
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Emzir, Muhammad; Lasanen, Sari; Purisha, Zenith; Roininen, Lassi; Särkkä, Simo
Non-stationary multi-layered Gaussian priors for Bayesian inversion

Published in:
Inverse Problems

DOI:
[10.1088/1361-6420/abc962](https://doi.org/10.1088/1361-6420/abc962)

Published: 03/12/2020

Document Version
Peer-reviewed accepted author manuscript, also known as Final accepted manuscript or Post-print

Please cite the original version:
Emzir, M., Lasanen, S., Purisha, Z., Roininen, L., & Särkkä, S. (2020). Non-stationary multi-layered Gaussian priors for Bayesian inversion. *Inverse Problems*, 37(1), Article 015002. <https://doi.org/10.1088/1361-6420/abc962>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

Non-Stationary Multi-layered Gaussian Priors for Bayesian Inversion

Muhammad Emzir¹⁾, Sari Lasanen²⁾, Zenith Purisha¹⁾, Lassi Roininen³⁾, Simo Särkkä¹⁾

1) Department of Electrical Engineering and Automation, Aalto University, P.O. Box 12200, FI-00076 Aalto, Finland

2) Sodankylä Geophysical Observatory, University of Oulu, P.O. Box 8000, FI-90014 University of Oulu, Finland

3) School of Engineering Science, Lappeenranta-Lahti University of Technology, P.O. Box 20, FI-53851 Lappeenranta, Finland

E-mail: muhammad.emzir@aalto.fi

Nov 2020

Abstract. In this article, we study Bayesian inverse problems with multi-layered Gaussian priors. The aim of the multi-layered hierarchical prior is to provide enough complexity structure to allow for both smoothing and edge-preserving properties at the same time. We first describe the conditionally Gaussian layers in terms of a system of stochastic partial differential equations. We then build the computational inference method using a finite-dimensional Galerkin method. We show that the proposed approximation has a convergence-in-probability property to the solution of the original multi-layered model. We then carry out Bayesian inference using the preconditioned Crank–Nicolson algorithm which is modified to work with multi-layered Gaussian fields. We show via numerical experiments in signal deconvolution and computerized X-ray tomography problems that the proposed method can offer both smoothing and edge preservation at the same time.

1. Introduction

The Bayesian approach provides a consistent framework to obtain solutions of inverse problems. By formulating the unknown as a random variable, the degree of information that is available can be encoded as a statistical prior. The ill-posedness of the problem is mitigated by reformulating the inverse problem as a well-posed extension in the space of probability distributions [1]. Among statistical priors that are commonly used in Bayesian inverse problem is the Gaussian prior which is relatively easy to manipulate, has a simple structure, and also has a close relation with traditional Tikhonov regularization. This approach has also got a growing interest from a machine learning community, where the use of Gaussian prior for Bayesian inference is known as Gaussian process regression [2].

When the unknown is a multivariate function, it is natural to model it as a random field. There is a vast amount of studies on Gaussian random fields and their applications where the random field is assumed to be stationary [3]. Stationary Gaussian fields have uniform spatial behavior. As a result, stationary Gaussian fields fail in the cases where the smoothness of the target varies spatially in unexpected ways [4, 5]. To model variable spatial behaviour, the covariance structure needs to be tuned appropriately. There are a number of proposals to increase flexibility of the non-stationary Gaussian fields. Some of the strategies offered in literature are to construct an anisotropic variant of an isotropic covariance function [6, 4, 7], warping of the inputs to a Gaussian process [8], to use transdimensional Gaussian processes, [9, 10], or to reformulate the fields as stochastic partial differential equations and let the coefficients vary in space [11]. In [12], a similar idea to [11] is used, where instead of a predetermined length-scale function, a random field is used. The Gaussian fields are chosen to have Matérn covariance functions and their length-scales are varied according to another Gaussian field. This approach is recently extended to allow flexibility in measurement noise model and hyperprior parameters [13]. A different approach is used in [14] where the model is formed as a cascaded composition of Gaussian fields (see also [15]).

There are some recent findings analyzing how adding more layers translate to the ability of the overall hierarchical Gaussian field to describe a random field with a complex structure. It has been demonstrated in [15] that as the number of layers increases, the density of the last Gaussian field shrinks to a one-dimensional manifold. This might prevent cascaded Gaussian fields to model phenomena where the underlying dimension is greater than one. Ergodicity and effective depth of a hierarchical Gaussian fields has also been analyzed in [16]. The consequence of their result is that there might only be little benefit from increasing number of layers after reaching a certain number.

