

EFFICIENT ALGORITHMS FOR BAYESIAN INVERSE PROBLEMS WITH WHITTLE--MATERN PRIORS*

HARBIR ANTIL[†] AND ARVIND K. SAIBABA[†]

Abstract. This paper tackles efficient methods for Bayesian inverse problems with priors based on Whittle--Matérn Gaussian random fields. The Whittle--Matérn prior is characterized by a mean function and a covariance operator that is taken as a negative power of an elliptic differential operator. This approach is flexible in that it can incorporate a wide range of prior information including nonstationary effects, but it is currently computationally advantageous only for integer values of the exponent. In this paper, we derive an efficient method for handling all admissible noninteger values of the exponent. The method first discretizes the covariance operator using finite elements and quadrature, and uses preconditioned Krylov subspace solvers for shifted linear systems to efficiently apply the resulting covariance matrix to a vector. This approach can be used for generating samples from the distribution in two different ways: by solving a stochastic partial differential equation, and by using a truncated Karhunen--Loève expansion. We show how to incorporate this prior representation into the infinite-dimensional Bayesian formulation, and show how to efficiently compute the maximum a posteriori estimate, and approximate the posterior variance. Although the focus of this paper is on Bayesian inverse problems, the techniques developed here are applicable to solving systems with fractional Laplacians and Gaussian random fields. Numerical experiments demonstrate the performance and scalability of the solvers and their applicability to model and real-data inverse problems in tomography and a time-dependent heat equation.

Key words. Bayesian inverse problems, Gaussian priors, Krylov subspace methods, noninteger powers

M S C codes. 65M60, 65F50, 65F10, 65N21

D O I. 10.1137/22M1494397

1. Motivation and introduction. Inverse problems involve the use of experimental data or measurements to make inferences about parameters (e.g., initial or boundary conditions), governing the model. Inverse problems have a wide range of applications such as in material science, geophysics, and medicine. The Bayesian approach to inverse problems is prevalent because of its ability to incorporate prior knowledge of the parameters and its ability to quantify the uncertainty associated with the parameter reconstructions. However, many computational challenges persist in the implementation of the Bayesian approach to large-scale inverse problems.

Stuart [30] advocated the use of Gaussian priors within Bayesian inverse problems with covariance operators of the form $C = A^{-\alpha}$, where A is an elliptic differential operator, the exponent $\alpha > d/2$, and d is the spatial dimension. This choice ensures that the covariance operator is trace-class and under appropriate conditions on the likelihood, the resulting inverse problem is well-posed. A computational framework for

*Received by the editors May 10, 2022; accepted for publication (in revised form) January 9, 2023; published electronically July 19, 2023.

<https://doi.org/10.1137/22M1494397>

Funding: The first author is partially supported by National Science Foundation (NSF) grants DMS-2110263, DMS-1913004 and the Air Force Office of Scientific Research under Award FA9550-22-1-0248. The second author is partially supported by NSF DMS-1745644 and DMS-2026830.

[†]Department of Mathematical Sciences and the Center for Mathematics and Artificial Intelligence (CMAI), George Mason University, Fairfax, VA 22030 USA (hantil@gmu.edu).

[†]Department of Mathematics, North Carolina State University, Raleigh, NC 27695 USA (asaibab@ncsu.edu).

infinite-dimensional Bayes inverse problems based on Stuart's framework was developed in [11]. For computational convenience, the authors used the covariance operator $C = A^{-2}$ since it satisfies the requirement $a > d/2$ for up to three spatial dimensions. This choice has two advantages: first, it is integer valued and, second, it is easy to obtain a factored form of the covariance operator. For these reasons, this approach is attractive from a computational standpoint, and has since become popular and has resulted in scalable software implementations [31, 21]. A similar approach can also be found in [25], in which the authors also used the Whittle--Matérn priors for Bayesian inverse problems. However, the choice of $a = 2$ for computational convenience makes it restrictive in practice, and in many instances, it can lead to oversmoothing of reconstructions. It is, therefore, desirable to have an efficient computational method for handling exponents $a > d/2$ that allows the user to choose the appropriate prior covariance based on prior knowledge or estimate it from data in a hierarchical Bayesian setting.

The use of Whittle--Matérn priors in Bayesian inverse problems has close connections with the developments in Gaussian random fields in spatial and computational statistics. In [24], the authors considered the stochastic partial differential equation (SPDE) approach to Gaussian random fields. They showed an explicit connection between Gaussian and Gauss--Markov random fields based on the Matérn covariance family and provide explicit expressions for the precision operator for integer values of a . Based on this connection, they provided extensions beyond the Matérn model and showed how the SPDE approach can be applied on manifolds, for nonstationary random fields. In a review paper [23], 10 years after the publication of [24], the authors trace the recent developments in the SPDE approach to Gaussian and non-Gaussian random fields. Recent work (e.g., [7, 8]) allows for noninteger values of the exponent; see the discussion in section 2.4.

Goals and main contributions. In this paper, we want to develop an efficient method for representing and computing with Whittle--Matérn Gaussian random priors in which the covariance operator is of the form $C = A^{-a}$ and $a > d/2$. Furthermore, we also want to incorporate this efficient representation of covariance operators while solving Bayesian inverse problems. The novel and noteworthy features of our contributions are the following:

1. We consider a quadrature based approach for the approximation of the covariance operator for fractional powers. The resulting approximation has the form of a rational function. The action of the covariance operator on a vector can be written as a sequence of shifted linear systems, which we accelerate using Krylov subspace methods. As a result, we have a method for efficiently applying the covariance operator for any exponent $a > d/2$. The method is scalable to large problems and the cost is nearly the same for values of a between consecutive integers. This approach is applicable for stationary (including isotropic and anisotropic) and nonstationary random fields.
2. We leverage the fast method for covariance matrices to efficiently generate samples in two different ways: by solving the underlying SPDE that defines the prior covariance and by efficiently computing the truncated Karhunen--Loève approximation.
3. We show how to use this efficient representation of the prior covariance operator in an infinite-dimensional Bayesian problem formulation. We perform careful discretization of all the relevant quantities and leverage the fast applications of the covariance operator to compute the maximum a posteriori (MAP) estimate and approximate posterior variance.

4. Although the focus of this paper is on Bayesian inverse problems, the computational techniques developed here are applicable to solving systems with fractional Laplacians [9, 32] and Gaussian random fields.

Through numerical experiments, we demonstrate the computational performance of the solvers including an m40X speedup over native approaches. We also demonstrate the feasibility of the approach on model and real-data problems from seismic and X-ray tomography and an application to a 2-dimensional (2D) PDE-based inverse diffusion equation.

2. Background. In section 2.1, we review the SPDE approach to Gaussian random fields which form the basis for the prior distributions in Bayesian inverse problems. In section 2.2, we review the appropriate definitions of fractional elliptic operators and show how to discretize them in section 2.3. Finally, section 2.4 focuses on the related work.

2.1. SPDE approach to Whittle--Matérn Gaussian random fields. We follow the Whittle--Matérn random field approach, which models the random fields as stationary Gaussian random fields with the covariance kernel

$$c(x, y) := a^2 \frac{2^{1-u}}{a(u)} (a \|x - y\|_2)^{-u} K_u(a \|x - y\|_2), \quad x, y \in \mathbb{R}^d,$$

where $u > 0$ is the smoothness parameter, the integer value of which controls the mean-square differentiability of the process, K_u is the modified Bessel function of the second kind of order u , and a is a scaling parameter that governs the length scale. When $u = 1/2$, this corresponds to the exponential covariance function, and as $u \rightarrow \infty$, this corresponds to the Gaussian kernel. On $\mathbb{R}^d$, samples from the random field with the covariance function correspond to the solution of the SPDE

$$(2.1) \quad (a^2 - \Delta)^{(u+d/2)/2} u(x) = W(x), \quad x \in \mathbb{R}^d,$$

where W is a spatial Gaussian white noise process with unit variance on $\mathbb{R}^d$, and the marginal variance is

$$a^2 = \frac{a(u)}{(4\pi)^{d/2} a^{(u+d/2)}}.$$

However, to implement this on a bounded region $\Omega \subset \mathbb{R}^d$, we need to impose additional boundary conditions. The choice of boundary conditions may affect the correlation near the boundaries. To mitigate this effect, one may solve the SPDE on a slightly larger domain that encloses Ω . Alternative approaches include the use of Robin boundary conditions with optimized Robin coefficient and pointwise rescaling of the covariance operators [25, 15]. The covariance operator corresponding to the covariance kernel $c(x, y)$ is $(a^2 - \Delta)^{-(u+d/2)}$. In this paper, we work with covariance operators of the form C_a ,

$$(2.2) \quad C_a := (a^2 - \Delta)^{-a},$$

and $a := u + d/2$ with corresponding zero Neumann boundary conditions. The definition ensures that the covariance operator C is a self-adjoint, trace-class operator [30].

