In TLR approximations, we choose $tlr_acc = 10^{-9}$, $opt_tol = 10^{-6}$

Table 7

The value of *nb* and the corresponding *tlr_max_rank* used in the estimation of the soil moisture data.

| n | nb | tlr_max_rank (ν unknown) | tlr_max_rank (v known) |
|--------|------|-----------------------------|---------------------------|
| 3600 | 450 | 210 | 210 |
| 14,400 | 900 | 310 | 320 |
| 32,400 | 1350 | 490 | 500 |
| 57,600 | 1800 | 430 | 430 |

Table 8

Estimation results, computational time and MSPE of the MLE and TLR estimation for soil moisture data, where ν is unknown or fixed at 0.5.

| n | tlr_acc | opt_tol | ν is unknown | | | | | ν is fixed at 0.5 | | | | |
|--------|------------------|------------------|--------------|--------|--------|-----------|--------|-----------------------|--------|-----------|--------|--|
| | | | σ^2 | α | ν | Times (s) | MSPE | σ^2 | α | Times (s) | MSPE | |
| | Exact N | ILE | 1.2488 | 0.4590 | 0.2970 | 2038.1 | 0.2283 | 1.0970 | 0.1105 | 663.1 | 0.2335 | |
| 3600 | 10 ⁻⁷ | 10 ⁻³ | 1.1163 | 0.3800 | 0.2971 | 83.5 | 0.2283 | 1.1842 | 0.1191 | 41.7 | 0.2337 | |
| | 10^{-9} | 10^{-6} | 1.2486 | 0.4587 | 0.2971 | 399.7 | 0.2283 | 1.0969 | 0.1106 | 142.9 | 0.2335 | |
| | Exact N | ILE | 1.1412 | 0.2358 | 0.3566 | 13550.4 | 0.1461 | 1.0046 | 0.0864 | 8075.8 | 0.1457 | |
| 14,400 | 10 ⁻⁷ | 10 ⁻³ | 0.8784 | 0.1740 | 0.3488 | 1331.5 | 0.1462 | 1.2210 | 0.1054 | 454.2 | 0.1458 | |
| | 10^{-9} | 10^{-6} | 1.1410 | 0.2356 | 0.3568 | 3248.5 | 0.1461 | 1.0046 | 0.0865 | 1863.5 | 0.1457 | |
| | Exact N | ILE | 1.0478 | 0.1263 | 0.4282 | 83 197.3 | 0.1067 | 0.9800 | 0.0797 | 36 503.7 | 0.1060 | |
| 32,400 | 10 ⁻⁷ | 10 ⁻³ | 1.4342 | 0.1898 | 0.4229 | 8687.8 | 0.1068 | 1.1183 | 0.0913 | 3183.4 | 0.1060 | |
| | 10 ⁻⁹ | 10^{-6} | 1.0475 | 0.1261 | 0.4285 | 25 414.0 | 0.1067 | 0.9801 | 0.0798 | 14840.5 | 0.1060 | |
| | Exact M | ILE | 0.9870 | 0.0774 | 0.5066 | 286729.5 | 0.0811 | 0.9935 | 0.0805 | 184842.0 | 0.0812 | |
| 57,600 | 10^{-7} | 10^{-3} | 1.1911 | 0.1002 | 0.4928 | 15807.4 | 0.0813 | 1.1260 | 0.0916 | 5680.1 | 0.0812 | |
| | 10 ⁻⁹ | 10^{-6} | 0.9868 | 0.0773 | 0.5071 | 74225.3 | 0.0811 | 0.9937 | 0.0806 | 21907.0 | 0.0812 | |

and $tlr_acc = 10^{-7}$, $opt_tol = 10^{-3}$, which are our recommendations for keeping the estimation and prediction performances, respectively. The tlr_max_rank value is determined using the procedure presented in Section 5.2. The settings of *nb* and tlr_max_rank are shown in Table 7, and results are given in Table 8. We also try $tlr_acc = 10^{-5}$, $opt_tol = 10^{-3}$ or 10^{-6} in this experiment, but these parameter settings cannot work for n = 14,400, 32,400 and 57,600, because the covariance matrix is numerically non positive-definite, similar as in Table 4. Therefore, we ignore the results with $tlr_acc = 10^{-5}$.