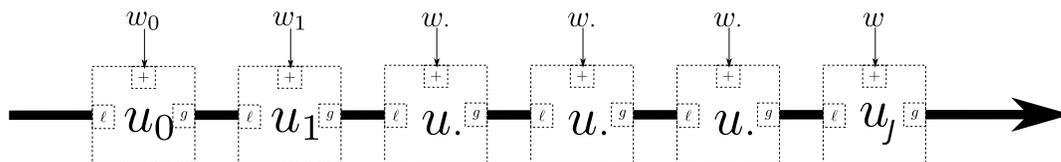


Figure 1. Illustration of a chain of Gaussian fields. Each node (field) has an independent white noise input field and a length-scale parameter ℓ . The length-scale ℓ is obtained as a function evaluated on the fields above it. These Gaussian fields are approximated in finite-dimensional Hilbert space using $\mathbf{L}(\mathbf{u}_{j-1})\mathbf{u}_j = \mathbf{w}_j$, see (4).

In many spatio-temporal inverse problems, it is often useful to start by working in an infinite-dimensional space. There are essentially two approaches to construct a Bayesian inference algorithm. The first approach is to discretize the forward map (e.g., via grid partitions or finite element meshes) and apply a Bayesian inference method to the finite-dimensional setting [1]. The second approach [17, 18] is to directly apply Bayesian inference to the infinite-dimensional problem, and afterward, apply a discretization method. The latter approach is possible through realizing that the posterior and prior

probability distribution can be related via the Radon–Nikodym derivative which can be generalized to function spaces [17].

For both of the aforementioned approaches, the sampling technique has to be carefully designed. The traditional Markov chain Monte Carlo (MCMC) algorithms suffer from slow mixing times upon the grid partitioning refinement [19]. These methods dictate to reduce the MCMC step size which becomes computationally expensive. There are several MCMC algorithms designed specifically to deal with infinite-dimensional problems so that mixing time will, up to some extent, be almost independent of the dimensionality [19, 20, 21, 22]. Among these algorithms, the preconditioned Crank–Nicolson (pCN) [19] is a very simple one to be implemented. The main idea of this algorithm is to design a random walk such that the discretization of this random walk is invariant for the target measure which can be designed to be the posterior measure. Assuming that the target measure has a density with respect to a Gaussian reference measure, the pCN algorithm takes advantage of a clever selection of the Markov transition kernel. It has been shown in [23] via spectral gap analysis that pCN has dimension-independent sampling efficiency. This benefit comes in contrast to the standard random walk proposal where the probability of the proposal acceptance will be almost zero in infinite-dimensional case. Recently, a non-centered version of this algorithm was introduced in [24]. Their work was developed using a non-centered reparametrization developed in [25]. This transformation is important in the hierarchical prior Bayesian inversion cases since it breaks the dependency between parameters in different levels which simplifies the calculation of posterior. There is also a generalization of the pCN algorithm to take into account the information of the measure, and an adaptive version of it [22, 26].

The main contribution of this article is to present a method for Bayesian inverse problems with multi-layered Gaussian field prior models via a Galerkin method. The motivation is to have enough complexity in the model to allow for both smoothing and edge preserving properties while keeping relatively low number of layers at the same time. By this approach, we essentially formulate a hierarchical prior so that the regularization becomes less subjective, while keeping the number of total parameters low [27]. In particular, we follow the approach recently described in [12], where the Gaussian fields are represented as stochastic partial differential equations (SPDEs), in which the length-scale parameters depend on the solution of SPDEs for the layer above. However, instead of using a grid partition, we will apply Galerkin method, on which we have already obtained preliminary results in [28]. Using this approach, we can avoid evaluating SPDE forward problem via finite difference equations, and the number of parameters is greatly reduced. This can be considered as a compromise between the accuracy and the computational complexity. Our approach is also related to the one described in [29], where Gaussian fields with stationary covariance functions are approximated in finite-dimensional Hilbert spaces. Such approach relies on the fact that the covariance function of a stationary Gaussian field is expandable via Mercer’s theorem.

In this work, we construct several Gaussian fields in a hierarchical structure. Using this structure, we show how to transform this model into a chain of Gaussian fields driven by white noise fields. This formulation enables us to use the Bayesian inference framework described in [16]. We then implement the proposed approximation into an MCMC algorithm. Our MCMC algorithm is based on the non-centered version of the preconditioned Crank–Nicolson algorithm modified to work with multi-layered Gaussian fields [16, 24]. An alternative approach to MCMC would be to use the variational Bayesian approach which approximates the posterior with a member of tractable distribution family [30]. The variational Bayesian approach could be a few orders faster than MCMC [31, 30]. In addition, there are some results that show the asymptotical limit of the variational Bayesian posterior [32]. However, it is not clear whether the variational Bayesian approach will be well defined for the infinite-dimensional case [33, 34, 35]. Nevertheless, the variational Bayesian approach has given promising results in certain high-dimensional inverse problems [36].

