Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

iournal homepage: www.elsevier.com/locate/csda

Likelihood approximation with hierarchical matrices for large spatial datasets

Alexander Litvinenko^{a,*}, Ying Sun^b, Marc G. Genton^b, David E. Keyes^b

^a RWTH Aachen, 52072, Aachen, Germany

^b King Abdullah University of Science and Technology, 23955-6900, Thuwal, Saudi Arabia

ARTICLE INFO

Article history: Received 5 September 2017 Received in revised form 27 January 2019 Accepted 4 February 2019 Available online 12 February 2019

Keywords: Computational statistics Hierarchical matrix Large dataset Matérn covariance Random field Spatial statistics

ABSTRACT

The unknown parameters (variance, smoothness, and covariance length) of a spatial covariance function can be estimated by maximizing the joint Gaussian log-likelihood function. To overcome cubic complexity in the linear algebra, the discretized covariance function is approximated in the hierarchical (\mathcal{H}_{-}) matrix format. The \mathcal{H}_{-} matrix format has a loglinear computational cost and $O(kn \log n)$ storage, where the rank k is a small integer, and *n* is the number of locations. The \mathcal{H} -matrix technique can approximate general covariance matrices (also inhomogeneous) discretized on a fairly general mesh that is not necessarily axes-parallel, and neither the covariance matrix itself nor its inverse has to be sparse. It is investigated how the \mathcal{H} -matrix approximation error influences the estimated parameters. Numerical examples with Monte Carlo simulations, where the true values of the unknown parameters are given, and an application to soil moisture data with unknown parameters are presented. The C, C++ codes and data are freely available.

© 2019 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

The number of measurements that must be processed for statistical modeling in environmental applications is usually very large, and these measurements may be located irregularly across a given geographical region. This makes the computing procedure expensive and the data difficult to manage. These data are frequently modeled as a realization from a stationary Gaussian spatial random field. Specifically, we let $\mathbf{Z} = \{Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_n)\}^{\top}$, where $Z(\mathbf{s})$ is a Gaussian random field indexed by a spatial location $\mathbf{s} \in \mathbb{R}^d$, $d \ge 1$. Then, we assume that \mathbf{Z} has mean zero and a stationary parametric covariance function $C(\mathbf{h}; \boldsymbol{\theta}) = \operatorname{cov}\{Z(\mathbf{s}), Z(\mathbf{s} + \mathbf{h})\}$, where $\mathbf{h} \in \mathbb{R}^d$ is a spatial lag vector and $\boldsymbol{\theta} \in \mathbb{R}^q$ is the unknown parameter vector of interest. Statistical inferences about θ are often based on the Gaussian log-likelihood function:

$$\mathcal{L}(\boldsymbol{\theta}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log|\mathbf{C}(\boldsymbol{\theta})| - \frac{1}{2}\mathbf{Z}^{\mathsf{T}}\mathbf{C}(\boldsymbol{\theta})^{-1}\mathbf{Z},\tag{1}$$

where the covariance matrix $C(\theta)$ has entries $C(s_i - s_i; \theta)$, i, j = 1, ..., n. The maximum likelihood estimator of θ is the value $\hat{\theta}$ that maximizes (1). When the sample size n is large, the evaluation of (1) becomes challenging, due to the computation of the inverse and log-determinant of the *n*-by-*n* covariance matrix $C(\theta)$. Indeed, this requires $O(n^2)$ memory and $O(n^3)$ computational steps. Hence, scalable methods that can process larger sample sizes are needed. A recent proposal for exact

https://doi.org/10.1016/j.csda.2019.02.002







Correspondence to: RWTH Aachen, Chair of Mathematics for Uncertainty Quantification, Kackertstrasse 9C, 52072 Aachen, Germany.

E-mail addresses: litvinenko@uq.rwth-aachen.de (A. Litvinenko), ying.sun@kaust.edu.sa (Y. Sun), marc.genton@kaust.edu.sa (M.G. Genton), david.keyes@kaust.edu.sa (D.E. Keyes).

^{0167-9473/© 2019} The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/ by/4.0/).



Fig. 1. Scheme of approximations to the underlying dense covariance matrix (and corresponding methods). Hierarchically low-rank matrices generalize purely structure-based or global schemes with analysis-based locally adaptive schemes.

computations using high performance computing (HPC) on manycore systems using a unified software called *ExaGeoStat* can be found in Abdulah et al. (2018).

Stationary covariance functions, discretized on a rectangular grid, have block Toeplitz structure. This structure can be further extended to a block circulant form and resolved with the Fast Fourier Transform (FFT) (Whittle, 1954; Dahlhaus and Künsch, 1987; Guinness and Fuentes, 2017; Stroud et al., 2017; Dietrich and Newsam, 1997). The computing cost, in this case, is $O(n \log n)$. However, this approach either does not work for data measured at irregularly spaced locations or requires expensive, non-trivial modifications.

During the past two decades, a large amount of research has been devoted to tackling the aforementioned computational challenge of developing scalable methods: for example, low-rank tensor methods (Litvinenko et al., 2018; Nowak and Litvinenko, 2013), covariance tapering (Furrer et al., 2006; Kaufman et al., 2008; Sang and Huang, 2012), likelihood approximations in both the spatial (Stein et al., 2004, 2013) and spectral (Fuentes, 2007) domains, latent processes such as Gaussian predictive processes (Banerjee et al., 2008) and fixed-rank kriging (Cressie and Johannesson, 2008), and Gaussian Markov random-field approximations (Rue and Tjelmeland, 2002; Rue and Held, 2005; Fuglstad et al., 2015); see Sun et al. (2012) for a review. A generalization of the Vecchia approach and a general Vecchia framework was introduced in Katzfuss and Guinness (2017) and Vecchia (1988). Each of these methods has its strengths and drawbacks. For instance, covariance tapering sometimes performs even worse than assuming independent blocks in the covariance (Stein, 2013); low-rank approximations have their own limitations (Stein, 2014); and Markov models depend on the observation locations, requiring irregular locations to be realigned on a much finer grid with missing values (Sun and Stein, 2016). A matrix-free approach for solving the multi-parametric Gaussian maximum likelihood problem was developed in Anitescu et al. (2012). To further improve on these issues, other methods that have been recently developed include the nearest-neighbor Gaussian process models (Datta et al., 2016), low-rank update (Saibaba and Kitanidis, 2015), multiresolution Gaussian process models (Nychka et al., 2015), equivalent kriging (Kleiber and Nychka, 2015), multi-level restricted Gaussian maximum likelihood estimators (Castrillon et al., 2016), and hierarchical low-rank approximations (Huang and Sun, 2018). Bayesian approaches to identify unknown or uncertain parameters could be also applied (Rosić et al., 2013, 2012; Matthies et al., 2012; Litvinenko and Matthies, 2013; Matthies et al., 2012; Pajonk et al., 2012).

In this paper, we propose using the so-called hierarchical (\mathcal{H} -) matrices for approximating dense matrices with numerical complexity and storage $\mathcal{O}(k^{\alpha} n \log^{\alpha} n)$, where n is the number of measurements; $k \ll n$ is the rank of the hierarchical matrix, which defines the quality of the approximation; and $\alpha = 1$ or 2. \mathcal{H} -matrices are suitable for general grids and are also working well for large condition numbers. Previous results (Hackbusch, 2015) show that the \mathcal{H} -matrix technique is very stable when approximating the matrix itself (Li et al., 2014; Saibaba et al., 2012; Ballani and Kressner, 2015; Harbrecht et al., 2015; Ambikasaran et al., 2013; Börm and Garcke, 2007; Saibaba and Kitanidis, 2012), its inverse (Ambikasaran et al., 2016, 2013; Bebendorf and Hackbusch, 2003), its Cholesky decomposition (Bebendorf et al., 2016; Bebendorf, 2007), and the Schur complement (i.e., the conditional covariance matrix) (Hackbusch, 2015; Litvinenko, 2006, 2017a). The \mathcal{H} -matrix technique has also been proposed for inference including parameter estimation and maximum likelihood estimator (MLE) in Gaussian process regression in Ambikasaran et al. (2013) and Ambikasaran et al. (2016), for uncertainty quantification and conditional realizations in Saibaba and Kitanidis (2012), and for high-dimensional Gaussian probability computations in Genton et al. (2018) and Cao et al. (2019).

Other motivating factors for applying the H-matrix technique include the following:

- 1. *H*-matrices are more general than other compressed matrix representations (see scheme in Fig. 1);
- 2. The \mathcal{H} -matrix technique allows us to compute not only the matrix–vector products (e.g., like fast multipole methods), but also the more general classes of functions, such as $\mathbf{C}(\boldsymbol{\theta})^{-1}$, $\mathbf{C}(\boldsymbol{\theta})^{1/2}$, $|\mathbf{C}(\boldsymbol{\theta})|$, $\exp{\{\mathbf{C}(\boldsymbol{\theta})\}}$, resolvents, Cholesky decomposition and, many others (Hackbusch, 2015);



Fig. 2. (a) The \mathcal{H} -matrix approximation of an $n \times n$ covariance matrix from a discretized exponential covariance function on the unit square, with n = 16,641, unit variance, and length scale of 0.1. The dimensions of the densest (dark) blocks are 32×32 and the maximal rank is k = 13. (b) An example of the corresponding Cholesky factor with maximal rank k = 14. (c) The inverse of the exponential covariance matrix (precision matrix) with maximal rank k = 14.

- 3. The *H*-matrix technique is well-studied and has a solid theory, many examples, and multiple sequential and parallel implementations;
- 4. The \mathcal{H} -matrix accuracy is controllable by the rank, k, or by the accuracy ε . The full rank gives an exact representation;
- 5. The *H*-matrix technique preserves the structure of the matrix after the Cholesky decomposition and the inverse have been computed (see Fig. 2);
- 6. There are efficient rank-truncation techniques to keep the rank small after the matrix operations; for instance, the Schur complement and matrix inverse can be approximated again in the \mathcal{H} -matrix format.

Fig. 2(a) shows an \mathcal{H} -matrix approximation of the covariance matrix from a discretized (n = 16,641) exponential covariance function on the unit square, its Cholesky approximation 2(b), and its \mathcal{H} -matrix inverse 2(c). The dark (or red) blocks indicate the dense matrices and the gray (green) blocks indicate the rank-k matrices; the number inside each block is its rank. The steps inside the blocks show the decay of the singular values in log scale. The white blocks are empty. In the last few years, there has been great interest in numerical methods for representing and approximating large covariance matrices in the applied mathematics community (Quinonero-Candela et al., 2005; Börm and Garcke, 2007; Saibaba et al., 2012; Nowak and Litvinenko, 2013; Ambikasaran et al., 2013, 2014; Si et al., 2014; Ballani and Kressner, 2015).

Recently, the maximum likelihood estimator for parameter-fitting Gaussian observations with a Matérn covariance matrix was computed via a framework for unstructured observations in two spatial dimensions, which allowed the evaluation of the log-likelihood and its gradient with computational complexity $O(n^{3/2})$; the method relied on the recursive skeletonization factorization procedure (Ho and Ying, 2016; Martinsson and Rokhlin, 2005). However, the consequences of the approximation on the maximum likelihood estimator were not studied.

In Börm and Garcke (2007), the authors computed the exact solution of a Gaussian process regression by replacing the kernel matrix with a data-sparse approximation, which they called the \mathcal{H}^2 -matrix technique, cf. Li et al. (2014). It is more complicated than the \mathcal{H} -matrix technique, but has computational complexity and storage cost of $\mathcal{O}(kn)$.

The same \mathcal{H}^2 -matrix technique for solving large-scale stochastic linear inverse problems with applications in subsurface modeling was demonstrated in Ambikasaran et al. (2013). The authors explored the sparsity of the underlying measurement operator, demonstrated the effectiveness by solving a realistic crosswell tomography problem, quantified the uncertainty in the solution and provided an optimal capture geometry for this problem. Their algorithm was implemented in C++ and is available online.