Furthermore, since it is more convenient to work with integer values of the exponent a , the constraint $a > d/2$ forces us to choose $a = 1$ for one spatial dimension and

$\alpha = 2$ for two or three spatial dimensions. However, in the context of Bayesian inverse problems this choice of α can result in oversmoothing of the solution and poor edge-preserving behavior of the solution. Furthermore, it is common to use parameters such as $u = 1/2, 3/2, 5/2$, etc.; for example, $u = 1/2$ corresponds to the well-known exponential covariance kernel. However, when $d = 2$, $\alpha = 3/2$ and $A^{-\alpha}$ cannot be computed using standard approaches. It is desirable to develop a numerical approach that can exploit the full power of these parameterized priors, without the constraints of integer powers. This motivates the use of fractional operators which we discuss in section 2.2.

Following [24, section 3.2] and [23, section 2.6], it is also straightforward to model nonstationary effects by considering the SPDE

$$(2.3) \quad (a^2(x) - a \operatorname{tr}(a(x)a))^{1/2} u = u(x)W, \quad x \in \mathbb{R}^d,$$

where $a^2(x)$, $u(x)$, and $a(x)$ are now spatially varying functions and W is a spatial Gaussian white noise process with unit variance. Assuming that $u = 1$, the corresponding covariance operator (with zero Neumann boundary conditions) is

$$(2.4) \quad C_a := (a^2(x) - a \operatorname{tr}(a(x)a))^{-1/2}.$$

Note that this assumes (2.2) as a special case.

2.2. Fractional Laplacian. For $\alpha^2 > 0$, let us use the notation $a^2 - a$ to denote the realization in $L^2(\Omega)$ of $a^2 - a$ with zero Neumann boundary conditions. Then there exists a sequence of eigenvalues $\{\lambda_k\}_{k=1}^\infty$ satisfying $0 < \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_k \leq \dots$ with $\lim_{k \rightarrow \infty} \lambda_k = \infty$. The corresponding eigenfunctions $\{\phi_k\}_{k=1}^\infty \subset H^1(\Omega)$. It is well known that $\{\phi_k\}_{k=1}^\infty$ forms the orthonormal basis of $L^2(\Omega)$. Towards, this end, for any $s \geq 0$, we can define the fractional order Sobolev space $H^s(\Omega)$ as

$$H^s(\Omega) := \left\{ u = \sum_{k=1}^\infty u_k \phi_k : \|u\|_{H^s(\Omega)}^2 := \sum_{k=1}^\infty \lambda_k^s u_k^2 < \infty \right\}.$$

We are now ready to define the fractional powers of $a^2 - a$, see [3, section 8].

Definition 2.1. Let $u \in C^0(\overline{\Omega})$ with $n \cdot \nabla u = 0$ on $\partial\Omega$, where n is the outward normal vector. Then the fractional power of $a^2 - a$ is given by

$$(a^2 - a)^s u := \sum_{k=1}^\infty \lambda_k^{-s} u_k \phi_k \quad \text{with} \quad u_k = \int_\Omega u \phi_k.$$

Notice that $(a^2 - a)^s$ can be extended to an operator mapping from $H^s(\Omega)$ to $H^{-s}(\Omega)$, where $H^{-s}(\Omega)$ is the dual of $H^s(\Omega)$. We further emphasize that the condition $\alpha^2 > 0$ can be dropped in the above definition; see [3, Def. 2.2]. We also refer to [3], for extensions to nonzero Neumann conditions.

Finally, we can extend the above definition to $(a^2(x) - a \operatorname{tr}(a(x)a))$ in (2.3). In this case, we need to assume $a^2 \in L^Y(\Omega)$ and a to be symmetric, bounded in $L^Y(\Omega)$, and elliptic.

2.3. Discretization. We show how to discretize the appropriate quantities in a finite-dimensional space. The discussion here closely mimics the derivation in [11].

Finite-dimensional parameter space. Let $\phi_1, \dots, \phi_n$ be a basis for the finite-dimensional space $V_h \subset L^2(\Omega)$ with $N_h = \dim(V_h)$, where the subscript h refers

to a mesh discretization parameter. Assume that the basis functions corresponding to the nodal points $\{x_j\}_{j=1}^{N_h}$ satisfy $i_j(x_i) = \delta_{ij}$ for $1 \leq i, j \leq N_h$. We can then approximate the inversion parameter within this finite-dimensional subspace as $m_h = \sum_{j=1}^{N_h} m_{h,j} i_j$, and we denote the vector of coefficients in the expansion $m = [m_{h,1} \ \dots \ m_{h,N_h}]^T \in \mathbb{R}^{N_h}$.

For two functions $u, v \in L^2(a)$, the L^2 inner product can be approximated as $(u, v)_{L^2(a)} \approx u^T M v$, where M is the mass matrix with entries

$$(2.5) \quad M_{ij} = \int_a^t i_i(x) i_j(x) dx, \quad 1 \leq i, j \leq N_h.$$

We denote $\mathbb{R}_M^{N_h}$ to be the vector space $\mathbb{R}^{N_h}$ with the inner product $\langle \cdot, \cdot \rangle_M$, to distinguish from the usual Euclidean space.

Let $B : L^2(a) \rightarrow L^2(a)$. Then the matrix approximation of $B_h : V_h \rightarrow V_h$, denoted $B : \mathbb{R}_M^{N_h} \rightarrow \mathbb{R}_M^{N_h}$, is obtained as follows. First, we construct a matrix S with entries

$$S_{ij} = \int_a^t i_i(x) B i_j(x) dx = \langle e_i, B e_j \rangle_M, \quad 1 \leq i, j \leq N_h.$$

Here i_i are the finite-dimensional basis vectors and e_i is the canonical basis vector for $\mathbb{R}^{N_h}$ corresponding to the basis function i_i . This gives us the matrix representation $B = M^{-1} S$. The matrix transpose of B , denoted B^T , has entries $[B^T]_{ij} = b_{ji}$, but the adjoint of $B : \mathbb{R}_M^{N_h} \rightarrow \mathbb{R}_M^{N_h}$, denoted B^t , satisfies $B^t = M^{-1} B^T M$.

Discretization of the covariance operator. In the following discussion, we first assume that $0 < s < 1$ and $A := a^2 - a \int_a^t (a(x)a) dx$, with zero Neumann boundary conditions. To apply the fractional operator $u = A^{-s} f$, we use sinc quadrature to approximate the integral [9]. Let $z_{h,j} \in V_h$ solve

$$\int_a^t ((a - a z_{h,j}) \tau a) v_h + (a^2 + e^{ja}) z_{h,j} v_h dx = \int_a^t f v_h dx \quad \forall v_h \in V_h, \quad -M_- \leq j \leq M_+,$$

and compute the approximation to u as

$$u_h = \sum_{j=-M_-}^{M_+} \frac{a \sin(\pi j a)}{i} e^{(1-s)ja} z_{h,j},$$

where $M_+ = \lfloor \frac{i^2}{4sa} \rfloor$, $M_- = \lfloor \frac{i^2}{4(1-s)a} \rfloor$, and $a = 1/\log(1/h)$. The number of terms in the expansion was chosen to balance the errors in the quadrature with the finite element discretization error. Alternatively, in matrix notation, we can write

$$u = \sum_{j=-M_-}^{M_+} w_j (K + z_j M)^{-1} M f,$$

where K is the stiffness matrix with entries

$$[K]_{ij} = \int_a^t (a - a i_i) \tau a i_j + a^2 i_i i_j dx, \quad 1 \leq i, j \leq N_h$$

and M is the mass matrix with entries given in (2.5). The quadrature weights and nodes are

$$(2.6) \quad w_j = \frac{a \sin(\pi j a)}{i} e^{(1-s)ja}, \quad z_j = e^{ja}, \quad -M_- \leq j \leq M_+.$$

Based on the preceding discussion, the matrix representation of the discretized version of $D = A^{-s}$ is

$$D := \sum_{j=-M}^M w_j (K + z_j M)^{-1} M.$$

Let N denote the set of natural numbers (excluding zero). Now consider the covariance operator $C_a = A^{-a}$, where $a > d/2$ and $a \notin N$. We write $a = r + s$, where $r = \text{rar}$ is the integer part of a and $0 < s < 1$ is the fractional part. Following [8, Remark 3.6], the discretized representation of the covariance operator $C_a = A^{-(u+d/2)}$ is

$$(2.7) \quad C_a := \sum_{j=-M}^M w_j (K + z_j M)^{-1} M [K^{-1} M]^r.$$

Note that we can interchange the order of the fractional and the integer part since terms like $(K + z_j M)^{-1} M$ and $K^{-1} M$ commute. Furthermore, the matrix C_a is self-adjoint, i.e., $C_a^t = C_a$. An important point to note here is that C_a is not formed explicitly. The application of the matrix to a vector can be obtained by $M_+ + M_+ + 1$ independent linear system solves which can be readily parallelized. However, we will show how to compute the action of C_a efficiently using Krylov subspace methods for shifted linear systems (see section 3.2).