Table 8 indicates that, when $tlr_{acc} = 10^{-9}$ and $opt_{tol} = 10^{-6}$, the TLR approximation can provide parameter estimates that are very close to the exact MLE, with a significantly shorter computational time. The computational times are further shortened when $tlr_{acc} = 10^{-7}$ and $opt_{tol} = 10^{-3}$. In this case, the prediction performances are similar to the exact MLE, though the estimates are no longer similar. Thus, our proposed tuning parameter suggestions work well for the soil moisture dataset, showing that our suggested MLOE, MMOM, and RMOM criteria are successfully used to choose the tuning parameters for the TLR approximation.

Besides the TLR method, we also compare the estimation and prediction performance for the soil moisture dataset with the composite likelihood method proposed by Vecchia (1988) and implemented by Guinness et al. (2020), and the Gaussian predictive process method proposed by Banerjee et al. (2008). We still consider n = 3600, 14,400, 32,400, 57,600 and use the same soil moisture dataset stated above.

The composite likelihood approximates the log-likelihood function $\ell(\theta)$ by

$$\tilde{\ell}_m(\boldsymbol{\theta}) := \sum_{i=1}^n \log p(Z(\boldsymbol{s}_i)|Z(\boldsymbol{s}_{i1}), \ldots, Z(\boldsymbol{s}_{im})),$$

where $\mathbf{s}_{i1}, \ldots, \mathbf{s}_{im}$ are min(i - 1, m) locations that are nearest to $Z(\mathbf{s}_i)$, p is the density of $Z(\mathbf{s}_i)$ conditional on the observations on these nearest locations. In this numerical experiment, we adopt the function vecchia_meanzero_loglik() in the GpGp package (Guinness et al., 2020) to compute $\tilde{\ell}_m$, which employs a Fisher scoring algorithm introduced by Guinness (2019). We use the function constrOptim() to compute the optimum. We set m = 20 and choose the initial value for the optimization by $(\sigma^2, \alpha, \nu) = (1, 0.2, 0.5)$. Denote by $\hat{\theta}_n = (\hat{\sigma}_n^2, \hat{\alpha}, \hat{\nu})$ the estimation results for different *n*, then we have $\hat{\theta}_{3600} = (1.1716, 0.3778, 0.3057)$; $\hat{\theta}_{14,400} = (1.0757, 0.2032, 0.3646)$; $\hat{\theta}_{32,400} = (1.0243, 0.1171, 0.4359)$; $\hat{\theta}_{57,600} = (0.9598, 0.0726, 0.5140)$. The MSPE results for n =3600, 14,400, 32,400, and 57,600 are 0.2283, 0.1459, 0.1066, and 0.0810, respectively, and the computational times (in seconds) are 123.1, 410.6, 1112.7, and 2805.1, respectively. One can check from Table 8 that, the TLR estimates for $tlr_acc = 10^{-9}$, $opt_tol = 10^{-6}$ are closer to the exact MLE results, compared to the composite likelihood estimates. Despite that the MSPE results of the TLR are slightly less competitive, it is clear that the TLR with our suggested tuning parameters for keeping the estimation performance reaches our goal for approximating the exact MLE estimation results, which serves our purpose better than the composite likelihood. Thus, with our suggested parameters, one can get more accurate information about the properties for the random field corresponding to the soil moisture dataset from a more accurate approximation of the MLE. Results for the case when ν is fixed at 0.5 are similar and not reported here.