Since translating the inverse problem into finite-dimensional formulation will introduce a discretization problem that could exacerbate the reconstruction errors [37, 38], we rigorously show that the proposed approximation enjoys a nice convergence-in-probability property to the weak solution of the forward model. To show the convergence result, we start by establishing an upper bound for the square root of the precision operator of the Gaussian field. From there, we develop another upper bound for the error between this operator and its finite-dimensional approximation. We then show that the original weak SPDE solution satisfies a Hölder continuity property. Using these results and additional tightness conditions, we finally show that the proposed approximation converges in probability to the original weak solution of the forward model. As a consequence, we can guarantee that the approximated prior and posterior probability distribution converges weakly to the original prior and posterior, respectively.

In this article, we also present an application of the proposed method to a computerized X-ray tomography problem. Computerized tomography problems are very challenging since they are ill-posed [39]. One way is to compute maximum a posteriori estimate for a general X-ray tomography problem using a Gaussian prior in finite-dimensional setting [40]. In [41], a Bayesian method using Gaussian fields with non-stationary covariance functions described in [4, 42] is developed for plasma fusion and soft X-ray tomography. Recently in [43], a Hilbert space approximation technique described in [29] is used in a sparse tomographic inverse problem. Bayesian tomographic reconstruction with non-Gaussian prior has been studied in [44, 45]. There are also some recent results in X-ray tomography using deep learning methods. However, contrary to the Bayesian and regularization approaches, these methods are prone to instabilities when exposed to a small perturbation and structural changes [46].

Previously, we have presented a subset of our contributions in [28]. In the present work, we extend the methods presented in [28] to Bayesian inverse problems and we have also added a throughout convergence analysis of the methods. The algorithm presented in this work also generalizes the algorithm presented in [28] to the case of multiple

hyperprior layers.

The outline of the article is the following. In Section 2, we present a formulation of Bayesian inverse problems using multi-layered Gaussian priors via Galerkin method. The convergence analysis is presented in Section 3. In Section 4, we propose a Markov chain Monte Carlo algorithm to sample the Fourier coefficients from their posterior distribution. In Section 5, we present an application of the proposed method to a one-dimensional example model and to a tomographic inverse problem. Finally, Section 6 concludes the article.

1.1. Notation

Let $\{\phi_l\}$ be a basis formed from the orthogonal eigenfunctions of the Laplace operator with respect to some domain Ω with suitable boundary conditions. Let N be the number of basis functions $\{\phi_l\}$ used in the Galerkin method and let H_N denote their span. For multi-layer Gaussian field priors, we use u_j to denote the random field at layer j , where the total number of layer is $J + 1$, that is, J is the number of hyperprior layers. The collection of $J + 1$ random fields (u_0, u_1, \dots, u_J) is denoted by u . The Fourier transform of any random field z is denoted by \widehat{z} , where the Fourier coefficient for index l is given by $\widehat{z}(l)$, that is, $\widehat{z}(l) = \langle z, \phi_l \rangle$, where $\langle \cdot, \cdot \rangle$ is the standard L^2 inner product on the domain Ω . We use bold characters to denote vectors or matrices with elements in \mathbb{C} or \mathbb{R} . Fourier coefficient for random field u_j for index $-N$ to N is given by \mathbf{u}_j , that is, $\mathbf{u}_j = (\widehat{u}_j(-N) \dots \widehat{u}_j(N))$. We denote $J + 1$ collections of the Fourier coefficients of random fields u as $\mathbf{u} = (\mathbf{u}_0, \dots, \mathbf{u}_J)$. The identity operator is denoted with I .

2. Finite-dimensional approximations

Consider a Bayesian inverse problem on a Gaussian field where the unknown is a real-valued random field $v(\mathbf{x}) : \Omega \rightarrow \mathbb{R}$ on a bounded domain $\Omega \subset \mathbb{R}^d$. We assume in the inverse problem that v belongs to a Hilbert space H , specifically, $v \in H := L^2(\Omega)$. To carry out the Bayesian inference, a set of measurements is taken (either direct or indirect in multiple locations). In this article, the measurement is assumed to be a linear operation on v corrupted with additive noises, that is, $y_k = \langle v, h_k \rangle + e_k$, where h_k is an element in H represents a real linear functional on H , and e_k is a zero-mean white noise with a covariance matrix \mathbf{E} .