2.4. Related work. As mentioned earlier, in Bayesian inverse problems, the exponent is typically chosen to be $a = 2$. To our knowledge, in Bayesian inverse problems, the problem of efficient computations with covariance operators for non-integer values of a is still an outstanding challenge. The fractional Laplacian with $0 < a < 1$ has been used as a regularization operator in a non-Bayesian setting [1, 2]. Notice that, [1] uses the Fourier approach which is limited to periodic settings and [2] uses an eigendecomposition of the Laplacian that is not scalable to large problems.

There are some techniques in computational and spatial statistics that are relevant to this discussion; however, for a detailed review, see [23]. In [28], a method was proposed for sampling from generalized Matérn fields on compact Riemannian manifolds using the so-called matrix transfer technique. In this method, after discretizing the differential operator, the samples are obtained by applying the matrix function $(f(x) = x^{-a})$ on a random vector. The application of the matrix function is accomplished using contour integral and Krylov subspace methods similar to ours. But our approach differs in how we approximate the fractional operator and also in the choice of Krylov subspace methods. Recent work in [7, 8] developed efficient methods for sampling from the SPDE for certain fractional exponents (in our notation, $d/4 < a/2 < 1$); these references also use the same quadrature scheme for the integral representation of the fractional inverse as ours. However, our approach generalizes it to all exponents $a > d/2$ and the use of Krylov subspace solvers substantially speeds up the solution of the SPDE. In a follow-up work [6], the authors developed the rational SPDE approach in which they use a rational approximation of the form (in our notation) $A^{-a} \approx p(A)q(A)^{-1}$, where p, q are polynomials. When the degrees of the polynomials are small, the discretized representations retain the computational benefits of the integer versions. However, the error analysis in [6, Remark 3.4] suggests in order to balance the errors due to the spatial discretization and the error in the rational approximation, the degrees of the polynomials p, q may be large, which leads to a loss in computational efficiency. The paper in [18] also works in a similar setting as ours in bounded domains but solves the linear systems in parallel using multilevel

preconditioned iterative solvers with linear complexity in the degrees of freedom; additionally, their approach is also applicable to compact metric spaces. Also relevant to this discussion are [22, 20] which deal with sampling from random fields on compact Riemannian manifolds. Our efficient approach for representing covariance matrices and sampling from the random field may be applicable to that setting as well. To our knowledge, these techniques have not been used within the context of Bayesian inverse problems, which is the primary focus of this paper.

3. Efficient computations with the covariance operator. In this section, we review the multipreconditioned GMRES method for shifted linear systems [4] (MPGMRES-Sh) for solving sequences for shifted linear systems (section 3.1). We use this solver to accelerate the computations with the covariance operator C_a and for generating samples from the SPDE (section 3.2). Finally, in section 3.3, we also describe how to efficiently compute the truncated Karhunen-Loève expansion.

3.1. Krylov methods for shifted systems. The dominant cost of applying the fractional operator is the solution of a system of shifted linear equations. We first briefly review the MPGMRES-Sh approach [4]. This is an efficient method for solving shifted linear systems of the form

$$(A_1 + a_j A_2)x_j = d, \quad 1 \leq j \leq N_a,$$

where $A_1, A_2 \in \mathbb{R}^{n_s \times n}$ and $\{a_j\}_{j=1}^{N_a}$ are a set of shifts. This solver can be used in multiple ways to accelerate computations involving Gaussian random fields.

In the MPGMRES-Sh approach, we select a set of shifts $\{u_j\}_{j=1}^{n_p}$ that determine the preconditioners $P_j = (A_1 + u_j A_2)$. The method uses multiple preconditioners to build a basis $Z_m \in \mathbb{R}^{n_s \times m n_p}$ that we will use to define a search space for all the shifted systems. In the discussion below, for simplicity, we drop the index on a . At the k th iteration, given the matrix $V^{(k)} \in \mathbb{R}^{n_s \times n_p^{k-1}}$ the method builds the matrix $Z^{(k)}$ as

$$Z^{(k)} = \begin{bmatrix} P_1^{-1} v_k & \dots & P_{n_p}^{-1} v_k \end{bmatrix} \in \mathbb{R}^{n_s \times n_p}, \quad v_k = V^{(k)} e_{n_p}.$$

The method is initialized with $V^{(1)} = d/\|d\|_2$. We define the matrices that collect the columns $V^{(k)}$ and $Z^{(k)}$ as $V_m = \begin{bmatrix} V^{(1)} & \dots & V^{(m)} \end{bmatrix}$ and $Z_m = \begin{bmatrix} Z^{(1)} & \dots & Z^{(m)} \end{bmatrix}$. Using these relationships we can derive the multipreconditioned Arnoldi relationship for a shift a :

$$(3.1) \quad (A_1 + a A_2) Z_m = V_m \begin{pmatrix} E_m \\ 0 \end{pmatrix} + H_m (a I - T_m) = V_{m+1} H_m(a; T_m).$$

Here the matrices T_m and E_m are defined as follows. Let s denote the Kronecker product and define

$$T^{(k)} = \text{blkdiag}(I_{n_p^{k-1}} s u_1, \dots, I_{n_p^{k-1}} s u_{n_p}) \in \mathbb{R}^{n_p^k \times n_p^k}$$

and $E^{(k)} = e_{n_p}^p s I_{n_p^{k-1}}$ for $1 \leq k \leq m$. Here blkdiag means a block diagonal matrix composed of the given subblocks. With these definitions in place, then $T_m = \text{blkdiag}(T^{(1)}, \dots, T^{(m)})$ and $E_m = \text{blkdiag}(E^{(1)}, \dots, E^{(m)})$. The matrix H_m is block upper-Hessenberg. The importance of (3.1) is that the basis Z_m is applicable across all the shifts and need not be derived from scratch for each new shifted system that needs to be solved. This is a consequence of the shift-invariant property of MPGMRES-Sh. Furthermore, the growth in the number of columns in the search space is linear in the

number of preconditioners rather than exponential; see [4], in particular Theorem 3.6 and the discussion in section 3.2 of that paper.

The main point here is that the solution for each shifted system corresponding to shift a is obtained as $x_m(a) = Z_m y_m(a)$, where $y_m(a) \in \mathbb{R}^{m \times p}$ minimizes the residual and is obtained by solving the least squares problem

$$\min_{y \in \mathbb{R}^{m \times p}} \|r_m(a)\|_2 = \| |d|_2 e_1 - A_m(a; T_m) y \|_2.$$

The details of this algorithm as well as the implementation details are given in [4, Algorithm 2].

Computational cost. We do assume that the cost of matrix-vector products (matvecs) involving A_1 and A_2 is T_{matvec} and the cost of applying the preconditioner is T_{prec} . If k_m iterations of the MPGMRES-Sh are performed, then the total cost is

$$T_{\text{cost}} = k_m n_p (T_{\text{matvec}} + T_{\text{prec}}) + O(N_a n k_m^2 n_p^2) + O(N_a k_m^3 n_p^3) > \text{flops}.$$

As we shall see in the numerical experiments, the number of iterations k_m is low (typically $q \leq 50$). Note that the cost of building the basis is independent of the number of shifted systems, whereas the cost of the projected system depends on the number of shifted systems N_a .

3.2. Applying the covariance operator. We now show how to compute the application of the covariance operator C_a to a vector f . Let the matrices K and M be as defined in section 2.3.

Let $a \in \mathbb{R}_+$ be a positive number. If $a \in \mathbb{N}$, then it is straightforward to compute an application of the covariance operator which requires a sequential applications of solves involving K . Henceforth, we assume that $a \in \mathbb{R}_+ \setminus \mathbb{N}$. Let r be the integer part of a and let $s \in (0, 1)$ be the fractional part. We can apply the covariance operator to a vector f ; this can be computed as

$$(3.2) \quad C_a f = \sum_{j=-M}^M \frac{w_j}{z_j} \left(M + z_j^{-1} K \right)^{-1} M \left[K^{-1} M \right]^r f.$$

In writing this expression, we have replaced the weights with $w_j z_j^{-1}$. This form is completely equivalent to (2.7) but makes it easier to precondition. To turn this expression into an efficient algorithm, we perform the following steps. In the offline phase, we determine the quadrature weights $\{w_j\}_{j=-M}^M$ and nodes $\{z_j\}_{j=-M}^M$, preconditioner shifts $\{u_j\}_{j=1}^q$. Furthermore, we compute factorizations of K and the preconditioners $\{P_j = M + u_j K\}_{j=1}^q$. In the online stage, we compute the action of $K^{-1} M$ on f r times to obtain $\sum_{j=1}^r [K^{-1} M]^r f$. We then apply MPGMRES-Sh with $A_1 = M$, $A_2 = K$, $d = M c$, and shifts $a_j = 1/z_j$ for $-M \leq j \leq M$ to obtain the solutions $\{x_j\}_{j=-M}^M$. Finally, we compute $C_a f = \sum_{j=-M}^M w_j x_j$. The details are given in Algorithm 3.1.

When the integer part r is large, the number of solves with K can become expensive; however, in practice, a value of $r > 3$ is typically not used since the resulting field is extremely smooth. If an application requires the use of a large value of r , techniques from [19, section 4.1] may be used.