The authors of Si et al. (2014) observed that the structure of shift-invariant kernels changed from low-rank to blockdiagonal (without any low-rank structure) when they varied the scale parameter. Based on this observation, they proposed a new kernel approximation algorithm, which they called the Memory-Efficient Kernel Approximation. That approximation considered both the low-rank and the clustering structure of the kernel matrix. They also showed that the resulting algorithm outperformed state-of-the-art low-rank kernel approximation methods regarding speed, approximation error, and memory usage. The BigQUIC method for a sparse inverse covariance estimation of a million variables was introduced in Hsieh et al. (2013). This method could solve ℓ_1 -regularized Gaussian Maximum Likelihood Estimation (MLE-) problems with dimensions of one-million. In Ballani and Kressner (2015), the authors estimated the covariance matrix of a set of normally distributed random vectors. To overcome large numerical issues in the high-dimensional regime, they computed the (dense) inverses of sparse covariance matrices using \mathcal{H} -matrices. This explicit representation enables them to ensure positive definiteness of each Newton-like iterate in the resulting convex optimization problem. In conclusion, the authors compared their new \mathcal{H} -QUIC method with the existing BIG-QUIC method (Hsieh et al., 2013). In Harbrecht et al. (2015), the authors offered to use \mathcal{H} -matrices for the approximation of random fields. In particular, they approximated Matérn covariance matrices in the \mathcal{H} -matrix format, suggested the pivoted \mathcal{H} -Cholesky method and provided an a posteriori error estimate in the trace norm. In Saibaba et al. (2012) and Saibaba and Kitanidis (2012), the authors applied \mathcal{H} -matrices to linear inverse problems in large-scale geostatistics. They started with a detailed explanation of the \mathcal{H} -matrix technique, then reduced the cost of dense matrix-vector products, combined \mathcal{H} -matrices with a matrix-free Krylov subspace and solved the system of equations that arise from the geostatistical approach. They illustrated the performance of their algorithm on an application, for monitoring CO₂ concentrations using crosswell seismic tomography. The largest problem size was $n = 10^6$. The code is available online.

In Ambikasaran et al. (2014), the authors considered the new matrix factorization for Hierarchical Off-Diagonal Low-Rank (HODLR) matrices. A HODLR matrix is an \mathcal{H} -matrix with the weak admissibility condition (see Appendix B). All off-diagonal sub-blocks of a HODLR matrix are low-rank matrices, and this fact significantly simplifies many algorithms. The authors showed that typical covariance functions can be hierarchically factored into a product of block low-rank updates of the identity matrix, yielding an $\mathcal{O}(n \log^2 n)$ algorithm for inversion, and used this product for evaluation of the determinant with the cost $\mathcal{O}(n \log n)$, with further direct calculation of probabilities in high dimensions. Additionally, they used this HODLR factorization to speed up prediction and to infer unknown hyper-parameters. The provided formulas for the factorization and the determinant hold only for HODLR matrices. How to extend this case to general \mathcal{H} -matrices, where off-diagonal block are allowed to be also \mathcal{H} -matrices, is not so clear. For instance, conditional covariance matrices may easily lose the HODLR structure. The largest demonstrated problem size was $n = 10^6$. The authors made their code freely available on GitHub.

Summarizing the literature review, we conclude that \mathcal{H} -matrices are already known to the statistical community. It is well-known that Matérn covariance functions could be approximated in the \mathcal{H} -matrix format (\mathcal{H} , \mathcal{H}^2 , HODLR). In Ballani and Kressner (2015), the authors used \mathcal{H} -matrices to compute derivatives of the likelihood function. Approximation of the inverse and Cholesky factorization is working in practice, but requires more theoretical analysis for error estimations (the existing theory is mostly developed for elliptic PDEs and integral equations). The influence of the \mathcal{H} -matrix approximation error on the quality of the estimated unknown parameters is not well studied yet. To research this influence with general \mathcal{H} -matrices (and not only with the simple HODLR format) is the main contribution of this work.

The rest of this paper is structured as follows. In Section 2, we review the \mathcal{H} -matrix technique and the \mathcal{H} -matrix approximation of Matérn covariance functions. Section 3 contains the hierarchical approximation of the Gaussian likelihood function, the algorithm for parameter estimation, and a study of the theoretical properties and complexity of the approximate log-likelihood (Guinness and Ipsen, 2015; Ipsen and Lee, 2005). Results from Monte Carlo simulations, where the true parameters are known, are reported in Section 4. An application of the hierarchical log-likelihood methodology to soil moisture data, where parameters are unknown, is considered in Section 5. We end the paper with a discussion in Section 6. The derivations of the theoretical results are provided in Appendix A, more details about \mathcal{H} -matrices in Appendix B.

2. Hierarchical approximation of covariance matrices

Hierarchical matrices have been described in detail (Hackbusch, 2015, 1999; Grasedyck and Hackbusch, 2003; Hackbusch et al., 2004; Litvinenko, 2006). Applications of the \mathcal{H} -matrix technique to the covariance matrices can be found in Ambikasaran et al. (2016), Saibaba and Kitanidis (2012), Saibaba et al. (2012), Börm and Garcke (2007), Harbrecht et al. (2015), Ambikasaran et al. (2013), Ambikasaran et al. (2014) and Khoromskij et al. (2009). There are many implementations. To our best knowledge, the HLIB library¹ is not supported anymore, but it could be used very well for academic purposes, the \mathcal{H}^2 -library² operates with \mathcal{H}^2 -matrices, is actively supported and contains some new ideas like parallel implementation on GPU. The HLIBPro library³ is actively supported commercial, robust, parallel, very tuned, well tested, but not open source library. It contains about 10 open source examples, which can be modified for the user's personal needs. There are some other implementations (e.g., from M. Bebendorf or Matlab implementations), but we do not have experience with them. To new users, we recommend to start with free open source HLIB or H2Lib libraries and then to move to HLIBPro, which is free for academic purposes.

2.1. Hierarchical matrices

In this section, we review the definition of \mathcal{H} -matrices and show how to approximate covariance matrices using the \mathcal{H} matrix format. The \mathcal{H} -matrix technique is defined as a hierarchical partitioning of a given matrix into sub-blocks followed by the further approximation of the majority of these sub-blocks by low-rank matrices (see Fig. 2). To define which sub-blocks can be approximated well by low-rank matrices and which cannot, a so-called admissibility condition is used (see more details in Appendix B). There are different admissibility conditions possible: weak, strong, domain decomposition based, and some others (Fig. 5). The admissibility condition may depend on parameters (e.g., the covariance length).

To define an \mathcal{H} -matrix some other notions are needed: index set *I*, clusters (denoted by τ and σ), cluster tree T_I , block cluster tree $T_{I\times I}$, admissibility condition (see Appendix B). We start with the index set $I = \{0, \ldots, n-1\}$, which corresponds to the available measurements in *n* locations. After the hierarchical decomposition of the index set *I* into sub-index sets has

¹ http://www.hlib.org/.

² https://github.com/H2Lib, developed by S. Börm and his group, Kiel, Germany.

³ https://www.hlibpro.com/, developed by R. Kriemann, Leipzig, Germany.

been completed (or in other words, a cluster tree T_I is constructed), the block cluster tree (denoted by $T_{I\times I}$, see more details in Appendix B) together with the admissibility condition decides which sub-blocks can be approximated by low-rank matrices. For definitions and example of cluster trees and corresponding block cluster trees (block partitioning) see Appendix B and Fig. 6.

On the first step, the matrix is divided into four sub-blocks. The hierarchical tree $T_{l \times l}$ tells how to divide. Then each (or some) sub-block(s) is (are) divided again and again until sub-blocks are sufficiently small. The resulting division is hierarchical. The procedure stops when either one of the sub-block sizes is n_{\min} or smaller ($n_{\min} \leq 128$), or when this sub-block can be approximated by a low-rank matrix.

Another important question is how to compute these low-rank approximations. The HLIB library uses a well-known method, called Adaptive Cross Approximation (ACA) algorithm (Goreinov et al., 1997; Bebendorf and Rjasanow, 2003; Börm et al., 2003), which performs the approximations with a linear complexity O(kn) in contrast to $O(n^3)$ by SVD.

Remark 2.1. Errors in the \mathcal{H} -matrix approximation may destroy the symmetry of the symmetric positive definite covariance matrix, causing the symmetric blocks to have different ranks. As a result, the standard algorithms used to compute the Cholesky decomposition may fail. A remedy to this is defining $\mathbf{C} := \frac{1}{2}(\mathbf{C} + \mathbf{C}^{\top})$.

Remark 2.2. Errors in the \mathcal{H} -matrix approximation may destroy the positive definiteness of the covariance matrix. Especially for matrices which have very small (very close to zero) eigenvalues. A remedy is: (1) to use a more robust, e.g., block \mathcal{H} -Cholesky, algorithm; (2) use **LDL**^{\top} factorization instead of **LL**^{\top} or (3) add a positive diagonal $\tau^2 \cdot \mathbf{I}$ to **C**.

Remark 2.3. We use both notations $\mathbf{C} \in \mathbb{R}^{n \times n}$ and $\mathbf{C} \in \mathbb{R}^{I \times I}$. In this work it is the same, since |I| = n. The notation $\mathbf{C} \in \mathbb{R}^{I \times I}$ is useful when it is important to differentiate between $\mathbf{C} \in \mathbb{R}^{I \times I}$ and $\mathbf{C} \in \mathbb{R}^{J \times J}$, where I and J are two different index sets of the same size.

To define the class of \mathcal{H} -matrices, we assume that the cluster tree T_l and block cluster tree $T_{l\times l}$ are already constructed.

Definition 2.1. We let *I* be an index set and $T_{I \times I}$ a hierarchical division of the index set product, $I \times I$, into sub-blocks. The set of \mathcal{H} -matrices with the maximal sub-block rank *k* is

 $\mathcal{H}(T_{I\times I}, k) := \{ \mathbf{C} \in \mathbb{R}^{I\times I} | \operatorname{rank}(\mathbf{C}|_b) \le k \text{ for all admissible blocks } b \text{ of } T_{I\times I} \},\$

where *k* is the maximum rank. Here, $\mathbf{C}|_{b} = (c_{ij})_{(i,j)\in b}$ denotes the matrix block of $\mathbf{C} = (c_{ij})_{(i,j)\in I}$ corresponding to the sub-block $b \in T_{I \times I}$.

Blocks that satisfy the admissibility condition can be approximated by low-rank matrices; see Hackbusch (1999). An \mathcal{H} -matrix approximation of **C** is denoted by $\widetilde{\mathbf{C}}$. The storage requirement of $\widetilde{\mathbf{C}}$ and the matrix–vector multiplication cost $\mathcal{O}(kn \log n)$, the matrix–matrix addition costs $\mathcal{O}(k^2n \log n)$, and the matrix–matrix product and the matrix inverse cost $\mathcal{O}(k^2n \log^2 n)$; see Hackbusch (1999).

Remark 2.4. There are two different \mathcal{H} -matrix approximation strategies, Börm et al. (2003) and Hackbusch (2015). (1) *The fixed rank strategy* (Fig. 2), i.e., each sub-block has a maximal rank k. It could be sufficient or not, but it has simple theoretical estimates for the computing time and storage. In this case we write $\widetilde{\mathbf{C}} \in \mathcal{H}(T_{I \times I}; k)$ (2) *The adaptive rank strategy* (Fig. 5, left sub-block), i.e., where absolute accuracy (in the spectral norm) for each sub-block **M** is ε or better (smaller):

$$k := \min\{k \in \mathbb{N}_0 \mid \exists \mathbf{M} \in \mathcal{R}(k, n, m) : \|\mathbf{M} - \mathbf{M}\| \le \varepsilon \|\mathbf{M}\|\},\$$

where $\mathcal{R}(\tilde{k}, n, m) := {\mathbf{M} \in \mathbb{R}^{n \times m} | \operatorname{rank}(\mathbf{M}) \le \tilde{k}}$ is a set of low-rank matrices of size $n \times m$ of the maximal rank \tilde{k} . The second strategy is better, since it allows to have an adaptive rank for each sub-block, but it makes it difficult to estimate the total computing cost and storage. In this case, we write $\tilde{\mathbf{C}} \in \mathcal{H}(T_{I \times I}; \varepsilon)$.

The fixed rank strategy is useful for a priori evaluations of the computational resources and storage memory. The adaptive rank strategy is preferable for practical approximations and is useful when the accuracy in each sub-block is crucial. Similarly, we introduce $\widetilde{\mathcal{L}}(\theta; k)$ and $\widetilde{\mathcal{L}}(\theta; \varepsilon)$ as \mathcal{H} -matrix approximations of the log-likelihood \mathcal{L} from (1). Sometimes we skip k (or ε) and write $\widetilde{\mathcal{L}}(\theta)$.