Next, we compare the TLR with a typical low-rank based approximation method, the Gaussian predictive process (GPP) method proposed by Banerjee et al. (2008). For some predetermined knots s_1^*, \ldots, s_m^* , the GPP method approximates the observed value Z(s) by its kriging prediction value with respect to the observations on the knots plus a nugget term. Denote the observations on the knots by $Z^* := \{Z(s_1^*), \ldots, Z(s_m^*)\}^{\top}$. The kriging prediction, which is treated as an approximation of Z(s), is

$$\mathbb{E}\{Z(\boldsymbol{s})|Z(\boldsymbol{s}_1^{\star}),\ldots,Z(\boldsymbol{s}_m^{\star})\}=\boldsymbol{c}^{\top}(\boldsymbol{s},\boldsymbol{\theta})(C^{\star})^{-1}(\boldsymbol{\theta})\boldsymbol{Z}^{\star},$$

so the Gaussian predictive process model is

$$\tilde{Z}(\boldsymbol{s}) = \boldsymbol{c}^{\top}(\boldsymbol{s}, \boldsymbol{\theta})(C^{\star})^{-1}(\boldsymbol{\theta})\boldsymbol{Z}^{\star} + \boldsymbol{\epsilon}(\boldsymbol{s}),$$
(13)

where $c(s, \theta) = [C(s, s_{j}^{\star}; \theta)]_{j=1}^{m}, C^{\star}(\theta) = [C(s_{i}^{\star}, s_{j}^{\star}; \theta)]_{i,j=1}^{m}, C(s_{1}, s_{2}; \theta) = Cov\{Z(s_{1}), Z(s_{2})\}, and$ $\epsilon(s)$ is the nugget effect term which has a normal distribution with mean zero and variance τ^2 . The approximated kriging prediction is computed with the covariance matrix of the approximated random field $\tilde{Z}(\mathbf{s})$. For the Gaussian predictive process model, the computation of the inverse covariance matrix only involves the inversion of the matrix of order m, so computational time can be saved, compared with directly inverting a matrix of order n. In our comparison, we first fit the Gaussian predictive process model (13) by maximum likelihood estimation, where the covariance function $C(h; \theta)$ is the Matérn covariance (1). The smoothness parameter v is either treated as unknown, or fixed at $\nu = 0.5$. Next, we compute the plug-in kriging prediction and the MSPE based on the GPP model, on the same prediction locations used in the computation of Table 8. We choose the knots as the 23 \times 23 regular grid, which is evenly distributed on the observed range $[33.0, 35.2] \times [-106.1, -103.9]$. Results show that the MSPE of the GPP model is significantly larger than the corresponding prediction results based on exact MLE and has no apparent change when the number of observation n increases. For instance, when v is unknown and n = 3600. 14,400, 32,400, and 57,600, the MSPE of the GPP model are 0.4292, 0.4328, 0.4298, and 0.4268, respectively, whereas the corresponding MSPE for the exact model are 0.2283, 0.1461, 0.1067, and 0.0811, respectively. We also consider the case with a larger number of knots m for n = 3600. The knots used in this computation are chosen as the $\sqrt{m} \times \sqrt{m}$ regular grids, evenly distributed on the observed range $[33.0, 35.2] \times [-106.1, -103.9]$. Results show that even when m = 3600, the MSPE of the GPP method (0.2406 when ν is unknown or fixed) is still larger than that of the exact MLE method (0.2283 when ν is unknown or fixed). Thus, for kriging prediction of our considered soil dataset, the Gaussian predictive process method is less efficient than the exact MLE and the TLR approximations.

In conclusion, our suggested settings of the tuning parameters for TLR approximation, obtained by using the MLOE, MMOM, and RMOM criteria, can maintain the estimation or prediction performances for the soil moisture data. Thus, we have successfully applied our suggested criteria to the TLR tuning parameter selection problem in applications. According to our comparison, the TLR approximation with our suggested parameters outperforms the Gaussian predictive process method in the soil dataset prediction problem. For this soil data, the TLR approximation also outperforms the composite likelihood in approximating the exact MLE, though the MSPE of the TLR is here slightly less competitive.