Alternative approaches. In the expression for the application of the covariance defined in (3.2), we used a slight reformulation for the fractional part, which is obtained

Algorithm 3.1 Applying the covariance operator $C_a f$.

Require: matrices K and M , vector f , scalar $a \neq 0$, preconditioner shifts $\{u_j\}_{j=1}^{n_p}$

- 1: {Offline stage: precomputation}
- 2: Split $a = r + s$, where $s \in (0, 1)$ and $r \in \mathbb{N}$
- 3: Determine quadrature weights w_j and nodes z_j for M . $1 \leq j \leq M$
- 4: Factorize $P_j = M + u_j K$ for $1 \leq j \leq n_p$ and K
- 5: {Online stage: solve}
- 6: Apply the integer part: $c = [K^{-1}M]^r f$
- 7: Solve the sequence $(M + z_j^{-1}K)x_j = Mc$ and compute $u = \sum_{j=M_-}^{M_+} w_j x_j$.
- 8: return $u = C_a f$

by factoring out z_j in each summand. We also investigated two other formulations for computing Db . The first is perhaps the more natural version,

$$Db = \sum_{j=-M_-}^{M_+} w_j (z_j M + K)^{-1} Mb.$$

While mathematically equivalent to the approach used in (3.2), we found that it was easier to precondition the formulation in (3.2). The second approach is to decompose

$$Db = \sum_{j=-M_-}^{M_+} w_j (z_j M + K)^{-1} Mb + \sum_{j=1}^{M_+} \frac{w_j}{z_j} (M + z_j^{-1}K)^{-1} Mb.$$

The advantage of this approach is that in this instance both the sets of weights, w_j and w_j/z_j , are at most one in their respective intervals $-M_- \leq j \leq 0$ and $1 \leq j \leq M_+$; similarly, the nodes z_j and z_j^{-1} are also at most 1 in their respective intervals. However, the downside is that to apply the fractional operator, we need two solves with MPGMRES-Sh. For this reason, we prefer the formulation in (3.2).

In the offline stage Algorithm 3.1, we factorize the preconditioners. However, this may not be feasible for very large-scale problems, especially in three spatial dimensions. In such cases, we can apply the preconditioners using iterative solvers. See [4, section 4.2.2].

3.2.1. Generating samples from the SPDE. We can adapt the efficient approach in Algorithm 3.1 to generate samples from the SPDE,

$$(3.3) \quad (a^2(x) - a^2(x))^{(u+d/2)/2} u(x) = W(x), \quad x \in \mathbb{R}^d,$$

with zero Neumann boundary conditions. We follow the approach in [7] to solve the SPDE. Let $LL^T = M$ be the Cholesky factorization of M and let $w \in \mathbb{N}(0, I)$. For $0 < a/2 < 1$, we can compute a solution to the SPDE as

$$u := \sum_{j=-M_-}^{M_+} \frac{2a \sin(i a/2)}{i} e^{(1-a/2)aj} (K + e^{ja} M)^{-1} Lw,$$

where $a = 1/\log(1/h)$, and the number of shifted systems $N_a = M_- + M_+ + 1$, where

$$M_+ = \lceil i^2/(2a^2) \rceil \quad \text{and} \quad M_- = \lceil i^2/(2(2-a)^2) \rceil.$$

We can readily modify Algorithm 3.1 to efficiently solve this problem. In particular, we can leverage the use of the MPGMRES-Sh approach to accelerate the solution to the SPDE. If we instead use zero Dirichlet boundary conditions for u , the error analysis in [7, Theorem 2.1] applies directly. If the Cholesky factorization of M is expensive to compute, we can instead use the square root of M applied to w using the MPGMRES-Sh technique; see [4, section 5].

3.3. Karhunen--Loève expansions. In this section, we discuss an approach to generate samples from the Gaussian measure $N(u, C)$ on a Hilbert space H , where $u \in H$ is the mean function and C is a self-adjoint, positive semidefinite covariance operator. This approach is based on the Karhunen--Loève (KL) expansion of the stochastic process. In this approach we consider an orthonormal set of eigenpairs (a_j, i_j) for $1 \leq j \leq \infty$, where the eigenvalues are arranged in decreasing order as $a_1 \geq a_2 \geq \dots$. Furthermore, let $\{i_j\}_{j=1}^\infty$ be an independent and identically distributed sequence of random variables with $i_1 \sim N(0, 1)$. Then by [30, Theorem 6.19], the sample u defined by the KL expansion

$$u := u + \sum_{k=1}^{\infty} \bar{a}_j i_j i_j,$$

is distributed according to $N(u, C)$. In practice, for certain applications, the eigenvalues exhibit rapid decay and, therefore, the expansion can be approximated by the truncated representation $u_N := u + \sum_{j=1}^N \bar{a}_j i_j i_j$. In the context of Bayesian inverse problems, rather than estimating the field, one can estimate the coefficients $\{i_j\}_{j=1}^N$ of the KL expansion. This can be computationally beneficial because of the input dimensionality reduction.

Galerkin approximation. In this discussion, we will derive the expressions for the truncated KL expansion for the Gaussian measure $N(0, C)$. We will then discuss implementation details for computing the truncated KL expansion for the Gaussian measure $N(0, C_a)$. Consider a finite-dimensional subspace $V_h \subset L^2(a)$ with $N_h = \dim(V_h)$ and let $\{i_j\}_{j=1}^{N_h}$ be a basis for the subspace. The eigenvalue problem $C i = a i$ is then replaced by the finite-dimensional eigenvalue problem $C_h i_h = a_h i_h$, where $i_h \in V_h$ and $C_h : V_h \rightarrow V_h$. We seek a solution $i_h = \sum_{j=1}^{N_h} i_{h,j} i_j$ by considering the Galerkin projection

$$\sum_{j=1}^{N_h} i_{h,j} (i_i, C_h i_j)_{L^2(a)} = \sum_{j=1}^{N_h} i_{h,j} (i_i, i_j)_{L^2(a)}, \quad 1 \leq i \leq N_h.$$

Following the discussion in section 2.3, we can write $(i_i, i_j)_{L^2(a)} = \delta_{ij}$ and $(i_i, C_h i_j)_{L^2(a)} = C_{ij}$, where C is the discretized covariance matrix. Therefore, we now have the generalized Hermitian eigenvalue problem (GHEP)

$$(3.4) \quad M C i_h = a_h M i_h,$$

where $i_h = [i_{h,1} \dots i_{h,N_h}]^T \in \mathbb{R}^{N_h}$ is the generalized eigenvector. Let $a_{h,1} \geq \dots \geq a_{h,N_h}$ be the generalized eigenvalues and let $i_{h,j}$ be the corresponding eigenvectors. Then a sample from the process can be generated as

$$(3.5) \quad u_{h,N} = \sum_{j=1}^{N_h} \bar{a}_{h,j} i_{h,j} i_{h,j}.$$

Efficient implementation. We want to compute the truncated KL for the measure $N(0, C_a)$. The GHEP (3.4) can be solved using any standard Krylov-based eigensolver (e.g., Lanczos) that does not form the matrix C_a or MC_a explicitly but instead relies on forming the matvecs involving the matrices. The matvecs involving C_a can be efficiently computed using the MPMRES-Sh approach as described in Algorithm 3.1.

An alternative approach to computing the truncated KL expansion is to work with the eigenpairs of the operator A , which is attractive at first glance since it avoids the complications with computing the fractional part. However, the approach based on the covariance operator C_a directly is advantageous for two reasons. First, to compute the truncated KL expansion, we need to target the eigenpairs of C_a corresponding to the largest eigenvalues, which corresponds to the eigenpairs of A corresponding to the smallest eigenvalues, which is typically much harder. Second, for $\alpha > 1$ the operator $A^{-\alpha}$ acts as a spectral transformation that enhances the eigenvalue gaps.

4. Application to Bayesian inverse problems. In this section, we show how to leverage the efficient representation of the covariance operator. In section 4.1, we review the necessary background for Bayesian inverse problems. In section 4.2, we derive the discretized posterior distribution that uses the covariance matrix in (2.7) as the prior covariance and in section 4.3, we show how to adapt generalized hybrid iterative methods (GenHyBR) appropriately to efficiently compute the MAP estimate. The material in sections 4.1 and 4.2 follow the discussions in [30, 11].

4.1. Bayesian inverse problems. Let $m \in L^2(\mathcal{A})$ be the inversion parameter that we wish to recover from the measurements $y \in \mathbb{R}^{N_y}$, which are related through the measurement equation

$$(4.1) \quad y = F(m) + a,$$

where $F : L^2(\mathcal{A}) \rightarrow \mathbb{R}^{N_y}$ is the parameter-to-observable map or the measurement operator, and a is the measurement noise which is assumed to be Gaussian with zero mean and covariance a_{noise} , which we write as $N(0, a_{\text{noise}})$. Therefore, the likelihood probability density function takes the form

$$i(y|m) \propto \exp \left(-\frac{1}{2} |y - F(m)|_{a_{\text{noise}}}^2 \right).$$

We assume that m is endowed with a Gaussian prior with mean function m_{pr} and covariance operator $a_c^{-2} C_{\text{pr}}$, i.e., $m \sim N(m_{\text{pr}}, a_c^{-2} C_{\text{pr}})$. Here a_c^2 acts as a regularization parameter that may be known, or can be estimated as part of the inversion process (this is what we do in this paper). For the operator C_{pr} we take it to be C_a as defined in section 2.1.