Thus, we conclude that different \mathcal{H} -matrix formats exist. For example, the weak admissible matrices (HODLR) result in a simpler block structure (Fig. 5, right), but the \mathcal{H} -matrix ranks could be larger than the ranks by standard admissible matrices. In Hackbusch et al. (2004), the authors observed a factor of \approx 3 between the weak and standard admissible sub-blocks. In Litvinenko (2006), the author observed that the HODLR matrix format might fail or result in very large ranks for 2D/3D computational domains or when computing the Schur complement. The \mathcal{H}^2 matrix format allows to get rid of the log factor in the complexity and storage but it is more complicated for understanding and implementation. We note that there are conversions from \mathcal{H} to \mathcal{H}^2 matrix formats implemented (see HLIB and H2Lib libraries).

Sequential and parallel complexity of the main linear algebra operations on <i>p</i> processors.								
Operation	Sequential complexity	Parallel complexity (Kriemann, 2005) (shared memory)						
Build $\widetilde{\mathbf{C}}$	$\mathcal{O}(n\log n)$	$\frac{\mathcal{O}(n\log n)}{p} + \mathcal{O}(V(T) \setminus L(T))$						
Store $\widetilde{\mathbf{C}}$	$\mathcal{O}(kn\log n)$	$\mathcal{O}(kn\log n)$						
$\widetilde{\textbf{C}}\cdot\textbf{z}$	$\mathcal{O}(kn\log n)$	$\frac{\mathcal{O}(kn\log n)}{p} + \frac{n}{\sqrt{p}}$						
$\widetilde{\mathbf{C}}^{-1}$	$\mathcal{O}(k^2 n \log^2 n)$	$rac{\mathcal{O}(n\log n)}{p} + \mathcal{O}(nn_{min}^2), 1 \leq n_{min} \leq 128$						
$\mathcal H$ -Cholesky \widetilde{L}	$\mathcal{O}(k^2n\log^2 n)$	$\frac{\mathcal{O}(n\log n)}{p} + \mathcal{O}(\frac{k^2 n\log^2 n}{n^{1/d}})$						
$ \widetilde{\mathbf{C}} $	$\mathcal{O}(k^2 n \log^2 n)$	$\frac{\mathcal{O}(n\log n)}{p} + \mathcal{O}(\frac{k^2 n \log^2 n}{n^{1/d}}), d = 1, 2, 3$						

2.2. Matérn covariance functions

Table 1

Among the many covariance models available, the Matérn family (Matérn, 1986) has gained widespread interest in recent years. The Matérn form of spatial correlations was introduced into statistics as a flexible parametric class, with one parameter determining the smoothness of the underlying spatial random field (Handcock and Stein, 1993). The varied history of this family of models can be found in Guttorp and Gneiting (2006).

The Matérn covariance depends only on the distance $h := \|\mathbf{s} - \mathbf{s}'\|$, where **s** and **s**' are any two spatial locations. The Matérn class of covariance functions is defined as

$$C(h;\boldsymbol{\theta}) = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)} \left(\frac{h}{\ell}\right)^{\nu} K_{\nu}\left(\frac{h}{\ell}\right),\tag{2}$$

with $\theta = (\ell, \nu, \sigma^2)^{\top}$; where σ^2 is the variance; $\nu > 0$ controls the smoothness of the random field, with larger values of ν corresponding to smoother fields; and $\ell > 0$ is a spatial range parameter that measures how quickly the correlation of the random field decays with distance, with larger ℓ corresponding to a faster decay (keeping v fixed) (Harbrecht et al., 2015). Here \mathcal{K}_{ν} denotes the modified Bessel function of the second kind of order ν . When $\nu = 1/2$, the Matérn covariance function reduces to the exponential covariance model and describes a rough field. The value $\nu = \infty$ corresponds to a Gaussian covariance model that describes a very smooth, infinitely differentiable field. Random fields with a Matérn covariance function are $|\nu - 1|$ times mean square differentiable.

2.3. Matérn covariance matrix approximation

By definition, covariance matrices are symmetric and positive semi-definite. The decay of the eigenvalues (as well as the \mathcal{H} -matrix approximation) depends on the type of the covariance matrix and its smoothness, covariance length, computational geometry, and dimensionality. In this section, we perform numerical experiments with \mathcal{H} -matrix approximations.

To underline the fact that \mathcal{H} -matrix approximations can be applied to irregular sets of locations, we use irregularly spaced locations in the unit square (only for Tables 2, 3):

$$\frac{1}{\sqrt{n}}\left(i-0.5+X_{ij}, \ j-0.5+Y_{ij}\right), \quad \text{for } i,j \in \{1,2,\dots,\sqrt{n}\},\tag{3}$$

where X_{ii} , Y_{ii} are i.i.d. uniform on (-0.4, 0.4) for a total of *n* observations; see Sun and Stein (2016). The observations are ordered lexicographically by the ordered pairs (i, j).

All of the numerical experiments herein are performed on a Dell workstation with 20 processors (40 cores) and total 128 GB RAM. The parallel \mathcal{H} -matrix library, HLIBPro (Kriemann, 2005), was used to build the Matérn covariance matrix, compute the Cholesky factorization, solve a linear system, calculate the determinant and the quadratic form. HLIBPro is fast and efficient; see the theoretical parallel scalability in Table 1. Here V(T) denotes the set of vertices, L(T) the set of leaves in the block-cluster tree $T = T_{1 \times 1}$, and n_{min} the size of a block when we stop further division into sub-blocks (see Section 2.1). Usually $n_{min} = \{32, 64, 128\}$, since a very deep hierarchy slows down computations.

Table 2 shows the \mathcal{H} -matrix approximation errors for log $|\widetilde{\mathbf{C}}|, \widetilde{\mathbf{C}}$ and the inverse $(\widetilde{\mathbf{L}}\widetilde{\mathbf{L}}^{\top})^{-1}$, where $\widetilde{\mathbf{L}}$ is the \mathcal{H} -Cholesky factor. The errors are computed in the Frobenius and spectral norms, where **C** is the exact Matérn covariance matrix with v = 0.5and $\sigma^2 = 1$. The local accuracy in each sub-block is ε . The number of locations is n = 16,641. The last column demonstrates the total compression ratio, c.r., which is equal to $1-\text{size}(\tilde{\mathbf{C}})/\text{size}(\mathbf{C})$. The exact values are log $|\mathbf{C}| = 2.63$ for $\ell = 0.0334$ and $\log |\mathbf{C}| = 6.36$ for $\ell = 0.2337$. The uniformly distributed mesh points are taken in the unit square and perturbed as in (3).

2.4. Convergence of the \mathcal{H} -matrix error vs. the rank k

In Table 3 we show the dependence of the Kullback–Leibler divergence (KLD) and two matrix errors on the \mathcal{H} -matrix rank k for the Matérn covariance function with parameters $\ell = \{0.25, 0.75\}$ and $\nu = 1.5$, computed on the domain $\mathcal{G} = [0, 1]^2$.

Table 2

The \mathcal{H} -matrix accuracy and compression rates (c.r.). Accuracy in each sub-block is ε ; n = 16,641, **C** is a Matérn covariance with $\nu = 0.5$ and $\sigma^2 = 1$. The spatial domain is the unit square with locations irregularly spaced as in (3).

ε	$\left \log \mathbf{C} - \log \widetilde{\mathbf{C}} \right $	$\Big \frac{\log \mathbf{C} - \log \widetilde{\mathbf{C}} }{\log \widetilde{\mathbf{C}} } \Big $	$\ \mathbf{C} - \widetilde{\mathbf{C}}\ _F$	$\frac{\ \mathbf{C} - \widetilde{\mathbf{C}}\ _2}{\ \widetilde{\mathbf{C}}\ _2}$	$\ \mathbf{I} - (\widetilde{\mathbf{L}}\widetilde{\mathbf{L}}^{\top})^{-1}\mathbf{C}\ _2$	c.r. (%)
$\ell = 0.0334$						
10 ⁻¹	$3.2\cdot10^{-4}$	$1.2\cdot10^{-4}$	$7.0 \cdot 10^{-3}$	$7.6 \cdot 10^{-3}$	2.9	91.8
10 ⁻²	$1.6 \cdot 10^{-6}$	$6.0 \cdot 10^{-7}$	$1.0 \cdot 10^{-3}$	$6.7 \cdot 10^{-4}$	$9.9 \cdot 10^{-2}$	91.6
10^{-4}	$1.8 \cdot 10^{-9}$	$7.0 \cdot 10^{-10}$	$1.0 \cdot 10^{-5}$	$7.3 \cdot 10^{-6}$	$2.0 \cdot 10^{-3}$	89.8
10 ⁻⁸	$4.7 \cdot 10^{-13}$	$1.8 \cdot 10^{-13}$	$1.3 \cdot 10^{-9}$	$6 \cdot 10^{-10}$	$2.1 \cdot 10^{-7}$	87.3
$\ell = 0.2337$						
10 ⁻⁴	$9.8 \cdot 10^{-5}$	$1.5 \cdot 10^{-5}$	$8.1 \cdot 10^{-5}$	$1.4 \cdot 10^{-5}$	$2.5 \cdot 10^{-1}$	91.5
10 ⁻⁸	$1.5 \cdot 10^{-9}$	$2.3\cdot10^{-10}$	$1.1 \cdot 10^{-8}$	$1.5 \cdot 10^{-9}$	$4 \cdot 10^{-5}$	88.7

Table 3

Convergence of the \mathcal{H} -matrix approximation error vs. the \mathcal{H} -matrix rank k of a Matérn covariance function with parameters $\ell = \{0.25, 0.75\}, \nu = 1.5,$ domain $\mathcal{G} = [0, 1]^2, n = 16,641.$

k	KLD		$\ \mathbf{C} - \widetilde{\mathbf{C}}\ _2$		$\ \mathbf{C}(\widetilde{\mathbf{C}})^{-1} - \mathbf{I}\ _2$	
	$\ell = 0.25$	$\ell = 0.75$	$\ell = 0.25$	$\ell = 0.75$	$\ell = 0.25$	$\ell=0.75$
20	0.12	2.7	5.3e-7	2e-7	4.5	72
30	3.2e-5	0.4	1.3e-9	5e-10	4.8e-3	20
40	6.5e-8	1e-2	1.5e-11	8e-12	7.4e-6	0.5
50	8.3e-10	3e-3	2.0e-13	1.5e-13	1.5e-7	0.1

The KLD was computed as follows:

$$D_{KL}(\mathbf{C},\widetilde{\mathbf{C}}) = \frac{1}{2} \left\{ \operatorname{trace}(\widetilde{\mathbf{C}}^{-1}\mathbf{C}) - n + \log\left(\frac{|\widetilde{\mathbf{C}}|}{|\mathbf{C}|}\right) \right\}.$$

We can bound the relative error $\|\mathbf{C}^{-1} - \widetilde{\mathbf{C}}^{-1}\| / \|\mathbf{C}^{-1}\|$ for the approximation of the inverse as

$$\frac{\|\mathbf{C}^{-1} - \widetilde{\mathbf{C}}^{-1}\|}{\|\mathbf{C}^{-1}\|} = \frac{\|(\mathbf{I} - \widetilde{\mathbf{C}}^{-1}\mathbf{C})\mathbf{C}^{-1}\|}{\|\mathbf{C}^{-1}\|} \le \|(\mathbf{I} - \widetilde{\mathbf{C}}^{-1}\mathbf{C})\|.$$

The spectral norm of $\|(\mathbf{I} - \widetilde{\mathbf{C}}^{-1}\mathbf{C})\|$ can be estimated by few steps of the power iteration method (HLIB library uses 10 steps). The rank $k \le 20$ is not sufficient to approximate the inverse, and the resulting error $\|\mathbf{C}(\widetilde{\mathbf{C}})^{-1} - \mathbf{I}\|_2$ is large. One remedy could be to increase the rank, but it may not help for ill-conditioned matrices. The spectral norms of $\widetilde{\mathbf{C}}$ are $\|\widetilde{\mathbf{C}}_{(\ell=0.25)}\|_2 = 720$ and $\|\widetilde{\mathbf{C}}_{(\ell=0.75)}\|_2 = 1068$.