7. Concluding remarks

In this article, we presented the Mean Loss of Efficiency (MLOE), Mean Misspecification of the MSE (MMOM), and Root mean square MOM (RMOM) criteria as tools to detect the difference of the prediction performance between the true and the approximated covariance models in simulation studies. We found that the suggested criteria are more appropriate than the commonly used Mean Square Prediction Error criterion, as the criteria can detect the efficiency loss when a smoother covariance model is misspecified as a rougher covariance model with a nugget effect in simulation studies, which the MSPE cannot do. Our suggested criteria are valuable tools for understanding the impact of the tuning parameters on the statistical performance of sophisticated approximation methods, which is crucial for selecting these inputs. To illustrate this, we compared the estimation and prediction performances of the Tile Low-Rank (TLR) approximation with different tuning parameters, and obtained a practical suggestion on how to choose these tuning parameters for different application requirements. We showed a case study in which our suggested tuning of parameters obtained by our criteria works well to keep the estimation or prediction performances of the TLR method, e.g., the TLR outperforms the typical Gaussian predictive process method in prediction efficiency, and outperforms the composite likelihood in estimation efficiency.

It is worth noting that the smoothness can affect the effectiveness of adopting the TLR approximation in spatial prediction and the proper value of tuning parameters in this approximation. For instance, if we can ensure that the process is not smoother than the Whittle covariance model, then we can set $tlr_acc = 10^{-7}$; else we may need a lower value, say $tlr_acc = 10^{-9}$, as is discussed in Remark 2. Thus, it would be appealing to introduce a suitable method for determining this kind of smoothness, such as determining the range of v in the Matérn covariance model, before estimation and prediction. However, as simulations have shown, the misspecification of the smoothness of a random field significantly worsens the spatial prediction performance. Thus the smoothness determination method should be accurate enough. One can use a rough estimate of the smoothness, e.g., the composite likelihood method, for determining the tuning parameters. However, some methods, such as the TLR, require a range for unknown parameters as the input, rather than an initial value. In this case, determining a range of the smoothness parameter is more favorable. In future work, we will develop suitable smoothness parameter determination methods, such as the hypothesis tests proposed by Hong et al. (2020) and the references therein, and apply the method in the parameter estimation process to further improve the computation performance.

Currently, the ExaGeoStatR framework for computing the TLR method is not available for estimating the unknown nugget effect. In future work, we will try to overcome this restriction. After that, we will investigate the tuning parameter selections for this case. The fitting result of the soil moisture dataset may be better when a nugget effect term is involved in the spatial model. It would also be interesting to compare the performance of the other tile-based approximation methods with the TLR method, using our suggested criteria, in order to determine the best method for different application cases.

Acknowledgments

The authors wish to thank the anonymous reviewers for their insightful comments and suggestions that substantially improved this paper. This work was supported by the King Abdullah University of Science and Technology (KAUST), Saudi Arabia and partially supported by the NSFC, China (Nos. 11771241 and 11931001).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j. spasta.2021.100517.