By the choice of the likelihood and the prior distribution, using the infinite-dimensional Bayes formula, the Radon-Nikodym derivative of the posterior probability measure u_{post} with respect to the prior measure $u_{\text{pr}} = N(m_{\text{pr}}, a_c^{-2} C_{\text{pr}})$, takes the form

$$\frac{du_{\text{post}}}{du_{\text{pr}}} = \frac{1}{Z} \exp \left(-\frac{1}{2} |y - F(m)|_{a_{\text{noise}}}^2 \right),$$

where $Z = \int i(y|m) du_{\text{pr}}$ is a normalization constant. The MAP estimator maximizes the posterior distribution and can be obtained by the solution to the optimization problem

$$\min_{m \in L^2(\mathcal{A})} \frac{1}{2} |y - F(m)|_{a_{\text{noise}}}^2 + \frac{\hat{\epsilon}^2}{2} |m - m_{\text{pr}}|_{C_{\text{pr}}^{-1}}^2.$$

If F is a linear operator, then the posterior distribution is Gaussian; we exploit this fact in this paper.

4.2. Discretized posterior distribution. The covariance matrix C_{pr} after discretization takes the same form as in (2.7). It is easy to verify that the resulting matrix C_{pr} is self-adjoint; i.e., $C_{pr}^t = C_{pr}$ and $C_{pr}M^{-1}$ is symmetric with respect to the standard Euclidean inner product and positive definite. This gives us a discretized representation of the prior distribution $m \sim N(m_{pr}, a_c^{-2}C_{pr}M^{-1})$, i.e.,

$$i_{pr}(m) \propto \exp \left(-\frac{a_c^2}{2} (m - m_{pr})^T C_{pr}^{-1} (m - m_{pr}) \right).$$

For completeness, we give expressions for the posterior distribution and the MAP estimate, for the discretized problem. This uses the components developed in section 2.3. The posterior probability density function takes the form

$$i_{post}(m|y) \propto \exp \left(-\frac{1}{2} \|f(m) - y\|_{a_n^{-1}C_{noise}}^2 - \frac{a_c^2}{2} \|m - m_{pr}\|_{C_{pr}^{-1}M}^2 \right).$$

The MAP estimator can be obtained by solving the optimization problem

$$\min_{m \in \mathbb{R}^{N_h}} \frac{1}{2} \|f(m) - y\|_{a_n^{-1}C_{noise}}^2 + \frac{a_c^2}{2} \|m - m_{pr}\|_{C_{pr}^{-1}M}^2.$$

If the forward operator is linear $f(m) = Fm$, then the posterior distribution is Gaussian with $m|y \sim N(m_{post}, C_{post}M^{-1})$, where

$$(4.2) \quad \begin{aligned} C_{post} &:= (F^P a_{noise}^{-1} F + a_c^2 C_{pr}^{-1})^{-1}, \\ m_{post} &:= C_{post} (F^P a_{noise}^{-1} y + a_c^2 C_{pr}^{-1} m_{pr}), \end{aligned}$$

and $F^P : \mathbb{R}^{N_y} \rightarrow \mathbb{R}^{N_h}$ is the adjoint operator of F , and takes the form $F^P = M^{-1}F^T$. Note that C_{post} is self-adjoint with respect to the $\langle \cdot, \cdot \rangle_M$ inner product, so that $C_{post}M^{-1}$ is symmetric positive definite.

4.3. GenHyBR for computing the MAP estimate. We now derive an alternate expression for the MAP estimate (4.2). Multiplying both sides of (4.2) with C_{post}^{-1} , we get

$$(F^P a_{noise}^{-1} F + a_c^2 C_{pr}^{-1}) m_{post} = F^P a_{noise}^{-1} y + a_c^2 C_{pr}^{-1} m_{pr}.$$

We use $F^P = M^{-1}F^T$, and the change of variables $\tilde{m}_{post} = MC_{pr}^{-1}(m_{post} - m_{pr})$. Then to obtain the MAP estimate, we solve the linear system for $\tilde{m}_{post}$,

$$(4.3) \quad (F^T a_{noise}^{-1} F C_{pr} M^{-1} + a_c^2 I) \tilde{m}_{post} = F^T a_{noise}^{-1} (y - F m_{pr}),$$

and then compute $m_{post} = m_{pr} + C_{pr}M^{-1}\tilde{m}_{post}$. This form of the MAP estimator will be the basis for our computationally efficient procedure. Note that $C_{pr}M^{-1}$ is a symmetric positive definite matrix with respect to the Euclidean inner product; we denote this by $Q = C_{pr}M^{-1}$ for convenience. The application of Q to a vector can be performed efficiently, using a slight modification of Algorithm 3.1. Furthermore, note that $F^T a_{noise}^{-1} F Q$ is symmetric with respect to the $\langle \cdot, \cdot \rangle_Q$ inner product.

We can reformulate the equations in such a way that we can use the generalized Golub--Kahan bidiagonalization (gen-GK) approach [13] for efficiently estimating the

MAP estimate. We initialize the iterations with $\mathbf{b} = \mathbf{y} - \mathbf{F}\mathbf{m}_{\text{pr}}$, $\mathbf{a}_1 = \|\mathbf{b}\|_{\mathbf{a}^{-1}}^{-1}$, $\mathbf{u}_1 = \mathbf{y}_1/\mathbf{a}_1$, and $\mathbf{a}_1\mathbf{v}_1 = \mathbf{F}^\top \mathbf{a}_{\text{noise}}^{-1} \mathbf{b}$. At step k in this approach,

$$\begin{aligned}\mathbf{a}_{k+1}\mathbf{u}_{k+1} &= \mathbf{F}\mathbf{Q}\mathbf{v}_k - \mathbf{a}_k\mathbf{u}_k, \\ \mathbf{a}_{k+1}\mathbf{v}_{k+1} &= \mathbf{F}^\top \mathbf{a}_{\text{noise}}^{-1} \mathbf{u}_{k+1} - \mathbf{a}_{k+1}\mathbf{v}_k,\end{aligned}$$

where $\mathbf{a}_i, \mathbf{a}_i \geq 0$ are chosen such that $\|\mathbf{u}_i\|_{\mathbf{a}^{-1}} = \|\mathbf{v}_i\|_{\mathbf{Q}} = 1$ for $1 \leq i \leq k+1$. These iterates can be collected to form the matrices $\mathbf{U}_{k+1} := [\mathbf{u}_1 \dots \mathbf{u}_{k+1}]$, $\mathbf{V}_k := [\mathbf{v}_1 \dots \mathbf{v}_k]$, and the bidiagonal matrix

$$\mathbf{B}_k := \begin{bmatrix} \mathbf{a}_1 & & & \\ & \mathbf{a}_2 & & \\ & & \ddots & \\ & & & \mathbf{a}_k \\ & & & & \mathbf{a}_{k+1} \end{bmatrix} \in \mathbb{R}^{(k+1) \times k}.$$

This can be rearranged to obtain the relations, which are typically accurate up to machine precision,

$$\begin{aligned}\mathbf{F}\mathbf{Q}\mathbf{V}_k &= \mathbf{U}_{k+1}\mathbf{B}_k, \\ \mathbf{F}^\top \mathbf{a}_{\text{noise}}^{-1} \mathbf{U}_{k+1} &= \mathbf{V}_k\mathbf{B}_k + \mathbf{a}_{k+1}\mathbf{v}_{k+1}\mathbf{e}_{k+1}^\top,\end{aligned}$$

and the orthogonality relations $\mathbf{U}_{k+1}^\top \mathbf{a}_{\text{noise}}^{-1} \mathbf{U}_{k+1} = \mathbf{I}_{k+1}$ and $\mathbf{V}_k^\top \mathbf{Q}\mathbf{V}_k = \mathbf{I}_k$. The columns of $\mathbf{V}_k$ form a basis for the Krylov subspace $\mathcal{K}_k(\mathbf{F}^\top \mathbf{a}_{\text{noise}}^{-1} \mathbf{F}\mathbf{Q}, \mathbf{F}^\top \mathbf{a}_{\text{noise}}^{-1} \mathbf{b})$, where the Krylov subspace is defined as $\mathcal{K}_k(\mathbf{J}, \mathbf{d}) := \text{span}\{\mathbf{d}, \mathbf{J}\mathbf{d}, \dots, \mathbf{J}^{k-1}\mathbf{d}\}$. To obtain the approximate solution, we search for linear combinations of the columns $\mathbf{V}_k$, i.e., $\mathbf{m}_{k,a} = \mathbf{V}_k \mathbf{z}_{k,a}$ that solve the optimization problem

$$\min_{\mathbf{z} \in \mathbb{R}^k} \frac{1}{2} \|\mathbf{B}_{k+1} \mathbf{z} - \mathbf{a}_1 \mathbf{e}_1\|_2^2 + \frac{\mathbf{a}_c^2}{2} \|\mathbf{z}\|_2^2.$$

To estimate the regularization parameter $\mathbf{a}_c$, we minimize the projected generalized cross validation (GCV). Other choices for estimating the regularization parameter are possible (e.g., discrepancy principle, unbiased predictive risk estimate), but we do not pursue them here. We refer the reader to [13] for additional details. To terminate the iterations we use a stopping criterion similar to [12, section 5.5]. Suppose the iterations terminate at step k , then the solution to the MAP estimate can be obtained by undoing the change of variables; that is, we compute the approximate solution $\mathbf{m}_{\text{post}}^{(k)} := \mathbf{m}_{\text{pr}} + \mathbf{C}_{\text{pr}} \mathbf{M}^{-1} \mathbf{m}_{k,a}$.