Remark 2.5. Very often, the nugget $\tau^2 \mathbf{I}$ is a part of the model or is just added to the main diagonal of **C** to stabilize numerical calculations, i.e., $\widetilde{\mathbf{C}} := \widetilde{\mathbf{C}} + \tau^2 \mathbf{I}$, (Litvinenko et al., 2017). By adding a nugget, we "push" all the singular values away from zero. Adding a nugget effect to the list of unknown parameters $\boldsymbol{\theta} = (\ell, \nu, \sigma^2, \tau)^{\top}$ is straightforward. See the modified procedure in the GitHub repository (Litvinenko, 2017b).

3. Hierarchical approximation of Gaussian likelihood

3.1. Parameter estimation

We use the \mathcal{H} -matrix technique to approximate the Gaussian likelihood function. The \mathcal{H} -matrix approximation of the exact log-likelihood $\mathcal{L}(\theta)$ defined in (1) is denoted by $\widetilde{\mathcal{L}}(\theta; k)$:

$$\widetilde{\mathcal{L}}(\boldsymbol{\theta};k) = -\frac{n}{2}\log(2\pi) - \sum_{i=1}^{n}\log\{\widetilde{L}_{ii}(\boldsymbol{\theta};k)\} - \frac{1}{2}\mathbf{v}(\boldsymbol{\theta})^{\top}\mathbf{v}(\boldsymbol{\theta}),$$
(4)

where $\widetilde{\mathbf{L}}(\theta; k)$ is the rank-*k* \mathcal{H} -matrix approximation of the Cholesky factor $\mathbf{L}(\theta)$ in $\mathbf{C}(\theta) = \mathbf{L}(\theta)\mathbf{L}(\theta)^{\top}$, and $\mathbf{v}(\theta)$ is the solution of the linear system $\widetilde{\mathbf{L}}(\theta; k)\mathbf{v}(\theta) = \mathbf{Z}$. Analogously, we can define $\widetilde{\mathcal{L}}(\theta; \varepsilon)$.

To maximize $\widetilde{\mathcal{L}}(\theta; k)$ in (4), we use Brent's method (Brent, 1973; Press et al., 2007), also known as Brent–Dekker.⁴ The Brent–Dekker algorithm first uses the fast-converging secant method or inverse quadratic interpolation to maximize

⁴ Implemented in GNU Scientific library, https://www.gnu.org/software/gsl/.

 $\widetilde{\mathcal{L}}(\theta; \cdot)$. If those do not work, then it returns to the more robust bisection method. We note that the maximization of the log-likelihood function is an ill-posed problem, since even very small perturbations in the covariance matrix $\mathbf{C}(\theta)$ may result in huge perturbations in the log-determinant and the log-likelihood. An alternative to $\mathbf{C}(\theta) = \mathbf{L}(\theta)\mathbf{L}(\theta)^{\top}$ is the $\mathbf{L}(\theta)\mathbf{D}(\theta)\mathbf{L}(\theta)^{\top}$ decomposition, which is more stable since it avoids extracting square roots of diagonal elements, i.e., $\mathbf{L}\mathbf{D}\mathbf{L}^{\top} = (\mathbf{L}\mathbf{D}^{1/2})(\mathbf{L}\mathbf{D}^{1/2})^{\top}$. Very small negative diagonal elements can appear due to, e.g., the rounding off error.

3.2. Computational complexity and accuracy

We let $\mathbf{C}(\theta) \in \mathbb{R}^{n \times n}$ be approximated by an \mathcal{H} -matrix $\widetilde{\mathbf{C}}(\theta; k)$ with a maximal rank k. The \mathcal{H} -Cholesky decomposition of $\widetilde{\mathbf{C}}(\theta; k)$ costs $\mathcal{O}(k^2 n \log^2 n)$. The solution of the linear system $\widetilde{\mathbf{L}}(\theta; k)\mathbf{v}(\theta) = \mathbf{Z} \operatorname{costs} \mathcal{O}(k^2 n \log^2 n)$. The log-determinant $\log |\widetilde{\mathbf{C}}(\theta; k)| = 2 \sum_{i=1}^{n} \log \{\widetilde{L}_{ii}(\theta; k)\}$ is available for free. The cost of computing the log-likelihood function $\widetilde{\mathcal{L}}(\theta; k)$ is $\mathcal{O}(k^2 n \log^2 n)$ and the cost of computing the MLE in m iterations is $\mathcal{O}(mk^2 n \log^2 n)$.

Once we observe a realization \mathbf{z} of the random vector \mathbf{Z} , we can quantify the accuracy of the \mathcal{H} -matrix approximation of the log-likelihood function. Our main theoretical result is formulated below in Theorem 1.

Theorem 1 (Accuracy of the Hierarchical Log-likelihood). We let $\widetilde{C}(\theta)$ be an \mathcal{H} -matrix approximation of the matrix $C(\theta) \in \mathbb{R}^{n \times n}$, and $\mathbf{Z} = \mathbf{z}$ a data vector. We let also the spectral radius $\rho(\widetilde{C}(\theta)^{-1}C(\theta) - \mathbf{I}) < \varepsilon < 1$. Then the following statements hold:

$$\begin{split} \left| \log |\mathbf{C}(\boldsymbol{\theta}; k)| - \log |\mathbf{C}(\boldsymbol{\theta})| \right| &\leq -n \log(1 - \varepsilon) \approx n\varepsilon \quad \text{for small} \quad \varepsilon \\ |\widetilde{\mathcal{L}}(\boldsymbol{\theta}; k) - \mathcal{L}(\boldsymbol{\theta})| &\leq \frac{1}{2} \cdot n\varepsilon + \frac{1}{2} \|\mathbf{z}\|_{2}^{2} \cdot \|\widetilde{\mathbf{C}}^{-1}\|_{2} \cdot \varepsilon. \end{split}$$

Proof. See in Appendix A.

. ~

The estimate in Remark A.1 states how fast the probability decays while increasing the norm of **Z**. For simplicity, further on, we will assume $\|\mathbf{z}\|_2 \le c_0$.

In Ballani and Kressner (2015), the authors observed that the bound *n* in the inequalities above is pessimistic and is hardly observed in numerical simulations. Though Theorem 1 is shown for the fixed rank arithmetics, it also holds for the adaptive rank arithmetics with $\widetilde{C}(\theta; \varepsilon)$ and $\widetilde{\mathcal{L}}(\theta; \varepsilon)$.

4. Monte Carlo simulations

We performed numerical experiments with simulated data to recover the true values of the parameters of the Matérn covariance matrix, known to be $\theta^* := (\ell^*, \nu^*, \sigma^*) = (0.7, 0.9, 1.0)$. In the first step, we constructed 50 independent data sets (replicates) of size $n \in \{128, 64, \ldots, 4, 2\} \times 1000$, by multiplying the Cholesky factor $\widetilde{\mathbf{L}}(\theta, 10^{-10})$ on a Gaussian random vector $\mathbf{W} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \widetilde{\mathbf{C}}(\theta^*) = \widetilde{\mathbf{L}}(\theta^*)\widetilde{\mathbf{L}}(\theta^*)^{\top}$. We took the locations (not the data) from the daily moisture example, which is described below in Section 5. After that we ran the optimization algorithm and tried to identify (recover) the "unknown" parameters (ℓ^*, ν^*, σ^*). The boxplots for each parameter over 50 replicates are plotted in Fig. 3

This simulation study (Fig. 3) shows boxplots vs *n*. We are able to estimate the unknown parameters of a Matérn covariance function on an irregular grid with a certain accuracy. The parameter v (Fig. 3(b)) was identified more accurately than the parameters ℓ (Fig. 3(a)) and σ^2 (Fig. 3(c)). It is difficult to say why we do not see clear patterns for ℓ and σ^2 . We believe there are a few reasons. First, the iterative optimization method had difficulties to converge (we need to significantly decrease the threshold and increase the maximal number of iterations) for large *n*. This is very expensive regarding the computing time. Second, the log-likelihood is often very flat, and there are multiple solutions which fulfill the threshold requirements. We think that the reason is not in the \mathcal{H} -matrix accuracy since we took a very fine accuracy $\varepsilon = 10^{-10}$ already. Third, we have fixed the locations for a given size *n*. Finally, under infill asymptotics, the parameters ℓ , v, σ^2 cannot be estimated consistently but the quantity σ^2/ℓ^{2v} can, as seen in Fig. 3(d).

In Fig. 4 we present functional boxplots (Sun and Genton, 2011) of the estimated parameters as a function of the accuracy ε , based on 50 replicates with $n = \{4000, 8000, 32000\}$ observations. The true values to be identified are $\theta^* = (\ell, \nu, \sigma^2, \sigma^2/\ell^{2\nu}) = (0.7, 0.9, 1.0, 1.9)$, denoted by the dotted green line. Functional boxplots on (a), (b), (c) identify ℓ ; on (d), (e), (f) identify ν ; on (g), (h), (i) identify σ^2 , and on (j), (k), (l) identify $\sigma^2/\ell^{2\nu}$. One can see that we are able to identify all parameters with a good accuracy: the median, denoted by the solid black curve is very close to the dotted green line. The parameter ν and the auxiliary coefficient $\sigma^2/\ell^{2\nu}$ are identified better than ℓ and σ^2 , similarly to Fig. 3. The size of the middle box in the boxplot indicates the variability of the estimates. It decreases when n increases from 4000 to 32,000.

The functional boxplots in Fig. 4 demonstrate that the estimates are fairly insensitive to the choice of the accuracy ε . Plots of the estimated parameters' curves for 20 replicates with n = 64,000 are presented in Fig. 7 in Appendix C. All these plots suggest when to stop decreasing ε further. The MLE estimates are already good enough for relatively large ε . In other words, the \mathcal{H} -matrix approximations are accurate enough, and to improve the estimates of the parameters, one should improve the optimization procedure (the initial guess, the threshold, the maximal number of iterations).



Fig. 3. Boxplots for (a) ℓ ; (b) ν ; (c) σ^2 ; and (d) $\sigma^2/\ell^{2\nu}$ for $n \in \{128, 64, 32, 16, 8, 4, 2\} \times 1000$. Simulated data with known parameters (ℓ^*, ν^*, σ^*) = (0.7, 0.9, 1.0). Boxplots are obtained from 50 replicates; the \mathcal{H} -matrix sub-block accuracy is $\varepsilon = 10^{-7}$. The green, dotted horizontal line represents the true value of the parameters.

5. Application to soil moisture data

First, we introduce the daily soil moisture data set from a numerical model. Then we demonstrate how to apply \mathcal{H} -matrices to fit a Matérn covariance model. We use just one replicate sampled at *n* locations. Then we perform numerical tests to see which *n* is sufficient and explain how to chose the appropriate \mathcal{H} -matrix accuracy. We provide the corresponding computational times and storage costs.

5.1. Problem and data description

In climate and weather studies, numerical models play an important role in improving our knowledge of the characteristics of the climate system and the causes behind climate variations. These numerical models describe the evolution of many variables, such as temperature, wind speed, precipitation, humidity, and pressure, by solving a set of equations. The process is usually very complicated, involving physical parameterization, initial condition configurations, numerical integration, and data output schemes. In this section, we use the proposed \mathcal{H} -matrix methodology to investigate the spatial variability of soil moisture data, as generated by numerical models. Soil moisture is a key factor in evaluating the state of the hydrological process and has a broad range of applications in weather forecasting, crop yield prediction, and early warnings of flood and drought. It has been shown that a better characterization of soil moisture can significantly improve weather forecasting. However, numerical models often generate very large datasets due to the high spatial resolutions of the measured data, which makes the computation of the widely used Gaussian process models infeasible. Consequently, the whole region of interest must be divided into blocks of smaller size in order to fit the Gaussian process models independently to each block; alternatively, the size of the dataset must be reduced by averaging to a lower spatial resolution. Compared to fitting a consistent Gaussian process model to the entire region, it is unclear how much statistical efficiency is lost by such an



Fig. 4. Functional boxplots of the estimated parameters as a function of the accuracy ε , based on 50 replicates with $n = \{4000, 8000, 32000\}$ observations (left to right columns). True parameters $\theta^* = (\ell, \nu, \sigma^2, \sigma^2/\ell^{2\nu}) = (0.7, 0.9, 1.0, 1.9)$ represented by the green dotted lines. Replicates on (a), (b), (c) identify ℓ ; on (d), (e), (f) identify ν ; on (g), (h), (i) identify σ^2 , and on (j), (k), (l) identify $\sigma^2/\ell^{2\nu}$.