References

- Abdulah, S., Akbudak, K., Boukaram, W., Charara, A., Keyes, D., Ltaief, H., Mikhalev, A., Sukkari, D., Turkiyyah, G., 2019a. Hierarchical computations on manycore architectures (HiCMA). available at http://github.com/ecrc/hicma.
- Abdulah, S., Li, Y., Cao, J., Ltaief, H., Keyes, D.E., Genton, M.G., Sun, Y., 2019b. ExaGeoStatR: A package for large-scale geostatistics in R. arXiv preprint, arXiv:1908.06936.
- Abdulah, S., Ltaief, H., Sun, Y., Genton, M.G., Keyes, D.E., 2018. Parallel approximation of the maximum likelihood estimation for the prediction of large-scale geostatistics simulations. In: IEEE International Conference on Cluster Computing (CLUSTER). IEEE, pp. 98–108.
- Abdulah, S., Ltaief, H., Sun, Y., Genton, M.G., Keyes, D.E., 2018a. ExaGeoStat: A high performance unified software for geostatistics on manycore systems. IEEE Trans. Parallel Distrib. Syst. 29 (12), 2771–2784.
- Abdulah, S., Ltaief, H., Sun, Y., Genton, M.G., Keyes, D.E., 2019. Geostatistical modeling and prediction using mixed precision tile Cholesky factorization. In: IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC). IEEE, pp. 152–162.
- Agullo, E., Demmel, J., Dongarra, J., Hadri, B., Kurzak, J., Langou, J., Ltaief, H., Luszczek, P., Tomov, S., 2009. Numerical linear algebra on emerging architectures: The PLASMA and MAGMA projects. In: Journal of Physics: Conference Series, Vol. 180. IOP Publishing, 012037.
- Akbudak, K., Ltaief, H., Mikhalev, A., Keyes, D.E., 2017. Tile low rank Cholesky factorization for climate/weather modeling applications on manycore architectures. In: Kunkel, J., Yokota, R., Balaji, P., Keyes, D. (Eds.), High Performance Computing. ISC 2017. In: Lecture Notes in Computer Science, Springer, pp. 22–40.
- Aminfar, A., Ambikasaran, S., Darve, E., 2016. A fast block low-rank dense solver with applications to finite-element matrices. J. Comput. Phys. 304, 170–188.
- Augonnet, C., Thibault, S., Namyst, R., Wacrenier, P.A., 2011. StarPU: a unified platform for task scheduling on heterogeneous multicore architectures. Concurr. Comput.: Pract. Exper. 23 (2), 187–198.
- Banerjee, S., Gelfand, A.E., Finley, A.O., Sang, H., 2008. Gaussian predictive process models for large spatial data sets. J. R. Stat. Soc. Ser. B Stat. Methodol. 70 (4), 825–848.
- Blackford, L.S., Demmel, J., Dongarra, J., Duff, I., Hammarling, S., Henry, G., Heroux, M., Kaufman, L., Lumsdaine, A., Petitet, A., Pozo, R., Remington, K., Whaley, R.C., 2002. An updated set of basic linear algebra subprograms (BLAS). ACM Trans. Math. Software 28 (2), 135–151.
- Borm, S., Christophersen, S., 2016. Approximation of integral operators by Green quadrature and nested cross approximation. Numer. Math. 133 (3), 409-442.
- cha, 2017. The Chameleon project. available at https://project.inria.fr/chameleon.
- Cressie, N.A., Johannesson, G., 2006. Spatial prediction for massive datasets. In: Mastering the Data Explosion in the Earth and Environmental Sciences: Australian Academy of Science Elizabeth and Frederick White Conference. pp. 1–11.
- Cressie, N.A., Johannesson, G., 2008. Fixed rank kriging for very large spatial data sets. J. R. Stat. Soc. Ser. B Stat. Methodol. 70 (1), 209–226.
- Curriero, F.C., Lele, S.R., 1999. A composite likelihood approach to semivariogram estimation. J. Agric. Biol. Environ. Stat. 4 (1), 9–28.
- Dai, L., Wei, H., Wang, L., 2007. Spatial distribution and risk assessment of radionuclides in soils around a coal-fired power plant: A case study from the city of Baoji, China. Environ. Res. 104 (2), 201–208.
- Du, J., Zhang, H., Mandrekar, V.S., 2009. Fixed-domain asymptotic properties of tapered maximum likelihood estimators. Ann. Statist. 37 (6A), 3330–3361.
- Finley, A.O., Sang, H., Banerjee, S., Gelfand, A.E., 2009. Improving the performance of predictive process modeling for large datasets. Comput. Statist. Data Anal. 53 (8), 2873–2884.
- Furrer, R., Genton, M.G., Nychka, D., 2006. Covariance tapering for interpolation of large spatial datasets. J. Comput. Graph. Statist. 15 (3), 502–523.
- Gerber, F., De Jong, R., Schaepman, M.E., Schaepmanstrub, G., Furrer, R., 2018. Predicting missing values in spatio-temporal remote sensing data. IEEE Trans. Geosci. Remote Sens. 56 (5), 2841–2853.
- Ghysels, P., Li, X.S., Rouet, F.H., Williams, S., Napov, A., 2016. An efficient multicore implementation of a novel HSS-structured multifrontal solver using randomized sampling. SIAM J. Sci. Comput. 38 (5), S358–S384.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. J. Amer. Statist. Assoc. 102 (477), 359–378.
- Gramacy, R.B., Apley, D.W., 2015. Local Gaussian process approximation for large computer experiments. J. Comput. Graph. Statist. 24 (2), 561–578.
- Guinness, J., 2019. Gaussian process learning via Fisher scoring of Vecchia's approximation. arXiv preprint, arXiv:1905.08 374.