In the Bayesian framework, the estimation problem is equivalent to exploring the posterior distribution of v given the measurements $\{y_k\}$. The Bayesian inversion approach for this problem starts with assuming that v is a Gaussian field with a certain mean (assumed zero for simplicity) and covariance function $C(\mathbf{x}, \mathbf{x}')$. In the case when $C(\mathbf{x}, \mathbf{x}')$ is a Matérn covariance function, the Gaussian field v can be generated from a stochastic partial differential equation of the form [11, 12]

$$(1 - \ell^2 \Delta)^{\alpha/2} v(\mathbf{x}) = \sqrt{\beta \ell^d} w(\mathbf{x}), \quad (1)$$

where $\alpha = \nu + d/2$, d is the dimension of the space, ν is a smoothness parameter, $w(\mathbf{x})$ is a white noise on \mathbb{R}^d , ℓ is the length-scale constant of the Matérn covariance function C , and $\beta = \sigma^2 2^d \pi^{d/2} \Gamma(\alpha) / \Gamma(\nu)$ with σ^2 being a scale parameter.

To obtain a non-stationary field, we modify SPDE (1) so that the length-scale ℓ is modeled via another Gaussian field u with Matérn covariance function. Namely, we select $\ell(\mathbf{x}) = g(u(\mathbf{x}))$, where g is a smooth positive function $g : \mathbb{R} \rightarrow \mathbb{R}_+$. As in [12], we also require that ℓ should satisfy $\sup_{\mathbf{x} \in \Omega} \ell(\mathbf{x}) < \infty$ and $\inf_{\mathbf{x} \in \Omega} \ell(\mathbf{x}) > 0$ with probability one. For notational and mathematical convenience we restrict $\alpha = 2$. The results of this article could also be extended to other cases. Introducing the spatially varying length-scale $\ell(\mathbf{x}) = g(u(\mathbf{x}))$ into (1), and since the length-scale is always greater than zero, with probability one, using $\kappa = 1/\ell$ we obtain the following SPDE

$$(\kappa(u(\mathbf{x}))^2 - \Delta) v(\mathbf{x}) = \sqrt{\beta} \kappa(u(\mathbf{x}))^\nu w(\mathbf{x}). \quad (2)$$

In order to facilitate easy operation with the Laplace operator, we choose to expand v using the eigenfunctions of the Laplacian. With a suitable boundary condition, the Laplace operator can be expressed $-\Delta v = \sum_{j=-\infty}^{\infty} \lambda_j \langle v, \phi_j \rangle \phi_j$, where ϕ_j is a complete set of orthonormal eigenfunctions of Δ and $\lambda_j > 0$, where $\lim_{j \rightarrow \infty} \lambda_j = \infty$ [47]. Observe that, in the sense of (2), $\kappa(u(\mathbf{x})) := 1/g(u(\mathbf{x})) \in L^\infty(\Omega)$ is a multiplication operator acting pointwise, that is, $(\kappa(u)v)(\mathbf{x}) = \kappa(u(\mathbf{x}))v(\mathbf{x}), \forall \mathbf{x} \in \Omega$ [48].

In what follows, we will first describe the matrix representation of a chain of Gaussian fields generated from SPDEs in the form of (2) for d -dimensional domain. Then in Section 3, we will develop a convergence result of the Galerkin method developed here to the weak solution of (2).

2.1. Matrix representation

Let us examine a periodic boundary condition on d -dimensional box Ω with side length 1. Within this boundary condition, it is useful to consider H as a complex Hilbert space, so that we can set $\phi_l = \exp(i c_d \mathbf{x}^\top \mathbf{k}(l))$ as Fourier complex basis functions for d dimensions, for some constant c_d and a multi-index $\mathbf{k}(l)$ which is unique for every l . Let the finite-dimensional Hilbert subspace H_N of H be the span of $\phi_{-N}, \dots, \phi_0, \dots, \phi_N$. In what follows, we will explicitly construct the multi-index $\mathbf{k}(\cdot)$.