Table 4

õ ĨĨT Parameters n $\hat{\sigma}^2$ Time Size kB/n Time Size $\|\mathbf{I} - (\mathbf{\widetilde{L}}\mathbf{\widetilde{L}}^{\top})^{-1}\mathbf{C}\|_{2}$ î ŵ MB MB S s 2.000 0.04 5.2 2.7 0.07 5.2 $6.7\cdot10^{-4}$ 271 0.257 1.07 4.000 0.07 13 $1.5\cdot10^{-4}$ 3.03 0.228 127 32 0 17 1 0 2 $3.1 \cdot 10^{-4}$ 8,000 0.12 29.8 3.8 0.36 32 3.62 0.223 1.03 $2.6 \cdot 10^{-4}$ 16 000 0.52 72.0 46 75 1.73 0 2 5 2 1.11 0.95 32,000 0.93 193 6.2 3.04 210 $7.4\cdot10^{-5}$ 1.93 0.257 1.03 $1.0\cdot 10^{-4}$ 64,000 3.3 447 71 11.0 486 1.48 0.272 1.01 128,000 7.7 1160 9.5 36.7 1310 $3.8\cdot10^{-5}$ 0.84 0.302 0.94 256,000 13 2550 10.5 64.0 2960 $7.1 \cdot 10^{-5}$ 1.41 0.327 1.24 $7.1 \cdot 10^{-4}$ 512,000 23 4740 97 128 5800 1.41 0.331 1.25 13910 $3.0\cdot10^{-4}$ 0.308 1,000,000 53 11260 11 361 1.29 1.00 124 23650 1001 $5.2 \cdot 10^{-4}$ 0.340 2.000.000 12.4 29610 1.38 1.10

Estimation of unknown parameters for various *n* (see columns 1,8,9,10). Computing times and storage costs are given for 1 iteration. The \mathcal{H} -matrix accuracy $\varepsilon = 10^{-7}$.

approximation. Since our proposed \mathcal{H} -matrix technique can handle large covariance matrix computations, and the parallel implementation of the algorithm significantly reduces the computational time, we are able to fit Gaussian process models to a set of large subsamples in a given spatial region. Therefore, next we explore the effect of sample size on the statistical efficiency.

We consider high-resolution soil moisture data from January 1, 2014, measured in the topsoil layer of the Mississippi River basin, U.S.A. The spatial resolution is 0.0083 degrees, and the distance of one-degree difference in this region is approximately 87.5 km. The grid consists of $1830 \times 1329 = 2,432,070$ locations with 2,000,000 observations and 432,070 missing values. Therefore, the available spatial data are not on a regular grid. We use the same model for the mean process as in Huang and Sun (2018), and fit a zero-mean Gaussian process model with a Matérn covariance function to the residuals; see Huang and Sun (2018) for more details on the data description and exploratory data analysis.

5.2. Estimation, computing times and required storage costs

In the next experiment, we estimated the unknown parameters ℓ , ν , and σ^2 for different sample sizes n. Each time n samples were chosen from the whole set randomly. Table 4 shows the computational time and storage for the \mathcal{H} -matrix approximations for $n = \{2000, \ldots, 200000\}$. All computations are done with the parallel \mathcal{H} -matrix toolbox, HLIBpro. The number of computing cores is 40, the RAM memory 128 GB. It is important to note that the computing time (columns 2 and 5) and the storage cost (columns 3 and 6) are growing nearly linearly with n. These numerical computations illustrate the theoretical formulas from Table 1. Additionally, we provide the accuracy of the \mathcal{H} -Cholesky inverse and estimated parameters. The choice of the starting value is important. First, we run the optimization procedure with n = 2000 randomly sampled locations, and with the starting value $(\ell_0, \nu_0, \sigma_0^2) = (3, 0.5, 1)$, $\varepsilon = 10^{-7}$, the residual threshold 10^{-7} and the maximal number of iterations 1000. After 78 iterations the threshold was achieved and the solution is $(\hat{\ell}, \hat{\nu}, \hat{\sigma}^2) = (2.71, 0.257, 1.07)$. This is the starting value for the experiment with n = 4000. The estimated parameters, calculated with n = 4000 and 80 iterations, are $(\hat{\ell}, \hat{\nu}, \hat{\sigma}^2) = (3, 0.228, 1.02)$. Second, we take these values as the new starting values for another randomly chosen n = 8000 locations. After 60 iterations, the residual threshold is again achieved, the new estimated values are $(\hat{\ell}, \hat{\nu}, \hat{\sigma}^2) = (1.38, 0.34, 1.1)$.

In Table 5 we list the computing time and the storage cost (total and divided by *n*) vs. \mathcal{H} -matrix accuracy ε in each sub-block. This table connects the \mathcal{H} -matrix sub-block accuracy, the required computing resources, and the inversion error $\|\mathbf{I} - (\mathbf{\tilde{L}} \mathbf{L}^{\top})^{-1} \mathbf{C}\|_2$. It provides some intuition about how much computational resources are required for each ε for fixed *n*. We started with $\varepsilon = 10^{-7}$ (the first row) and then multiplied it each time by a factor of 4. The last row corresponds to $\varepsilon = 6.4 \cdot 10^{-3}$, which is sufficient to compute $\mathbf{\tilde{C}}$, but is not sufficient to compute the \mathcal{H} -Cholesky factor $\mathbf{\tilde{L}}$ (the procedure for computing \mathcal{H} -Cholesky crashes). We see that the storage and time increase only moderately with decreasing ε .

In Table 6 we use the whole daily moisture data set with $2 \cdot 10^6$ locations to estimate the unknown parameters. We provide computing times to set up the matrix \tilde{C} and to compute its \mathcal{H} -Cholesky factorization \tilde{L} . Additionally, we list the total storage requirement and the storage divided by n for both \tilde{C} and \tilde{L} . The initial point for the optimization algorithm is $(\ell_0, \nu_0, \sigma_0^2) = (1.18, 0.42, 1.3)$. We note that the choice of the initial point is very important and may have a strong influence on the final solution (due to the fact that the log-likelihood could be very flat around the optimum). We started our computations with a rough \mathcal{H} -matrix accuracy $\varepsilon = 10^{-4}$ and the \mathcal{H} -Cholesky procedure crashes. We tried $\varepsilon = 10^{-5}$ and received $(\hat{\ell}, \hat{\nu}, \hat{\sigma}) = (1.39, 0.325, 1.01)$. We took them as the initial condition for $\varepsilon = 10^{-5}$ and received very similar values, and so on. To find out which \mathcal{H} -matrix accuracy is sufficient, we provide tests for different values of $\varepsilon \in \{10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}\}$. One can see that $\varepsilon = 10^{-4}$ is not sufficient (the \mathcal{H} -Cholesky procedure crashes), and $\varepsilon = 10^{-5}$ provides similar results as $\varepsilon \in \{10^{-6}, 10^{-7}, 10^{-8}\}$. Therefore, our recommendation would be to start with some rather low \mathcal{H} -matrix accuracy, to use the obtained $\hat{\theta}$ as the new initial guess, and then decrease ε and repeat this a few times.

Table 5

ε	ĩ			$\widetilde{\mathbf{L}}\widetilde{\mathbf{L}}^{ op}$	$\widetilde{\mathbf{L}}\widetilde{\mathbf{L}}^{ op}$		
	Time	Size, MB	kB/n	Time	Size	$\ \mathbf{I} - (\widetilde{\mathbf{L}}\widetilde{\mathbf{L}}^{\top})^{-1}\mathbf{C}\ _2$	
	S			S	kB/dof		
$1.0 \cdot 10^{-7}$	5.8	1110	9.16	21.5	10.0	$8.4 \cdot 10^{-5}$	
$4.0 \cdot 10^{-7}$	5.4	1021	8.16	16.2	9.0	$3.5 \cdot 10^{-4}$	
$1.6 \cdot 10^{-6}$	4.5	913	7.3	13.3	8.1	$1.3 \cdot 10^{-3}$	
$6.4 \cdot 10^{-6}$	4.0	800	6.4	10.7	7.3	$4.4 \cdot 10^{-3}$	
$2.6 \cdot 10^{-5}$	3.5	698	5.6	8.7	6.3	$2.0 \cdot 10^{-2}$	
$1.0 \cdot 10^{-4}$	3.1	625	5.0	7.0	5.5	$6.4 \cdot 10^{-2}$	
$4.1 \cdot 10^{-4}$	2.7	550	4.4	6.2	5.1	$2.8 \cdot 10^{-1}$	
$1.6 \cdot 10^{-3}$	2.4	467	3.7	5.0	4.5	2.5	
$6.4 \cdot 10^{-3}$	2.3	403	3.2	-	-	-	

Computing time and storage cost vs. accuracy ε for parallel \mathcal{H} -matrix approximation; number of cores is 40, $\hat{\nu} = 0.302$, $\hat{\ell} = 0.84$, $\hat{\sigma}^2 = 0.94$, n = 128,000. \mathcal{H} -matrix accuracy in each sub-block for both $\widetilde{\mathbf{C}}$ and $\widetilde{\mathbf{L}}$ is ε .

Table 6

Estimating unknown parameters for various \mathcal{H} -matrix accuracies ε ; number of cores is 40, n = 2,000,000. \mathcal{H} -matrix accuracy in each sub-block for both $\widetilde{\mathbf{C}}$ and $\widetilde{\mathbf{L}}$ is ε , max. number of iterations is 60, threshold in the iterative solver is 10^{-4} .

ε	Estimat	Estimated parameters			Costs, \tilde{c}			Costs, $\widetilde{\mathbf{L}}$		
	ê	ŷ	$\hat{\sigma}^2$	Time Set up	Size GB	Size/n kB	Time s	Size GB	Size/n kB	$\ \mathbf{I} - (\mathbf{\widetilde{L}}\mathbf{\widetilde{L}}^{\top})^{-1}\mathbf{C}\ _2$
10^4	-	-	-	58	11	6	-	-	-	-
10^{-5}	1.39	0.325	1.01	61	12	6	241	15.4	8.11	$8.1 \cdot 10^{-2}$
10^{-6}	1.41	0.323	1.04	66	15	7.7	379	20.1	10.3	$5.7 \cdot 10^{-2}$
10^{-7}	1.39	0.330	1.05	114	22	11	601	27.4	14.0	$2.7 \cdot 10^{-3}$
10 ⁻⁸	1.38	0.329	1.06	138	29	15	1057	34.4	18.0	$1.5\cdot10^{-4}$

5.3. Reproducibility of the numerical results

To reproduce the presented numerical simulations, the first step is downloading the HLIB (www.hlib.org) or HLIBPro (www.hlibpro.com) libraries. Both are free for academic purposes. HLIB is a sequential, open-source C-code library, which is relatively easy to use and understand. HLIBPro is the highly tuned parallel library, where only the header and object files are accessible.

After signing the \mathcal{H} -matrix (HLIB or HLIBPro) license, downloading, and installing the libraries, our modules for approximating covariance matrices and computing log-likelihoods, and the soil moisture data can be downloaded from https://github.com/litvinen/HLIBCov.git.

HLIB requires not only the set of locations, but also the corresponding triangulation of the computing domains, see Litvinenko (2006) and Khoromskij and Litvinenko (2008). The vertices of the triangles should be used as the location points. To construct a triangulation, we recommend employing the routines in MATLAB, R, or any other third-party software. HLIBPro does not require triangulation; only the coordinates of the locations are needed (Litvinenko, 2017c).

6. Conclusion and discussion

We have applied a well-known tool from linear algebra, the \mathcal{H} -matrix technique, to spatial statistics. This technique makes it possible to work with large datasets of measurements observed on unstructured grids, in particular, to estimate the unknown parameters of a covariance model. The statistical model considered yields Matérn covariance matrices with three unknown parameters, ℓ , ν , and σ^2 . We applied the \mathcal{H} -matrix technique in approximating multivariate Gaussian log-likelihood functions and Matérn covariance matrices. The \mathcal{H} -matrix technique allowed us to drastically reduce the required memory and computing time. Within the \mathcal{H} -matrix approximation, we can increase the spatial resolution, take more measurements into account, and consider larger regions with larger datasets. \mathcal{H} -matrices require neither axes-parallel grids nor homogeneity of the covariance function.