- Guinness, J., Katzfuss, M., Fahmy, Y., 2020. GpGp: Fast Gaussian process computation using Vecchia's approximation. R package version 0.3.1.
- Hackbusch, W., 1999. A sparse matrix arithmetic based on H-matrices. Part I: Introduction to H-matrices. Computing 62 (2), 89–108.
- Heaton, M.J., Datta, A., Finley, A.O., Furrer, R., Guinness, J., Guhaniyogi, R., Gerber, F., Gramacy, R.B., Hammerling, D., Katzfuss, M., et al., 2019. A case study competition among methods for analyzing large spatial data. J. Agric. Biol. Environ. Stat. 24 (3), 398–425.
- Hengl, T., Heuvelink, G.B.M., Stein, A., 2004. A generic framework for spatial prediction of soil variables based on regression-kriging. Geoderma 120 (1), 75–93.
- Hong, Y., Zhou, Z., Yang, Y., 2020. Hypothesis testing for the smoothness parameter of Matérn covariance model on a regular grid. J. Multivariate Anal. 177, 104597.
- Huang, H., Sun, Y., 2018. Hierarchical low rank approximation of likelihoods for large spatial datasets. J. Comput. Graph. Statist. 27 (1), 110–118.
- Kaufman, C.G., Schervish, M.J., Nychka, D.W., 2008. Covariance tapering for likelihood-based estimation in large spatial data sets. J. Amer. Statist. Assoc. 103 (484), 1545–1555.
- Lindgren, F., Rue, H., Lindstrom, J., 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. J. R. Stat. Soc. Ser. B Stat. Methodol. 73 (4), 423–498.
- Minsker, S., Srivastava, S., Lin, L., Dunson, D.B., 2014. Robust and scalable Bayes via a median of subset posterior measures. arXiv preprint arXiv:1403.2660.
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., Sain, S., 2015. A multiresolution Gaussian process model for the analysis of large spatial datasets. J. Comput. Graph. Statist. 24 (2), 579–599.
- Pichon, G., Darve, E., Faverge, M., Ramet, P., Roman, J., 2017. Sparse supernodal solver using block low-rank compression. In: 18th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC 2017). Orlando, United States.
- Stein, M.L., 1999. Interpolation of Spatial Data: Some Theory for Kriging. Springer-Verlag New York.
- Stein, M.L., 2014. Limitations on low rank approximations for covariance matrices of spatial data. Spat. Stat. 8, 1-19.
- Stein, M.L., Chi, Z., Welty, L.J., 2004. Approximating likelihoods for large spatial data sets. J. R. Stat. Soc. Ser. B Stat. Methodol. 66 (2), 275–296.
- Sun, Y., Li, B., Genton, M.G., 2012. Geostatistics for large datasets. In: Porcu, E., Montero, J.M., Schlather, M. (Eds.), Advances and Challenges in Space-Time Modelling of Natural Events, Vol. 207. Springer, Berlin, Heidelberg, pp. 55–77, Chapter 3.
- Sun, Y., Stein, M.L., 2016. Statistically and computationally efficient estimating equations for large spatial datasets. J. Comput. Graph. Statist. 25 (1), 187–208.
- Sushnikova, D., Oseledets, I., 2016. Preconditioners for hierarchical matrices based on their extended sparse form. Russian J. Numer. Anal. Math. Modelling 31 (1), 29–40.
- Vecchia, A.V., 1988. Estimation and model identification for continuous spatial processes. J. R. Stat. Soc. Ser. B Stat. Methodol. 50 (2), 297–312.
- Yan, Y., Genton, M.G., 2018. Gaussian likelihood inference on data from trans-Gaussian random fields with Matérn covariance function. Environmetrics 29 (5–6), e2458.