It is worth mentioning that each iteration of gen-GK requires two matvecs with $\mathbf{Q}$, one with $\mathbf{F}$, and one with $\mathbf{F}^\top$. There is an additional cost of $\mathcal{O}(k^3)$ for solving the projected least-squares problem and estimating the regularization parameter $\mathbf{a}_c^2$, and an additional $\mathcal{O}(k(m+n))$ for orthogonalization. The additional cost of the extra matvec with $\mathbf{Q}$ is offset by the fact that we can efficiently estimate the regularization parameter during the iterative scheme.

4.4. Posterior variance. Methods to compute the approximate posterior variance have been developed in several references [16, 11, 31, 29, 26]. However, these are not directly applicable here and have to be suitably modified. We follow the approach in [26], that reutilizes the Gen-GK basis vectors computed during the computation of the MAP estimate, but we need to modify it in two different ways: first, since we are

approximating infinite-dimensional quantities, we have to work with additional mass matrices and, second, we no longer have ready access to the diagonals of Q .

The variance field corresponding to a covariance operator C can be obtained using the discussion in [11]. Once again suppose that we are given a basis $\{i_j\}_{j=1}^{N_h}$. Given the discretized representation C of the covariance operator C , then the discretized covariance function

$$c_h(x, y) = [i(x)]^T C M^{-1} i(y),$$

where $i(x) = [i_1(x), \dots, i_{N_h}(x)]^T$ is a vector containing the finite element basis functions. To visualize the covariance function, we consider the discretized covariance function evaluated at the nodal points $\{x_j\}_{j=1}^{N_h}$. Therefore, it is sufficient to consider the diagonals of the matrix $C M^{-1}$. For the posterior variance, therefore, we have to compute the diagonals of $C_{\text{post}} M^{-1}$; that is, we need to compute the diagonals of

$$C_{\text{post}} M^{-1} = (F^T a_{\text{noise}}^{-1} F + a_{\text{pr}}^2 C_{\text{pr}}^{-1})^{-1} M^{-1} = (F^T a_{\text{noise}}^{-1} F + a_{\text{pr}}^2 Q^{-1})^{-1}.$$

This can be efficiently estimated using the intermediate computations in the Gen-GK process [14, 26]. More precisely, we can approximate

$$Q F^T a_{\text{noise}}^{-1} F Q \approx Q V_k B_k^T B_k V_k^T Q = Z_k \Lambda_k Z_k^T,$$

where $B_k^T B_k = W_k \Lambda_k W_k^T$ is the eigendecomposition of $B_k^T B_k$ and $Z_k = Q V_k W_k$. Then we have the following approximation to the posterior covariance:

$$C_{\text{post}} \approx \mathbb{C}_{\text{post}} := Q(Z_k \Lambda_k Z_k^T + a_{\text{pr}}^2 Q)^{-1} Q M.$$

Applying the Woodbury identity, we get

$$\mathbb{C}_{\text{post}} = a_{\text{pr}}^{-2} Q M - Z_k \Lambda_k Z_k^T M,$$

where $\Lambda_k \in \mathbb{R}^{k \times k}$ is a diagonal matrix with the i th diagonal $a_{\text{pr}}^{-2} i_i / (i_i + a_{\text{pr}}^2)$ for $1 \leq i \leq k$. Finally, we can get the approximate posterior variance from the diagonals of

$$\mathbb{C}_{\text{post}} M^{-1} = a_{\text{pr}}^{-2} Q - Z_k \Lambda_k Z_k^T.$$

Since $Z_k \Lambda_k Z_k^T$ is low rank, its diagonals can be easily computed. Finally, since in floating point arithmetic, the vectors V_k lose orthogonality, we use full reorthogonalization which adds additional cost; see [14, 26]. To estimate the diagonals of Q we follow the Diag++ approach in [5, Algorithm 1]. This is a Monte Carlo-based approach that only uses matvecs with Q (see Algorithm 3.1) to estimate the diagonals.

5. Numerical experiments. We present a suite of numerical experiments that demonstrate the performance of the proposed methods. The timing results are reported on a Mac Mini (M1, 2020) with 16 GB memory and running macOS Big Sur 11.2.3 and MATLAB 2021A.

5.1. Application of the covariance operator. We perform some numerical experiments demonstrating the efficiency of the MPMRES-Sh solver in computing the action of the covariance operators. We take the domain to be $\Omega = (0, 1)^2$ and the number of grid points to vary from $33 \leq N \leq 513$. We take $C_a = A^{-s}$,




Fig. 1. (left) Timing results with increasing discretization for computing $C_a f$, where $a = 0.5$. Direct refers to solving the sequence of shifted linear systems individually using a direct solver, Preconditioner refers to the time required to construct the preconditioners $\{P_j\}_{j=1}^{n_p}$. (right) The number of iterations at convergence taken by MPGMRES-Sh for each shift; note that a single basis was used across all the systems. The vertical black lines denote the values of u_j used for preconditioners.

where $A = a^2 - a$, $a = \sin(0, 1)$. We also take $a^2 = 100$ and the Laplacian has zero Neumann boundary conditions. Although this operator is not trace-class, the goal here is merely to study the performance of MPGMRES-Sh and the results are applicable to other values of $a > d/2$.

Since the shifts $a_j = 1/z_j$ are on the positive real axis, we use three preconditioners $n_p = 3$ with $u \in \{10^{-8}, 10^{-4}, 10^{-2}\}$. The preconditioners P_j for $j = 1, \dots, n_p$ are factorized and stored ahead of the MPGMRES-Sh iterations. We stopped the MPGMRES-Sh iterations when the relative residuals of each shifted system were smaller than 10^{-8} . We investigated different numbers of shifts and with different shifts, but found the present setup to strike a balance between the number of iterations and the cost per iteration.

Varying mesh discretization. In this experiment, we study the effect of mesh refinement on the performance of the MPGMRES-Sh solver. We vary the grid sizes from 33×33 to 513×513 . In the largest problem instance, the number of degrees of freedom was 263169. For the MPGMRES-Sh algorithm, we still used the same three shifts to determine the preconditioners. We see that even though the number of shifted systems to be solved increases with mesh refinement, the number of MPGMRES-Sh iterations only rises by a modest amount. The highest dimension of the basis for the search space is $22 \times 3 = 66$. In Figure 1, we compare the timing cost of the MPGMRES-Sh with other approaches with mesh refinement. The "Direct" approach refers to solving the sequence of shifted systems using a direct solver. We also report the time taken in factorizing the preconditioners $\{P_j\}_{j=1}^{n_p}$, labeled "Preconditioner," and the time taken for solving the shifted linear systems with MPGMRES-Sh is labeled "MPGMRES-Sh." We see that the cost of MPGMRES-Sh is much smaller than the direct method and is comparable to the cost of factorizing the preconditioners. For the largest problem size, there is an $m40X$ speedup of MPGMRES-Sh compared to direct. The computational gains are more pronounced as the system size gets larger since the number of shifted systems grows significantly but the size of the basis only exhibits mild growth. The number of shifted systems N_a and the number of iterations taken by MPGMRES-Sh with varying mesh discretization is given in Table 1.

Next, we compute the accuracy of the covariance approximation. We follow the method of manufactured solutions and let $f = \cos(2i x_1) \cos(2i x_2)$, so that $u = C_a f = (a^2 + 8i^2)^{-a} \cos(2i x_1) \cos(2i x_2)$. We compute the approximation using the formula in (2.7) computed using the approach in Algorithm 3.1. We then compute the absolute




Fig. 2. Accuracy of the covariance approximation for different values of α and mesh discretization h .

Table 1

Effect of the mesh discretization on the number of MPGMRES-Sh iterations. Also reported are the number of shifted systems to be solved.