From the \mathcal{H} -matrix approximation of the log-likelihood, we computed the \mathcal{H} -Cholesky factorization, the KLD, the logdeterminant, and a quadratic form, $\mathbf{Z}^{\mathsf{T}}\mathbf{C}^{-1}\mathbf{Z}$ (Tables 2, 3). We demonstrated the computing time, storage requirements, relative errors, and convergence of the \mathcal{H} -matrix technique (Tables 4, 5, 6).

We reduced the cost of computing the log-likelihood from cubic to log-linear, by employing the \mathcal{H} -matrix approximations, without significantly changing the log-likelihood. We considered both simulated examples, where we identified the known parameters, and a large, real data (soil moisture) example, where the parameters were unknown. We were able to calculate the maximum likelihood estimate $(\hat{\ell}, \hat{\nu}, \hat{\sigma}^2)$ for all examples. We researched the impact of the \mathcal{H} -matrix approximation error and the statistical error (see Fig. 7) to the total error. For both examples we repeated the calculations for 50 replicates and computed box plots. We analyzed the dependence of $(\hat{\ell}, \hat{\nu}, \hat{\sigma}^2)$ on ε and on the sample size *n*.

With the parallel \mathcal{H} -matrix library HLIBPro, we were able to compute the log-likelihood function for 2,000,000 locations in a few minutes (Fig. 4) on a desktop machine, which is \approx 5 years old and cost nowadays \approx 5.000 USD. At the same time, computation of the maximum likelihood estimates is much more expensive and depends on the number of iterations in the optimization and can take from few hours to few days. In total, the algorithm may need 100–200 iterations (or more, depending on the initial guess and the threshold). If each iteration takes 10 min, then we may need 24 h to get ($\hat{\ell}$, $\hat{\nu}$, $\hat{\sigma}^2$). Possible extension of this work are (1) to reduce the number of required iterations by implementing the first and second derivatives of the likelihood; (2) add nugget to the set of unknown parameters; (3) combine the current framework with the domain decomposition method.

Finally, we conjecture that the \mathcal{H} -matrix approximation may have some links with the Vecchia approximation (Katzfuss and Guinness, 2017; Vecchia, 1988) but this seems to be a difficult question. Indeed, the \mathcal{H} -matrix approximation is tuned with respect to either accuracy or rank in sub-blocks of the covariance matrix, whereas the Vecchia approximation is tuned based on the selection and the number of neighbors. Hence, without a clear understanding of the link between these two methods, it is not possible to tune them in a way that makes them comparable at this time. Further research is needed to investigate the possible relationship between these two methods.

Acknowledgments

The research reported in this publication was supported by funding from King Abdullah University of Science and Technology (KAUST). Additionally, we would like to express our special thanks of gratitude to Ronald Kriemann (for the HLIBPro software library) as well as to Lars Grasedyck and Steffen Börm for the HLIB software library. Alexander Litvinenko was supported by the Bayesian Computational Statistics & Modeling group (KAUST), the SRI-UQ group (KAUST), the Extreme Computing Research Center (KAUST), and the Alexander von Humboldt Foundation.

Appendix A. Error estimates

The Lemma 3.3. (page 5 in Ballani and Kressner (2015)) gives the following result. Let $\mathbf{C} \in \mathbb{R}^{n \times n}$, and $\mathbf{E} := \mathbf{C} - \widetilde{\mathbf{C}}$, $\widetilde{\mathbf{C}}^{-1}\mathbf{E} := \widetilde{\mathbf{C}}^{-1}\mathbf{C} - \mathbf{I}$, and for the spectral radius $\rho(\widetilde{\mathbf{C}}^{-1}\mathbf{E}) = \rho(\widetilde{\mathbf{C}}^{-1}\mathbf{C} - \mathbf{I}) < \varepsilon < 1$. Then

 $|\log |\mathbf{C}| - \log |\widetilde{\mathbf{C}}|| \le -n \log(1 - \varepsilon).$

To prove Theorem 1 we use the result above and

$$\begin{split} |\widetilde{\mathcal{L}}(\boldsymbol{\theta}; k) - \mathcal{L}(\boldsymbol{\theta})| &= \frac{1}{2} \log \frac{|\mathbf{C}|}{|\mathbf{\widetilde{C}}|} - \frac{1}{2} |\mathbf{z}^{\top} \left(\mathbf{C}^{-1} - \mathbf{\widetilde{C}}^{-1}\right) \mathbf{z}| \\ &\leq -\frac{1}{2} \cdot n \log(1 - \varepsilon) - \frac{1}{2} |\mathbf{z}^{\top} \left(\mathbf{C}^{-1}\mathbf{C} - \mathbf{\widetilde{C}}^{-1}\mathbf{C}\right) \mathbf{C}^{-1}\mathbf{z}| \\ &\leq -\frac{1}{2} \cdot n \log(1 - \varepsilon) - \frac{1}{2} |\mathbf{z}^{\top} \left(\mathbf{I} - \mathbf{\widetilde{C}}^{-1}\mathbf{C}\right) \mathbf{C}^{-1}\mathbf{z}| \\ &\leq \frac{1}{2} \cdot n\varepsilon + \frac{1}{2} ||\mathbf{z}||^2 \cdot c_1 \cdot \varepsilon, \end{split}$$

where $-\log(1-\varepsilon) \approx \varepsilon$ for small ε .

Remark A.1. Let $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. Theorem 2.2.7 in Gine and Nickl (2015) provides estimates for a norm of a Gaussian vector, $\|\mathbf{Z}\|$, for instance for $\sup_{t \in T} |Z(t)|$,

$$\Pr\left\{\left|\sup_{t\in T}|Z(t)|-M\right|>u\right\}\leq 2\left(1-\Phi\left(\frac{u}{\sigma}\right)\right)\leq e^{-\frac{u^2}{2\sigma^2}},\quad M>0.$$

Remark A.2. The assumption $\|\mathbf{C}^{-1}\| \le c_1$ is strong. This estimation depends on the matrix **C**, the smoothness parameter ν , the covariance length ℓ , and the \mathcal{H} -matrix rank k. First, we find an appropriate block-decomposition for the covariance matrix **C**. Second, we estimate the ranks of $\widetilde{\mathbf{C}}$. Third, we prove that the inverse/Cholesky can also be approximated in the \mathcal{H} -matrix format. Then we estimate the ranks for the inverse $\widetilde{\mathbf{C}}^{-1}$ and the Cholesky factor $\widetilde{\mathbf{L}}$. Finally, we estimate the \mathcal{H} -matrix approximation accuracies; see (Hackbusch, 2015). In the worst case, the rank k will be of order n. We also note that some covariance matrices are singular, so that $\widetilde{\mathbf{C}}^{-1}$ and $\widetilde{\mathbf{L}}$ may not exist. The computation of log $|\widetilde{\mathbf{C}}|$ could be an ill-posed problem, in the sense that small perturbations in **C** result in large changes in log $|\mathbf{C}|$.

Below we estimate the error caused by the usage of \mathcal{H} -matrices to generate simulated data. Let $\widetilde{\mathbf{C}} = \mathbf{C} + \varepsilon \mathbf{C}$ be an \mathcal{H} -matrix approximation, and $\mathbf{Z} = \mathbf{LW}$ be the data generated without \mathcal{H} -matrix approximation.

Lemma A.1. Let $\widetilde{\mathbf{Z}} = \widetilde{\mathbf{L}}\mathbf{W}$, and $\|\widetilde{\mathbf{L}} - \mathbf{L}\| \le \varepsilon_L$, where ε_L tends to zero when the \mathcal{H} -matrix ranks are growing. Let $\mathbf{Z} = \mathbf{z}$, $\widetilde{\mathbf{Z}} = \widetilde{\mathbf{z}}$, $\mathbf{W} = \mathbf{w}$ denote the data realizations. Then

$$\|\widetilde{\mathbf{z}} - \mathbf{z}\| \le \|\mathbf{L} - \mathbf{L}\| \|\mathbf{w}\| \le \varepsilon_L \|\mathbf{w}\|$$

\mathcal{H} -matrix block-wise accuracy is $\varepsilon = 10^{-5}$.								
n	$\ \widetilde{\mathbf{C}}\ _{F}$	$\ \widetilde{\mathbf{C}}\ _2$	$\ \widetilde{\mathbf{L}}\ _F$	$\ \widetilde{\mathbf{L}}\ _2$	$\ \widetilde{\mathbf{C}}^{-1}\ _{2}$			
$32 \cdot 10^3$	241	7.70	182	3.2	2.90			
$64 \cdot 10^3$	410	13.5	262	4.5	2.97			
$128 \cdot 10^{3}$	739	24.0	379	6.4	3.05			
$256 \cdot 10^{3}$	1388	46.0	552	9.2	3.08			

Norms of Matérn covariance matrix \tilde{C} and its Cholesky factor \tilde{L} , $\ell = 0.089$, $\nu = 0.22$, $\sigma^2 = 1$, relative \mathcal{H} -matrix block-wise accuracy is $\varepsilon = 10^{-5}$.

The error ε_L depends on the condition number of **C** and can be large, for instance, for ill-conditioned matrices. The next lemma estimates the error in the quadratic form by replacing the original data set **Z** = **z** by its \mathcal{H} -matrix approximation.

Lemma A.2. Let $\mathbf{L} = \widetilde{\mathbf{L}} + \varepsilon_L \mathbf{L}$ be an \mathcal{H} -matrix approximation of \mathbf{L} , then

$$|\widetilde{\mathbf{z}}^{\top}\mathbf{C}^{-1}\widetilde{\mathbf{z}} - \mathbf{z}^{\top}\mathbf{C}^{-1}\mathbf{z}| \leq 2\varepsilon |\widetilde{\mathbf{z}}^{\top}\mathbf{C}^{-1}\widetilde{\mathbf{z}}| + \mathcal{O}(\varepsilon^2)$$

Table 7

Proof. Let $\widetilde{Z} = \widetilde{L}W$, $W \sim \mathcal{N}(0, I)$, and $Z = LW = (\widetilde{L} + \varepsilon_L \widetilde{L})W = \widetilde{Z} + \varepsilon_L \widetilde{Z}$. Let Z = z, and $\widetilde{Z} = \widetilde{z}$ denote the data realizations. Then, after simple calculations, we obtain

$$|\widetilde{\mathbf{z}}^{\top}\mathbf{C}^{-1}\widetilde{\mathbf{z}} - (\widetilde{\mathbf{z}} + \varepsilon_{L}\widetilde{\mathbf{z}})^{\top}\mathbf{C}^{-1}(\widetilde{\mathbf{z}} + \varepsilon_{L}\widetilde{\mathbf{z}})| \leq 2\varepsilon_{L}|\widetilde{\mathbf{z}}^{\top}\mathbf{C}^{-1}\widetilde{\mathbf{z}}| + \varepsilon_{L}^{2}|\widetilde{\mathbf{z}}^{\top}\mathbf{C}^{-1}\widetilde{\mathbf{z}}| \to 0, \text{ if } \varepsilon_{L} \to 0. \quad \blacksquare$$

Lemma A.3. Let $\mathbf{z} \in \mathbb{R}^n$ and $\|\widetilde{\mathbf{C}}^{-1} - \mathbf{C}^{-1}\| \le \varepsilon_c$, then $|\mathbf{z}^\top \widetilde{\mathbf{C}}^{-1} \mathbf{z} - \mathbf{z}^\top \mathbf{C}^{-1} \mathbf{z}| \le \|\mathbf{z}\|^2 \varepsilon_c$.