Without losing generality, let us assume that every entry in $\mathbf{k}(l)$ is between $-n$ to n . For any $-N \leq l \leq N$, where $N = \frac{(2n+1)^d - 1}{2}$, we would like to construct $\mathbf{k}(l) = (k_1(l), \dots, k_d(l))$ such that it is unique for each $-N \leq l \leq N$ and $\mathbf{k}(l+m) = \mathbf{k}(l) + \mathbf{k}(m)$, $-N \leq l, m \leq N$ and $\max(|k_r(l+m)|) < n, r \leq d$. The construction of $\mathbf{k} : [-N, N] \rightarrow [-n, n]^d$ is as follows. Let the matrix $\mathbf{K}(n)$ be given as

$$\mathbf{K}(n) = \begin{pmatrix} \mathbf{z}_n \otimes \mathbf{e}_n^{\otimes d-1} \\ \mathbf{e}_n \otimes z_n \otimes \mathbf{e}_n^{\otimes d-2} \\ \vdots \\ \mathbf{e}_n^{\otimes d-1} \otimes \mathbf{z}_n \end{pmatrix},$$

where $\mathbf{z}_n = (-n, -n + 1, \dots, n - 1, n)$, $\mathbf{e}_n = (1, 1, \dots, 1, 1)$, and $\mathbf{e}_n^{\otimes d}$ is a Kronecker product of \mathbf{e}_n repeated for d times. The multi index $\mathbf{k}(l)$ is given by selecting $l + (N + 1)$ -th column of \mathbf{K} . The linear relation is defined only if $-N \leq l + m \leq N$ and every element of the summation $\mathbf{k}(l) + \mathbf{k}(m)$ has values in $[-n, n]$. As an example let $d = 2$ and $n = 1$. This gives us:

$$\mathbf{K}(1) = \begin{pmatrix} -1 & -1 & -1 & 0 & 0 & 0 & 1 & 1 & 1 \\ -1 & 0 & 1 & -1 & 0 & 1 & -1 & 0 & 1 \end{pmatrix}.$$

It can be verified that $\mathbf{k}(l)$ selected this way is both unique and linear, given that the summation result is inside the range. For example, $\mathbf{k}(1) + \mathbf{k}(2) = \mathbf{k}(3)$. However, $\mathbf{k}(1) + \mathbf{k}(1)$ is not defined since the summation is outside the limit. In the following, if the context is clear, we will also use the multiple index $\mathbf{k}(l)$ for Fourier component of u , that is $\hat{u}(\mathbf{k}(l)) := \hat{u}(l)$.

Since we have no information outside of the frequencies of interest, under the periodic boundary condition, for $-N \leq l, m \leq N$, $\langle \phi_m u, \phi_l \rangle = \hat{u}(l - m)$, when $\max(|\mathbf{k}(l) - \mathbf{k}(m)|) \leq n$, and zero elsewhere.

Let us denote with \mathbf{u} , \mathbf{v} , and \mathbf{w} finite-dimensional representations of u , v , and w , respectively. Let $M_N(u)$ be the matrix representation of the multiplication operator $u(\mathbf{x})$ on H_N . For a random field r with Fourier coefficients $\mathbf{r} = (\hat{r}(-m) \cdots \hat{r}(m))^\top \in \mathbb{C}^{2m+1}$, let us write a Toeplitz matrix $T \in \mathbb{C}^{(m+1) \times (m+1)}$ with elements from \mathbf{r} as follows:

$$T(\mathbf{r}) = \begin{pmatrix} \hat{r}(0) & \cdots & \hat{r}(-m) \\ \vdots & \ddots & \vdots \\ \hat{r}(m) & \cdots & \hat{r}(0) \end{pmatrix}.$$

The previous discussion allows us to write for $d = 1$, $M_N(u) := T(\tilde{\mathbf{u}}) \in \mathbb{C}^{(2N+1) \times (2N+1)}$, where $\tilde{\mathbf{u}} = (\mathbf{0}_{1 \times N}, \mathbf{u}^\top, \mathbf{0}_{1 \times N})^\top \in \mathbb{C}^{(4N+1) \times 1}$. It is also possible to construct $M_N(u)$ for $d > 1$. However, instead of working directly on \mathbf{u} , we need to work on the frequency indices. Let $\mathbf{J} \in \mathbb{R}^{(2n+1) \times (2n+1)}$ be a square matrix where all of its entry equal to one. Let also $\mathbf{Z}^{(1)} = T(\mathbf{z}_{2n}) \otimes \mathbf{J}^{\otimes d-1}$, $\mathbf{Z}^{(2)} = \mathbf{J} \otimes T(\mathbf{z}_{2n}) \otimes \mathbf{J}^{\otimes d-2}$, \dots , $\mathbf{Z}^{(d)} = \mathbf{J}^{\otimes d-1} \otimes T(\mathbf{z}_{2n})$, respectively. Using these matrices, the (l, m) -th entry of $M_N(u)$ is given by

$$M_N(u)_{l,m} = \hat{u}(\tilde{\mathbf{k}}(l, m)), \quad (3)$$

where $\tilde{\mathbf{k}}(l, m) = (\mathbf{Z}_{l,m}^{(1)}, \dots, \mathbf{Z}_{l,m}^{(d)})$. In this equation, we assign $\hat{u}(\tilde{\mathbf{k}}(l, m)) = 0$ when $\max(|\tilde{\mathbf{k}}(l, m)|) > n$.