Grid size	N_a	Iters
33 s 33	123	9
65 s 65	173	10
129 s 129	235	11
257 s 257	305	16
513 s 513	387	22

error in the $L^2(\Omega)$ norm and plot the error in Figure 2 for different values of α and different values of h , the mesh discretization parameter. We see that for all the values of α , the error decreases with decreasing values of h on the order of h^2 . Furthermore, the absolute error also decreases with increasing values of α . The error was analyzed for $0 < \alpha < 1$ in [9] for the Dirichlet boundary case. In order to apply such an analysis to our setting, this analysis needs to be first extended to the Neumann boundary setting and then to the case $\alpha > 1$. We leave this as part of future work.

Effect of different values of α . In this experiment, we take the grid size to be 257×257 and vary the exponent s from 0.1 to 0.9 in increments of 0.1. The values of the number of shifted systems and the number of iterations taken by MPGMRES-Sh are given in Table 2. We see that the number of iterations is relatively small and it does not depend on the value of s . Furthermore, while the number of shifted systems can be large for each value of s , the number of iterations is relatively small. Since the same basis is used across all the shifted systems, the cost is amortized. Furthermore, the same basis can be used across multiple values of s since the behavior is essentially independent of s . This can be very advantageous in a multiquery setting, where we need to apply the covariance operator to the same vector for multiple values of s .

5.2. Samples from the SPDE. We provide numerical experiments for generating samples from the Gaussian process defined by the SPDE (2.1) with zero Neumann boundary conditions. We take the domain to be $\Omega = (0, 1)^2$ and set the grid size to 129×129 . We solve the SPDE using the approach described in section 3.2 with MPGMRES-Sh used to accelerate the solutions of the shifted linear systems. We use the same preconditioners as in the previous experiment.

Table 2

The number of shifted systems and the number of iterations required by MPGMRES-Sh. Three preconditioners are used corresponding to the shifts $\{10^{-8}, 10^{-4}, 10^{-2}\}$.

s	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
N_a	846	476	364	318	305	318	364	476	846
Iters	16	16	16	16	16	16	16	16	16




Fig. 3. Samples from the SPDE (3.3) with $u \in \{\frac{1}{4}, \frac{2}{4}, \frac{3}{4}\}$ and $a = u + d/2$.

We choose three different values of $u \in \{\frac{1}{4}, \frac{2}{4}, \frac{3}{4}\}$ and take $a^2 = 100$. Note that for these choices of u , we have $1/2 < a/2 < 1$. The samples are visualized in Figure 3; note that we have used the same realization w across all the values of a . It is readily seen that for larger values of a , the sample realizations appear smoother. The number of MPGMRES-Sh iterations for each value of a is the same and is 11.

We also compute samples from the SPDE (3.3) with zero Neumann boundary conditions, where we take $a^2(x) = 100$ and

$$(5.1) \quad a(x) = \begin{bmatrix} \cos(a) & \sin(a) \\ -\sin(a) & \cos(a) \end{bmatrix} \begin{bmatrix} l_1^2 & \\ & l_2^2 \end{bmatrix} \begin{bmatrix} \cos(a) & -\sin(a) \\ \sin(a) & \cos(a) \end{bmatrix}.$$

We take $a = i/4$, $l_1^2 = 10$, and $l_2 = 1$. Note that the corresponding random field is stationary and anisotropic. As before we pick $u \in \{\frac{1}{4}, \frac{2}{4}, \frac{3}{4}\}$ and the grid size is 129×129 . The samples are plotted in Figure 4; note that we have used the same realization w across all the values of a . The number of MPGMRES-Sh iterations taken for each value of a is the same and is 15.

5.3. Computing the truncated KL expansion. We consider the anisotropic covariance operator (2.4) where $a^2 = 80$ and a is defined in (5.1) with $l_1^2 = 4$, $l_2 = 1$, and $a = -i/4$. We compute the truncated KL expansion for three different values of $a \in \{\frac{3}{2}, \frac{5}{2}, \frac{7}{2}\}$. To compute the approximate eigenpairs, we used the two-pass randomized algorithm for GHEP [27, Algorithm 6]. We computed 200 eigenvalues and used an oversampling parameter of 20. The eigenvalues are plotted in the left panel of Figure 5; on the right panel, we plot 6 different samples computed using the truncated KL expansion with $a = 5/2$.

5.4. Tomography examples. We consider a test problem from the IRTTools package [17]. The number of grid points is $n_x \times n_y$. First, we use the PRseismic test problem with the number of sources $\lfloor 0.4n_x \rfloor$ and the number of receivers $\lfloor 0.6n_x \rfloor$. In total, there were 4004 measurements to which we add 2% Gaussian noise. We take the prior distribution to be Gaussian with zero mean and the covariance matrix obtained from the discretization of A^{-a} with $A = a^2 - a$. Next, we consider a test problem




Fig. 4. Samples from the anisotropic SPDE with $u \in \{\frac{1}{4}, \frac{2}{4}, \frac{3}{4}\}$ and $a = u + d/2$.




Fig. 5. (left) Eigenvalues of the KL expansion for $a \in \{\frac{3}{2}, \frac{5}{2}, \frac{7}{2}\}$, (right) 6 samples from the truncated KL expansion with $a = \frac{5}{2}$.

Table 3
Problem settings for the tomography test problems in section 5.4.

Image	Smooth	Cheese
Application	Seismic	X-ray
m	4004	14835
$n_x \times n_y$	128 $\times$ 128	128 $\times$ 128
a^2	80	300

with real data from the Finnish Inverse Problems society [10]. In this instance of the test problem, the number of grid points is 128×128 and the number of measurements is 14835. We use the same prior distribution as before. A summary of the settings for the test problems is given in Table 3. In Figure 6, we plot the results for the seismic test problem. In the top plot, we plot the reconstructions for three different values of $a \in \{1.5, 2.5, 3.5\}$; since $d = 2$, these correspond to $u = 1/2, 3/2, 5/2$. The relative error for each value of a is given in the title of each image. Since the images have smooth features, we see that the error has a slight decrease for increasing values of a . In the bottom plot of the same figure, we plot the iteration history of the relative error computed using the GenHyBR method. We also highlight, in black, the iteration at which the algorithm has stopped. However, we plot the relative error until the maximum number of iterations 100 to show that the relative error has stabilized and no semiconvergence is seen.

In the next experiment, we consider the seismic tomography problem but investigate the impact of an anisotropic covariance operator. We take the covariance operator




Fig. 6. Seismic tomography: (Top) reconstructions for different values of α ; (bottom) iteration history of relative error. Black circles denote the stopping criterion.

(2.4) with $a^2 = 80$ and a defined in (5.1); we take $l_1 = 4, l_2 = 1$, and $a = -i/4$. For the isotropic covariance, we take $a^2 = 80$ and $a = I$. For both covariance operators, we take the value of α to be 2.5. The true image, reconstruction with isotropic covariance and anisotropic covariance operators are displayed in Figure 7. The relative reconstruction errors are given on the titles; for isotropic covariance, GenHyBR took 56 iterations and for the anisotropic covariance, GenHyBR took 125 iterations. It is clearly seen that by choosing the anisotropic covariance as the prior covariance, we can reduce the reconstruction error. This suggests that the anisotropic covariance operators may be beneficial under certain circumstances; however, to pick the correct parameters one has to use expert knowledge or they have to be estimated from data.

Finally, in Figure 8, we plot the result of the reconstructions corresponding to the real data problem using X-ray tomography. Since the true image is not available, we do not plot the relative error history. We choose $\alpha \in \{1.1, 1.5, 1.9\}$ since, for higher values of α , we get poor reconstructions. We have used the value of $a^2 = 300$ and limited the maximum number of iterations to 150; the number of iterations taken was 150, 83, and 75 for $\alpha \in \{1.1, 2.1, 3.1\}$, respectively.

5.5. PDE-based example. In this application, we consider a forward problem that is PDE based. The underlying PDE is a 2D time-dependent heat equation

$$\begin{aligned}
 (5.2) \quad & \frac{\partial u}{\partial t} = a \nabla^2 u, & x \in \Omega, t \in (0, T], \\
 & u|_{\partial \Omega} = 0, & x \in \partial \Omega, t \in (0, T], \\
 & u(t=0) = u_0, & x \in \Omega.
 \end{aligned}$$




Fig. 7. Seismic test problem: (left) true image, (center) reconstruction using isotropic covariance, (right) reconstruction using anisotropic covariance.




Fig. 8. X-ray tomography: Reconstructions for different values of α .

Here $\mathbf{a} = (0, 1)^T$, $T = 0.01$, and $\mathbf{n}$ is the outward normal vector. The forward problem is discretized using Galerkin finite elements and Crank--Nicholson time differences; we use a grid size of 65×65 with 100 time steps. The inverse problem involves reconstructing the initial conditions from a discrete set of measurements at the final time point $t = T$. We add 2% Gaussian noise to simulate measurement error. We use the covariance operator (2.4) with $\mathbf{a}^2 = 80$, $\mathbf{a} = \mathbf{I}$, and $\mathbf{a} = 2.5$.