Proof.

$$|\mathbf{z}^{\top}\widetilde{\mathbf{C}}^{-1}\mathbf{z} - \mathbf{z}^{\top}\mathbf{C}^{-1}\mathbf{z}| = |\mathbf{z}^{\top}(\widetilde{\mathbf{C}}^{-1} - \mathbf{C}^{-1})\mathbf{z}| \le \|\mathbf{z}\|^{2}\varepsilon_{c}. \quad \blacksquare$$

Lemma A.4. Let $\mathbf{z}, \widetilde{\mathbf{z}} \in \mathbb{R}^n$, and conditions of Lemmas A.2–A.3 hold, then

 $|\widetilde{\mathbf{z}}^{\top}\widetilde{\mathbf{C}}^{-1}\widetilde{\mathbf{z}} - \mathbf{z}^{\top}\mathbf{C}^{-1}\mathbf{z}| \leq \mathcal{O}(\varepsilon_L) + \mathcal{O}(\varepsilon_c).$

Proof. Adding and subtracting an additional term and applying the previous Lemmas A.2–A.3, we obtain

$$|\widetilde{\mathbf{z}}^{\top}\widetilde{\mathbf{C}}^{-1}\widetilde{\mathbf{z}} - \mathbf{z}^{\top}\mathbf{C}^{-1}\mathbf{z}| \leq |\widetilde{\mathbf{z}}^{\top}\widetilde{\mathbf{C}}^{-1}\widetilde{\mathbf{z}} - \mathbf{z}^{\top}\widetilde{\mathbf{C}}^{-1}\mathbf{z}| + |\mathbf{z}^{\top}\widetilde{\mathbf{C}}^{-1}\mathbf{z} - \mathbf{z}^{\top}\mathbf{C}^{-1}\mathbf{z}| \leq 2\varepsilon_{L}|\widetilde{\mathbf{z}}^{\top}\widetilde{\mathbf{C}}^{-1}\widetilde{\mathbf{z}}| + \|\mathbf{z}\|^{2}\varepsilon_{c}. \quad \blacksquare$$

Thus, if the approximations of the matrices **C**, C^{-1} and **L** are accurate (the error tends to zero relatively fast with increasing ranks), then the \mathcal{H} -matrix approximation error in the quadratic form also tends to zero with increasing k or decreasing ε_L . A more rigorous proof calculating the order of convergence is out of the scope of this paper. Dependence on the condition number is researched in Bebendorf et al. (2016) and Bebendorf (2007) (only for matrices that origin from elliptic partial differential equations and integral equations).

In Table 7 we list the Frobenius and spectral norms of \tilde{C} , \tilde{L} , and \tilde{C}^{-1} vs. *n*. This table shows that $\|\tilde{C}\|_2$ is growing with *n* and $\|\tilde{C}^{-1}\|_2$ almost not growing with *n*.

Appendix B. Admissibility condition and cluster and block-cluster trees

The admissibility condition (criteria) is used to divide a given matrix into sub-blocks and to define which sub-blocks can be approximated well by low-rank matrices and which not. There are many different admissibility criteria (see Fig. 5). The typical criteria are: the strong, weak and domain decomposition-based (Hackbusch, 2015). The user can also develop his own admissibility criteria, which may depend, for example, on the covariance length. It regulates, for instance, block partitioning, the size of the largest block, the depth of the hierarchical block partitioning. Large low-rank blocks are good, but they may require also larger ranks. Fig. 5 shows three examples of \mathcal{H} -matrices with three different admissibility criteria: (left) standard, (middle) domain-decomposition based and (right) weak (HODLR). Matrices are taken from different applications and illustrate only the diversity of block partitioning.

Fig. 6 shows examples of cluster trees (1–2) and block cluster trees (3–4). Let *I* be an index set of all locations. Denote for each index $i \in I$ corresponding to a basis function b_i (e.g., the "hat" function) the support $\mathcal{G}_i := \text{supp } b_i \subset \mathbb{R}^d$, where $d \in \{1, 2, 3\}$ is spatial dimension. Now we define two trees which are necessary for the definition of hierarchical matrices. These trees are labeled trees where the label of a vertex *t* is denoted by \hat{t} .

Definition B.1 (*Cluster Tree T₁ Hackbusch, 1999; Grasedyck and Hackbusch, 2003*). A finite tree *T₁* is a cluster tree over the index set *I* if the following conditions hold:

• *I* is the root of T_I and a subset $\hat{t} \subseteq I$ holds for all $t \in T_I$.



Fig. 5. Examples of H-matrices with three different admissibility criteria: (left) standard, (middle) domain-decomposition based and (right) weak (HODLR).



Fig. 6. (1-2) Two examples of cluster trees: (1st) standard, (2nd) domain-decomposition based; (3rd) an example of block cluster tree, (4th) a block cluster tree, constructed from domain decomposition based cluster tree shown in (2nd).

- If $t \in T_I$ is not a leaf, then the set of sons, sons(t), contains disjoint subsets of I and the subset \hat{t} is the disjoint union of its sons, $\hat{t} = \bigcup_{s \in SONS(t)} \hat{s}$.
- If $t \in T_I$ is a leaf, then $|\hat{t}| \le n_{min}$ for a fixed number n_{min} .

We generalize \mathcal{G}_i to clusters $\tau \in T_i$ by setting $\mathcal{G}_{\tau} := \bigcup_{i \in \tau} \mathcal{G}_i$, i.e., \mathcal{G}_{τ} is the minimal subset of \mathbb{R}^d that contains the supports of all basis functions b_i with $i \in \tau$.

Definition B.2 (*Block Cluster Tree* $T_{I \times I}$ *Hackbusch*, 1999; *Grasedyck and Hackbusch*, 2003). Let T_I be a cluster tree over the index set *I*. A finite tree $T_{I \times I}$ is a block cluster tree based on T_I if the following conditions hold:

- $root(T_{I \times I}) = I \times I$.
- Each vertex *b* of $T_{I \times I}$ has the form $b = (\tau, \sigma)$ with clusters $\tau, \sigma \in T_I$.
- For each vertex (τ, σ) with sons $(\tau, \sigma) \neq \emptyset$, we have

$$\operatorname{sons}(\tau,\sigma) = \begin{cases} (\tau,\sigma') : \sigma' \in \operatorname{sons}(\sigma), & \text{if } \operatorname{sons}(\tau) = \varnothing \wedge \operatorname{sons}(\sigma) \neq \varnothing \\ (\tau',\sigma) : \tau' \in \operatorname{sons}(\tau), & \text{if } \operatorname{sons}(\tau) \neq \varnothing \wedge \operatorname{sons}(\sigma) = \varnothing \\ (\tau',\sigma') : \tau' \in \operatorname{sons}(\tau), \sigma' \in \operatorname{sons}(\sigma), & \text{otherwise} \end{cases}$$

• The label of a vertex (τ, σ) is given by $\widehat{(\tau, \sigma)} = \widehat{\tau} \times \widehat{\sigma} \subseteq I \times I$.

We can see that $root(T_{I \times I}) = I \times I$. This implies that the set of leaves of $T_{I \times I}$ is a partition of $I \times I$.

Definition B.3. The standard admissibility condition (Adm_{η}) for two domains B_{τ} and B_{σ} (which actually correspond to two clusters τ and σ) is defined as follows

 $\min\{\operatorname{diam}(B_{\tau}), \operatorname{diam}(B_{\sigma})\} \leq \eta \operatorname{dist}(B_{\tau}, B_{\sigma}),$

where B_{τ} , $B_{\sigma} \subset \mathbb{R}^d$ are axis-parallel bounding boxes of the clusters τ and σ such that $\mathcal{G}_{\tau} \subset B_{\tau}$ and $\mathcal{G}_{\sigma} \subset B_{\sigma}$. diam and dist are usual diameter and distance, by default $\eta = 2.0$.

Appendix C. Accuracy stability plots

See Fig. 7.



Fig. 7. Estimated parameters as a function of the accuracy ε , based on 20 replicates (black solid curves) with n = 64,000 observations. True parameters $\theta = (\ell, \nu, \sigma^2) = (0.7, 0.9, 1.0)$ represented by the green dotted lines. Replicates on (a) identify ℓ ; on (b) identify ν ; and on (c) identify σ^2 .

References

Abdulah, S., Ltaief, H., Sun, Y., Genton, M.G., Keyes, D.E., 2018. Exageostat: A high performance unified software for geostatistics on manycore systems. IEEE Trans. Parallel Distrib. Syst. 29, 2771–2784.

Ambikasaran, S., Foreman-Mackey, D., Greengard, L., Hogg, D.W., ONeil, M., 2014. Fast direct methods for Gaussian processes and the analysis of NASA Kepler mission data. ArXiv preprint, https://arxiv.org/abs/14036015v1.

Ambikasaran, S., Foreman-Mackey, D., Greengard, L., Hogg, D.W., ONeil, M., 2016. Fast direct methods for Gaussian processes. IEEE Trans. Pattern Anal. Mach. Intell. 38 (2), 252–265. http://dx.doi.org/10.1109/TPAMI.2015.2448083, (ISSN 0162-8828).

Ambikasaran, S., Li, J.Y., Kitanidis, P.K., Darve, E., 2013. Large-scale stochastic linear inversion using hierarchical matrices. Comput. Geosci. 17 (6), 913–927.
Anitescu, M., Chen, J., Wang, L., 2012. A matrix-free approach for solving the parametric Gaussian process maximum likelihood problem. SIAM J. Sci. Comput. 34 (1), 240–262. (ISSN 1064-8275).

Ballani, J., Kressner, D., 2015. Sparse inverse covariance estimation with hierarchical matrices. http://sma.epfl.ch/anchpcommon/publications/quic_ballani_ kressner_2014.pdf.

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. http://dx.doi.org/10.1111/j.1467-9868.2008.00663.x, (ISSN 1467-9868).

Bebendorf, M., 2007. Why approximate LU decompositions of finite element discretizations of elliptic operators can be computed with almost linear complexity. SIAM J. Numer. Anal. 45, 1472–1494.

Bebendorf, M., Bollhöfer, M., Bratsch, M., 2016. On the spectral equivalence of hierarchical matrix preconditioners for elliptic problems. Math. Comp. 85 (302), 2839–2861.

Bebendorf, M., Hackbusch, W., 2003. Existence of H-matrix approximants to the inverse FE-matrix of elliptic operators with L[∞]-coefficients. Numer. Math. 95 (1), 1–28.

Bebendorf, M., Rjasanow, S., 2003. Adaptive low-rank approximation of collocation matrices. Computing 70 (1), 1–24, (ISSN 0010-485X).

Börm, S., Garcke, J., 2007. Approximating Gaussian processes with H²-matrices. In: Kok, J.N., Koronacki, J., de Mantaras, R.L., Matwin, S., Mladen, D., Skowron, A. (Eds.), Proceedings of 18th European Conference on Machine Learning, Warsaw, Poland. In: ECML 2007, vol. 4701, pp. 42–53.

Börm, S., Grasedyck, L., Hackbusch, W., 2003. Hierarchical Matrices. In: Lecture Note, vol. 21, Max-Planck Institute for Mathematics, Leipzig, www.mis.mpg. de.

Brent, R.P., 1973. Chapter 4: An Algorithm with Guaranteed Convergence for Finding a Zero of a Function, Algorithms for Minimization without Derivatives. Englewood Cliffs, NJ: Prentice-Hall.

Cao, J., Genton, M.G., Keyes, D.E., Turkiyyah, G., 2019. Hierarchical-block conditioning approximations for high-dimensional multivariate normal probabilities. Stat. Comput. (in press).

Castrillon, J.E., Genton, M.G., Yokota, R., 2016. Multi-level restricted maximum likelihood covariance estimation and kriging for large non-gridded spatial datasets. Spat. Stat. 18, 105–124.

Cressie, N., Johannesson, G., 2008. Fixed rank kriging for very large spatial data sets. J. R. Stat. Soc. Ser. B Stat. Methodol. 70 (1), 209–226. http://dx.doi.org/ 10.1111/j.1467-9868.2007.00633.x, (ISSN 1467-9868).

Dahlhaus, R., Künsch, H., 1987. Edge effects and efficient parameter estimation for stationary random fields. Biometrika 74 (4), 877–882. http://dx.doi.org/ 10.1093/biomet/74.4.877.

Datta, A., Banerjee, S., Finley, A.O., Gelfand, A.E., 2016. Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. J. Amer. Statist. Assoc. 111 (514), 800–812. http://dx.doi.org/10.1080/01621459.2015.1044091, URL https://doi.org/101080/0162145920151044091.

Dietrich, C., Newsam, G.N., 1997. Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. SIAM J. Sci. Comput. 18(4), 1088–107.