The sparsity of M_N as n approaches infinity is $(\frac{3}{4})^d$. The weak solution to (2) in the span of H_N is equivalent to the following equation,

$$\mathbf{L}(u)\mathbf{v} = \mathbf{w}, \quad (4)$$

where $\mathbf{L}(u) := \frac{1}{\sqrt{\beta}}(M_N(\kappa(u)^{d/2}) - M_N(\kappa(u)^{-\nu})\mathbf{D})$ is the square root of the precision operator corresponds to v in matrix form, \mathbf{D} is a diagonal matrix, and \mathbf{v} , \mathbf{w} are complex

vectors with appropriate dimensions. The diagonal entries of \mathbf{D} are given by $\mathbf{D}_{i,i} = -\lambda_i$. Upon computing $M_N(\kappa(u)^\gamma)$, we approximate u by its projection onto H_N if $u \notin H_N$. An important numerical issue to note is that we cannot use an approximation of $M_N(\kappa(u^N)^\gamma)$ obtained by spectral decomposition, that is, $M_N(\kappa(u^N)^\gamma) \approx \mathbf{U}^\top \kappa(\mathbf{D}_u)^\gamma \mathbf{U}$, for a diagonal matrix \mathbf{D}_u with the diagonal entries being the eigenvalues of $M_N(u_N)$, and where \mathbf{U} is an orthonormal matrix. The reason is that the resulting matrix will not be in the form of (3). Instead, we can obtain $M_N(\kappa(u^N)^\gamma)$ matrix by applying Fourier transform directly to $\kappa(u^N)^\gamma$ and making use of (3). Writing (4) as $\mathbf{v} = \mathbf{L}(u)^{-1} \mathbf{w}$, we obtain a composition of a Gaussian field from a unit Gaussian field given in [16].

In what follows, for simplicity, with a slight abuse of notation for $\mathbf{L}(u)$, if $r \in H_N$ we also use $\mathbf{L}(\mathbf{r}) := \mathbf{L}(\sum_{l=-N}^N \widehat{r}(l) \phi_l) = \mathbf{L}(r)$. Using this notation, the J Gaussian field hyperpriors with zero mean assumption can be written in the following form:

$$\mathbf{L}(\mathbf{u}_{j-1}) \mathbf{u}_j = \mathbf{w}_j. \quad (5)$$

The hyperprior layers consist of the random fields u_0, \dots, u_{J-1} , where the random field u_0 will be stationary. Within this multi-layer hyperprior setting, the unknown field v is equivalent to u_J . With the assumption that each random field involved is real-valued, the number of element in \mathbf{u}_i is only $N + 1$ since the remaining element can be obtained by complex conjugation.

3. Convergence analysis

In this section, we will show that the solution of the original SPDE system can be approximated with Galerkin methods on d -dimensional torus \mathbb{T}^d . In our convergence analysis, we will restrict to $d \leq 3$, since we will rely on continuity of the sample paths (see Lemma 3.3). We start by carefully presenting the notation of Galerkin method for analysis purposes. We denote the constant κ_0 with $\exp(u_{-1})$ for notational ease. The function κ is taken to be a smooth function, which is bounded from below and above by exponential functions. That is, $c_1 \exp(-a_1|x|) \leq \kappa(x), \kappa'(x) \leq c_2 \exp(a_2x)$ for some $c_1, c_2, a_1, a_2 > 0$ for all $x \in \mathbb{R}$.

Almost surely bounded functions u_0, \dots, u_J are a weak solution of SPDE system

$$-\Delta u_i + \kappa^2(u_{i-1})u_i = \beta_i^{1/2} \kappa^\nu(u_{i-1})w_i, \quad \text{where } i = 0, \dots, J, \quad (6)$$

if they satisfy

$$-\langle u_i, \Delta \phi \rangle + \langle \kappa^2(u_{i-1})u_i, \phi \rangle = \beta_i^{1/2} \sum_{p=-\infty}^{\infty} \widehat{w}_i(p) \langle \kappa^\nu(u_{i-1})\phi_p, \phi \rangle, \quad \text{where } i = 0, \dots, J,$$

for all $\phi \in C^2(\mathbb{T}^d)$. Assuming that the functions u_i are almost surely bounded guarantees that the inner product of u_i and $\Delta \phi$ is well-defined as compared to the inner product of ∇u_i and $\nabla \phi$, which may not be well-defined. Here we expressed independent white

noises w_i on \mathbb{T}^d with the help of an orthonormal basis $\{\phi_p\}$ in $L^2(\mathbb{T}^d)$ as random series $w_i = \sum_{p=-\infty}^{\infty} \widehat{w}_i(p) \phi_p$, where the random coefficients $\widehat{w}_i(p) \sim N(0, 1)$ are independent.