The reconstructions were performed using GenHyBR as described in section 4.3 and the posterior variance was approximated using the technique in section 4.4. The diagonals of $\mathbf{Q}$ were estimated using the Diag++ method [5, Algorithm 1] with 300 samples. GenHyBR converged in 42 iterations and it produced a relative error of 0.15. However, we used a basis size of $k = 75$ for the subsequent uncertainty estimates. Figure 9 displays the true field and the reconstruction (top row). The white marks indicate the sensor locations at which data are collected. In the bottom row of the same figure, we plot the prior variance (diagonals of $\mathbf{a}^{-2} \mathbf{Q} = \mathbf{C}_{pr} \mathbf{M}^{-1}$), the update $\mathbf{Z}_k \mathbf{a}_k \mathbf{Z}_k^T$, and the approximate posterior variance, i.e., diagonals of $\mathbf{a}_{post} \mathbf{M}^{-1} = \mathbf{a} \mathbf{Q}^{-1} \mathbf{Z}_k \mathbf{a}_k \mathbf{Z}_k^T$. We see that the prior variance is high at the four corners which shows artifacts due to the boundary conditions. The (approximate) posterior variance shows the reduction in the uncertainty, which is especially prominent in and around the sensor coverage.

6. Conclusions and discussion. In this paper, we presented efficient methods for Whittle--Matérn Gaussian priors in the context of Bayesian inverse problems.




Fig. 9. 2D inverse diffusion. Top row: (left) true image, (right) reconstruction. Bottom row: (left) prior variance $\alpha_c^{-2}Q$, (center) diagonals of $-Z_k a_k Z_k^T$, (right) approximate posterior variance.

However, the techniques developed here may be of interest beyond inverse problems in fractional PDEs and Gaussian random fields. Several extensions of our work are possible. First, while we focused the derivation and numerical experiments on zero Neumann boundary conditions, this is not a limitation of the framework, and it is easy to extend the framework to zero Dirichlet and Robin boundary conditions. The latter may be particularly suitable to mitigate the effects of boundary conditions (see [15]). Second, the techniques in this paper may be extended to generalized Matérn fields on compact Riemannian fields; see [24, 23, 28, 22]. Finally, in Bayesian inverse problems, it is straightforward to extend the techniques to nonlinear forward problems within a Newton-based solver for the MAP estimate [31]. Another interesting line of research is to extend these techniques to dynamic inverse problems, in which the parameters of interest change in time and we need to use spatiotemporal priors.

Acknowledgments. A.K.S. would like to thank Daniel Szyld, Daniel Sanz-Alonso, Georg Stadler, and Alen Alexanderian for helpful discussions.

REFERENCES

- [1] H. Antil and S. Bartels, Spectral approximation of fractional PDEs in image processing and phase field modeling, *Comput. Methods Appl. Math.*, 17 (2017), pp. 661–678, <https://doi.org/10.1515/cmam-2017-0039>.
- [2] H. Antil, Z. W. Di, and R. Khatri, Bilevel optimization, deep learning and fractional Laplacian regularization with applications in tomography, *Inverse Problems*, 36 (2020), 064001.
- [3] H. Antil, J. Pfefferer, and S. Rogovs, Fractional operators with inhomogeneous boundary conditions: Analysis, control, and discretization, *Commun. Math. Sci.*, 16 (2018), pp. 1395–1426, <https://doi.org/10.4310/CMS.2018.v16.n5.a11>.

- [4] T. Bakhos, P. K. Kitanidis, S. Ladenheim, A. K. Saibaba, and D. B. Szyld, Multipreconditioned GMRES for shifted systems, *SIAM J. Sci. Comput.*, 39 (2017), pp. S222--S247.
- [5] R. A. Baston and Y. Nakatsukasa, Stochastic Diagonal Estimation: Probabilistic Bounds and an Improved Algorithm, preprint, arXiv:2201.10684, 2022.
- [6] D. Bolin and K. Kirchner, The rational SPDE approach for Gaussian random fields with general smoothness, *J. Comput. Graph. Statist.*, 29 (2020), pp. 274--285.
- [7] D. Bolin, K. Kirchner, and M. Kovács, Weak convergence of Galerkin approximations for fractional elliptic stochastic PDEs with spatial white noise, *BIT*, 58 (2018), pp. 881--906.
- [8] D. Bolin, K. Kirchner, and M. Kovács, Numerical solution of fractional elliptic stochastic PDEs with spatial white noise, *IMA J. Numer. Anal.*, 40 (2020), pp. 1051--1073.
- [9] A. Bonito and J. Pasciak, Numerical approximation of fractional powers of elliptic operators, *Math. Comp.*, 84 (2015), pp. 2083--2110.
- [10] T. A. Bubba, M. Juvonen, J. Lehtonen, M. März, A. Meaney, Z. Purisha, and S. Silta-nen, Tomographic X-ray Data of Carved Cheese, preprint, arXiv:1705.05732, 2017.
- [11] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems part I: The linearized case, with application to global seismic inversion, *SIAM J. Sci. Comput.*, 35 (2013), pp. A2494--A2523.
- [12] J. Chung, J. G. Nagy, and D. P. O'leary, A weighted GCV method for Lanczos hybrid regularization, *Electron Trans. Numer. Anal.*, 28 (2008), pp. 149--167.
- [13] J. Chung and A. K. Saibaba, Generalized hybrid iterative methods for large-scale Bayesian inverse problems, *SIAM J. Sci. Comput.*, 39 (2017), pp. S24--S46.
- [14] J. Chung, A. K. Saibaba, M. Brown, and E. Westman, Efficient generalized Golub-Kahan based methods for dynamic inverse problems, *Inverse Problems*, 34 (2018), 024005.
- [15] Y. Daou and G. Stadler, Mitigating the influence of the boundary on PDE-based covariance operators, *Inverse Probl. Imaging*, 12 (2018), pp. 1083--1102.
- [16] H. P. Flath, L. C. Wilcox, V. Akcelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas, Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations, *SIAM J. Sci. Comput.*, 33 (2011), pp. 407--432.
- [17] S. Gazzola, P. C. Hansen, and J. G. Nagy, IR tools: A MATLAB package of iterative regularization methods and large-scale test problems, *Numer. Algorithms*, 81 (2019), pp. 773--811, <https://doi.org/10.1007/s11075-018-0570-7>.
- [18] L. Herrmann, K. Kirchner, and C. Schwab, Multilevel approximation of Gaussian random fields: Fast simulation, *Math. Models Methods Appl. Sci.*, 30 (2020), pp. 181--223, <https://doi.org/10.1142/s0218202520500050>.
- [19] N. J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, 2008.
- [20] E. Jansson, M. Kovács, and A. Lang, Surface finite element approximation of spherical Whittle--Matérn Gaussian random fields, *SIAM J. Sci. Comput.*, 44 (2022), pp. A825--A842.
- [21] K.-T. Kim, U. Villa, M. Parno, Y. Marzouk, O. Ghattas, and N. Petra, HIPPyLib-MUQ: A Bayesian inference software framework for integration of data with complex predictive models under uncertainty, *ACM Trans. Math. Software*, to appear.
- [22] A. Lang and M. Pereira, Galerkin-Chebyshev Approximation of Gaussian Random Fields on Compact Riemannian Manifolds, preprint, arXiv:2107.02667, 2021.
- [23] F. Lindgren, D. Bolin, and H. Rue, The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running, *Spat. Stat.*, 50 (2022), 100599.
- [24] F. Lindgren, H. Rue, and J. Lindström, 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 (2011), pp. 423--498.
- [25] L. Roininen, J. M. Huttunen, and S. Lasanen, Whittle--Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography, *Inverse Probl. Imaging*, 8 (2014), pp. 561--586.
- [26] A. K. Saibaba, J. Chung, and K. Petroske, Efficient Krylov subspace methods for uncertainty quantification in large Bayesian linear inverse problems, *Numer. Linear Algebra Appl.*, 27 (2020), e2325.
- [27] A. K. Saibaba, J. Lee, and P. K. Kitanidis, Randomized algorithms for generalized Hermitian eigenvalue problems with application to computing Karhunen-Loève expansion, *Numer. Linear Algebra Appl.*, 23 (2016), pp. 314--339.
- [28] D. P. Simpson, *Krylov Subspace Methods for Approximating Functions of Symmetric Positive Definite Matrices with Applications to Applied Statistics and Anomalous Diffusion*, Ph.D. thesis, Queensland University of Technology, Brisbane, Australia, 2008.

- [29] A. Spantini, A. Solonen, T. Cui, J. Martin, L. Tenorio, and Y. Marzouk, Optimal low-rank approximations of Bayesian linear inverse problems, *SIAM J. Sci. Comput.*, 37 (2015), pp. A2451--A2487.
- [30] A. M. Stuart, Inverse problems: A Bayesian perspective, *Acta Numer.*, 19 (2010), pp. 451--559.
- [31] U. Villa, N. Petra, and O. Ghattas, hippylib: An extensible software framework for large-scale inverse problems, *J. Open Source Softw.*, 3 (2018), 940.
- [32] C. J. Weiss, B. G. van Bloemen Waanders, and H. Antil, Fractional operators applied to geophysical electromagnetics, *Geophys. J. Int.*, 220 (2020), pp. 1242--1259.