Fuentes, M., 2007. Approximate likelihood for large irregularly spaced spatial data. J. Amer. Statist. Assoc. 102, 321-331.

Fuglstad, G.-A., Simpson, D., Lindgren, F., Rue, H., 2015. Does non-stationary spatial data always require non-stationary random fields? Spat. Stat. 14, 505– 531.

Furrer, R., Genton, M.G., Nychka, D., 2006. Covariance tapering for interpolation of large spatial datasets. J. Comput. Graph. Statist. 15 (3), 502–523.

Genton, M.G., Keyes, D.E., Turkiyyah, G., 2018. Hierarchical decompositions for the computation of high-dimensional multivariate normal probabilities. J. Comput. Graph. Statist. 27 (2), 268–277.

Gine, E., Nickl, R., 2015. Gaussian processes. In: Mathematical Foundations of Infinite-Dimensional Statistical Models. In: Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, pp. 15–108. http://dx.doi.org/10.1017/CB09781107337862.003.

Goreinov, S.A., Tyrtyshnikov, E.E., Zamarashkin, N.L., 1997. A theory of pseudoskeleton approximations. Linear Algebra Appl. 261, 1–21, (ISSN 0024-3795). Grasedyck, L., Hackbusch, W., 2003. Construction and arithmetics of \mathcal{H} -matrices. Computing 70 (4), 295–334, (ISSN 0010-485X).

Guinness, J., Fuentes, M., 2017. Circulant embedding of approximate covariances for inference from Gaussian data on large lattices. J. Comput. Graph. Statist. 26 (1), 88–97.

Guinness, J., Ipsen, I., 2015. Efficient computation of Gaussian likelihoods for stationary Markov random field models. http://www4ncsu.edu/ipsen/ps/ MRFlikelihoods.pdf.

- Guttorp, P., Gneiting, T., 2006. Studies in the history of probability and statistics XLIX: On the Matérn correlation family. Biometrika 93, 989-995. http://dx.doi.org/10.1093/biomet/93.4.989.
- Hackbusch, W., 1999. A sparse matrix arithmetic based on H-matrices. I. Introduction to H-matrices. Computing 62 (2), 89–108, (ISSN 0010-485X).
- Hackbusch, W., 2015. Hierarchical matrices: Algorithms and Analysis. In: Springer Series in Comp. Math., vol. 49, Springer, ISBN: 978-3-662-47323-8.
- Hackbusch, W., Khoromskij, B.N., Kriemann, R., 2004. Hierarchical matrices based on a weak admissibility criterion. Computing 73 (3), 207-243, (ISSN 0010-485X).
- Handcock, M.S., Stein, M.L., 1993. A Bayesian analysis of kriging. Technometrics 35, 403-410.
- Harbrecht, H., Peters, M., Siebenmorgen, M., 2015. Efficient approximation of random fields for numerical applications. Numer. Linear Algebra Appl. 22 (4), 596-617
- Ho, K.L., Ying, L., 2016. Hierarchical interpolative factorization for elliptic operators: differential equations. Comm. Pure Appl. Math. 69, 1415–1451.
- Hsieh, C.-J., Sustik, M.A., Dhillon, I.S., Ravikumar, P.K., Poldrack, R., 2013. Big & amp; QUIC: Sparse inverse covariance estimation for a million variables. In: Burges, C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K. (Eds.), Advances in Neural Information Processing Systems 26. Curran Associates, Inc., pp. 3165-3173.
- Huang, H., Sun, Y., 2018. Hierarchical low rank approximation of likelihoods for large spatial datasets. J. Comput. Graph. Statist. 27 (1), 110–118. http: //dx.doi.org/10.1080/10618600.2017.1356324. URL https://doi.org/101080/10618600.2017.1356324.
- Ipsen, I.C.F., Lee, D.J., 2005. Determinant approximations. http://arxiv.org/abs/11050437v1.
- Katzfuss, M., Guinness, J., 2017. A general framework for Vecchia approximations of Gaussian processes. ArXiv e-prints.
- Kaufman, C.G., Schervish, M.J., Nychka, D.W., 2008. Covariance tapering for likelihood-based estimation in large spatial datasets. J. Amer. Statist. Assoc. 103 (484), 1545-1555.
- Khoromskij, B.N., Litvinenko, A., 2008. Domain decomposition based H-matrix preconditioners for the skin problem. In: Domain Decomposition Methods in Science and Engineering XVII. pp. 175-182.
- Khoromskij, B.N., Litvinenko, A., Matthies, H.G., 2009. Application of hierarchical matrices for computing the Karhunen-Loève expansion. Computing 84 (1-2), 49-67
- Kleiber, W., Nychka, D.W., 2015. Equivalent kriging. Spat. Stat. 12, 31-49.
- Kriemann, R., 2005. Parallel H-matrix arithmetics on shared memory systems. Computing 74 (3), 273–297. http://dx.doi.org/10.1007/s00607-004-0102-2, (ISSN 0010-485x)
- Li, J.Y., Ambikasaran, S., Darve, E.F., Kitanidis, P.K., 2014. A Kalman filter powered by H2 matrices for guasi-continuous data assimilation problems. Water Resour. Res. 50 (5), 3734–3749. http://dx.doi.org/10.1002/2013WR014607, URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/ 2013WR014607.
- Litvinenko, A., 2006. Application of hierarchical matrices for solving multiscale problems. https://www.researchgate.net/publication/315607074_ Application_of_Hierarchical_Matrices_for_Solving_Multi-Scale_Problems.
- Litvinenko, A., 2017a. Partial inversion of elliptic operator to speed up computation of likelihood in Bayesian inference, arXiv preprint. https://arxiv.org/ abs/1708 02207
- Litvinenko, A., Library: Log-likelihood approximation with hierarchical matrices, 2017b, URL https://github.com/litvinen/HLIBCov.git.
- Litvinenko, A., HLIBCov: Parallel Hierarchical Matrix Approximation of Large Covariance Matrices and Likelihoods with Applications in Parameter Identification, 2017c, ArXiv preprint, http://arxiv.org/abs/170908625, Sept. 2017.
- Litvinenko, A., Keyes, D., Khoromskaia, V., Khoromskij, B.N., Matthies, H.G., 2018. Tucker tensor analysis of matern functions in spatial statistics. Comput. Methods Appl. Math. http://dx.doi.org/10.1515/cmam-2018-0022.
- Litvinenko, A., Matthies, H.G., 2013. Inverse problems and uncertainty quantification. arXiv preprint, https://arxiv.org/abs/1312.5048.
- Litvinenko, A., Sun, Y., Genton, M.G., Keyes, D., Likelihood Approximation With Hierarchical Matrices For Large Spatial Datasets. ArXiv preprint, http: //arxiv.org/abs/1709.04419, Sept. 2017.
- Martinsson, P., Rokhlin, V., 2005. A fast direct solver for boundary integral equations in two dimensions. J. Comput. Phys. 205 (1), 1–23.

Matérn, B., 1986, Spatial Variation, second ed. In: Lecture Notes in Statistics, vol. 36, Springer-Verlag, Berlin: New York,

- Matthies, H., Litvinenko, A., Pajonk, O., Rosić, B.V., Zander, E., 2012. Parametric and uncertainty computations with tensor product representations. In: Dienstfrey, A.M., Boisvert, R.F. (Eds.), Uncertainty Quantification in Scientific Computing. In: IFIP Advances in Information and Communication Technology, vol. 377, Springer Berlin Heidelberg, ISBN: 978-3-642-32676-9, pp. 139-150.
- Nowak, W., Litvinenko, A., 2013. Kriging and spatial design accelerated by orders of magnitude: combining low-rank covariance approximations with FFT-techniques. Math. Geosci. 45 (4), 411-435.
- 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.
- Pajonk, O., Rosić, B.V., Litvinenko, A., Matthies, H.G., 2012. A deterministic filter for non-Gaussian Bayesian estimation applications to dynamical system estimation with noisy measurements. Physica D 241 (7), 775-788.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 2007. Section 9.3. Van Wijngaarden-Dekker-Brent Method. Numerical Recipes: The Art of Scientific Computing, Vol. 3rd ed., New York: Cambridge University Press, ISBN: 978-0-521-88068-8.
- Quinonero-Candela, J., Rasmussen, C.E., Herbrich, R., 2005. A unifying view of sparse approximate Gaussian process regression. J. Mach. Learn. Res. 6, 1939-1959.
- Rosić, B.V., Kučerová, A., Sýkora, J., Pajonk, O., Litvinenko, A., Matthies, H.G., 2013. Parameter identification in a probabilistic setting. Eng. Struct. 50, 179–196.
- Rosić, B.V., Litvinenko, A., Pajonk, O., Matthies, H.G., 2012. Sampling-free linear Bayesian update of polynomial chaos representations. J. Comput. Phys. 231 (17), 5761-5787. Rue, H., Held, L., 2005. Gaussian Markov Random Fields: Theory and Applications. CRC Press.
- Rue, H., Tjelmeland, H., 2002. Fitting Gaussian Markov random fields to Gaussian fields. Scand. J. Stat. 29 (1), 31-49.
- Saibaba, Å., Ambikasaran, S., Yue Li, J., Kitanidis, P.K., Darve, E., 2012. Application of hierarchical matrices to linear inverse problems in geostatistics. Oil & Gas Sci. Technol. Rev. IFP Ener. Nouvelles 67 (5), 857-875.
- Saibaba, A., Kitanidis, P., 2012. Efficient methods for large-scale linear inversion using a geostatistical approach. Water Resour. Res. 48 (5), http://dx.doi. org/10.1029/2011WR011778, URL https://agupubs.onlinelibrary.wiley.com/doi/abs/101029/2011WR011778.
- Saibaba, A., Kitanidis, P., 2015, Fast computation of uncertainty quantification measures in the geostatistical approach to solve inverse problems. Adv. Water Resour. 82, 124-138. http://dx.doi.org/10.1016/j.advwatres.2015.04.012, (ISSN 0309-1708).
- Sang, H., Huang, J.Z., 2012. A full scale approximation of covariance functions for large spatial data sets. J. R. Stat. Soc. Ser. B Stat. Methodol. 74 (1), 111–132. Si, S., Hsieh, C.-J., Dhillon, I.S., 2014. Memory efficient Kernel approximation. In: International Conference on Machine Learning (ICML).
- Stein, M.L., 2013. Statistical properties of covariance tapers. J. Comput. Graph. Statist. 22 (4), 866-885.
- Stein, M.L., 2014. Limitations on low rank approximations for covariance matrices of spatial data. Spat. Stat. 8, 1-19.
- Stein, M.L., Chen, J., Anitescu, M., 2013. Stochastic approximation of score functions for Gaussian processes. Ann. Appl. Stat. 7 (2), 1162–1191. http: //dx.doi.org/10.1214/13-AOAS627, URL https://doi.org/101214/13-AOAS627.
- 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. http: //dx.doi.org/10.1046/j.1369-7412.2003.05512.x, (ISSN 1467-9868).

Stroud, J.R., Stein, M.L., Lysen, S., 2017. Bayesian and maximum likelihood estimation for Gaussian processes on an incomplete lattice. J. Comput. Graph. Statist. 26 (1), 108–120.

Sun, Y., Genton, M.G., 2011. Functional boxplots. J. Comput. Graph. Statist. 20 (2), 316–334. http://dx.doi.org/10.1198/jcgs.2011.09224, URL https://doi.org/ 101198/jcgs.2011.09224.

Sun, Y., Li, B., Genton, M.G., 2012. Geostatistics for large datasets. In: Porcu, E., Montero, J.M., Schlather, M. (Eds.), Space-Time Processes and Challenges Related to Environmental Problems. Springer, pp. 55–77.

Sun, Y., Stein, M.L., 2016. Statistically and computationally efficient estimating equations for large spatial datasets. J. Comput. Graph. Statist. 25 (1), 187–208. Vecchia, A.V., 1988. Estimation and model identification for continuous spatial processes. J. R. Stat. Soc. Ser. B Stat. Methodol. 50 (2), 297–312, (ISSN 00359246).

Whittle, P., 1954. On stationary processes in the plane. Biometrika 41 (3-4), 434-449. http://dx.doi.org/10.1093/biomet/41.3-4.434.