In the first approximation step for the SPDE system (6), we will approximate white noise w_i with its projections $w_i^N(\mathbf{x})$ onto the subspace H_N . That is, $w_i^N(\mathbf{x}) = \sum_{p=-N}^N w_i \phi_p(\mathbf{x})$. Then the Galerkin approximations u_i^N of u_i satisfy the system

$$\langle \nabla u_i^N, \nabla \phi \rangle + \langle \kappa^2(u_{i-1}^N) u_i^N, \phi \rangle = \beta_i^{1/2} \langle \kappa^\nu(u_{i-1}^N) w_i^N, \phi \rangle, \quad i = 0, \dots, J \quad (7)$$

for all $\phi \in H_N$, where $u_{-1}^N(\mathbf{x}) := \ln(\kappa_0)$. Here we are allowed to use to use inner products of ∇u_i^N and $\nabla \phi$, since the approximated white noise w_i^N belongs to H_N .

We will denote with $L(u_{i-1})^{-1}$ the solution operator, which maps f from the negatively indexed Sobolev space $H^{-1}(\mathbb{T}^d)$ to the weak solution of $-\Delta u + \kappa^2(u_{i-1})u = \beta_i^{1/2} \kappa^\nu(u_{i-1})f$. Similarly, we will denote with $L_N(u_{i-1}^N)^{-1}$ the solution operator, which maps $f \in H^{-1}(\mathbb{T}^d)$ to the Galerkin approximation $u^N \in H_N$ of the equation $-\Delta u + \kappa^2(u_{i-1}^N)u = \beta_i^{1/2} \kappa^\nu(u_{i-1}^N)f$. The matrix form of $L_N(u_{i-1}^N)^{-1}$ is given by $\mathbf{L}(\mathbf{u}_{i-1})^{-1}$ from Equation (5). The solution operators $L(u_{i-1})^{-1}$ and $L_N(u_{i-1}^N)^{-1}$ satisfy the following elementary norm estimates. For simplicity, we will take $\beta_i = 1$ from now on.

Lemma 3.1. *Let u_{i-1} and u_{i-1}^N be bounded functions and let κ be a positive continuous function. The mappings $L(u_{i-1}) : L^2(\mathbb{T}^d) \rightarrow H^2(\mathbb{T}^d)$ and $L_N(u_{i-1}^N) : L^2(\mathbb{T}^d) \rightarrow H^2(\mathbb{T}^d)$ satisfy norm estimates*

$$\begin{aligned} \|L(u_{i-1})^{-1}\|_{L^2, H^2} &\leq C \|\kappa(u_{i-1})\|_\infty^\nu \frac{\max(1, \|\kappa^2(u_{i-1})\|_\infty)}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))^2} \quad \text{and} \\ \|L_N(u_{i-1}^N)^{-1}\|_{L^2, H^2} &\leq C \|\kappa(u_{i-1}^N)\|_\infty^\nu \frac{\max(1, \|\kappa^2(u_{i-1}^N)\|_\infty)}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}^N(\mathbf{x})))^2}, \end{aligned}$$

respectively.

Proof. By the Lax–Milgram theorem [49], $\|(-\Delta + \kappa^2(u_{i-1})I)^{-1}\|_{H^{-1}, H^1} \leq C / \min(1, \inf \kappa^2(u_{i-1}))$ and the multiplication operator has norm $\|\kappa^\nu(u_{i-1})\|_{L^2, L^2} \leq \|\kappa(u_{i-1})\|_\infty^\nu$. Hence, the solution operator has norm

$$\|L(u_{i-1})^{-1}\|_{L^2, H^1} \leq C \frac{\|\kappa(u_{i-1}(\mathbf{x}))\|_\infty^\nu}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))}.$$

We rewrite the PDE in the form

$$(-\Delta + a)u_i^f = (a - \kappa^2(u_{i-1})u_i^f + \kappa^\nu(u_{i-1})f),$$

where the right-hand side belongs now to $L^2(\mathbb{T}^d)$. By inverting the operator $-\Delta + aI$, we obtain an equation for the solution u_i^f , which leads to the norm estimate

$$\begin{aligned} \|u_i^f\|_{H^2} &\leq \frac{2}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))} \left(\frac{C \|\kappa^2(u_{i-1})\|_\infty \|\kappa(u_{i-1})\|_\infty^\nu}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))} + \|\kappa(u_{i-1})\|_\infty^\nu \right) \|f\|_{L^2} \\ &\leq C \|\kappa(u_{i-1})\|_\infty^\nu \frac{\max(1, \|\kappa^2(u_{i-1})\|_\infty)}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))^2} \|f\|_{L^2} \end{aligned}$$

after choosing $a = \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x}))/2$. Similar procedure leads to the desired estimate for $L_N(u_{i-1})$. \square

We will later need the following technical lemma to establish convergence.

Lemma 3.2. *Let u_{i-1} and u_{i-1}^N be bounded functions and let κ be a continuously differentiable positive function. The mappings $L_N(u_{i-1}^N)^{-1}$ and $L(u_{i-1})^{-1}$ satisfy*

$$\|L_N(u_{i-1}^N)^{-1} - L(u_{i-1})^{-1}\|_{L^2, L^2} \leq \frac{1}{N} G_1(u_{i-1}) + G_2(u_{i-1}, u_{i-1}^N) \|u_{i-1}^N - u_{i-1}\|_{L^1}^{1/6},$$

where

$$G_1(u_{i-1}) = C \frac{\max(1, \|\kappa^2(u_{i-1})\|_\infty)^2 \max(\|\kappa^\nu(u_{i-1})\|_\infty, \|\kappa^{-\nu}(u_{i-1})\|_\infty)}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))^3}$$

$$G_2(u_{i-1}, u_{i-1}^N) = C \frac{\max(1, \|\kappa^\nu(u_{i-1})\|_\infty) \max(\|\kappa(u_{i-1}^N)\|_\infty, \|\kappa(u_{i-1})\|_\infty)^{5/3+5\nu/6}}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}^N(\mathbf{x}))) \min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))}$$

$$\times \max_{t \in B} (\kappa'(t))^{1/6}$$

and the set B is the interval $[\min(\inf_{\mathbf{x}} u^N(\mathbf{x}), \inf_{\mathbf{x}} u(\mathbf{x})), \max(\|u^N\|_\infty, \|u\|_\infty)]$.

Proof. We partition $T := L_N(u_{i-1}^N)^{-1} - L(u_{i-1})^{-1}$ into two parts

$$T = (L_N(u_{i-1}^N)^{-1} - L_N(u_{i-1})^{-1}) + (L_N(u_{i-1})^{-1} - L(u_{i-1})^{-1}) = T_1 + T_2.$$

By Cea's lemma [49], the term T_2 has an upper bound

$$\|T_2\|_{L^2, L^2} \leq C \frac{\max(1, \|\kappa^2(u_{i-1})\|_\infty)}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))} \|(I - \tilde{P}_N)L(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1})\|_{L^2, H^1} \|\kappa^\nu(u_{i-1})\|_\infty,$$

where \tilde{P}_N is the orthogonal projection onto H_N in H^1 and I is the identity operator. Let $f \in L^2$ and denote $g := L(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1})\tilde{P}_N f$. Then

$$\begin{aligned} \|(I - \tilde{P}_N)g\|_{H^1}^2 &= \sum_{|k_1|, |k_2| > n} (1 + k_1^2 + k_2^2) |\hat{g}_{k_1, k_2}|^2 = \sum_{|k_1|, |k_2| > n} (1 + k_1^2 + k_2^2)^{-1} |(-\Delta + 1)g|_{k_1, k_2}^2 \\ &\leq \frac{1}{n^2} \|(-\Delta + 1)g\|_{L^2}^2 \leq \frac{C}{N^2} \|L(u_{i-1})^{-1}\|_{L^2, H^2}^2 \|\kappa(u_{i-1}(\mathbf{x}))^{-\nu}\|_\infty^2 \|f\|_{L^2}^2, \end{aligned}$$

where $-\Delta + 1 : H^2 \rightarrow L^2$ is continuous, g^\wedge denotes the Fourier transform and $N = (2n + 1)d$. Hence,

$$\|T_2\|_{L^2, L^2} \leq \frac{C}{N} \frac{\max(1, \|\kappa^2(u_{i-1})\|_\infty)^2}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))^3} \max(\|\kappa^\nu(u_{i-1})\|_\infty, \|\kappa^{-\nu}(u_{i-1})\|_\infty).$$

In the term T_1 , we need to tackle the difference of κ -terms. We aim to use L^p -estimates in order to later allow induction with respect to different layers. We partition T_1 into simpler terms

$$\begin{aligned} \|T_1\|_{L^2, L^2} &\leq \|L_N(u_{i-1}^N)^{-1}\kappa^{-\nu}(u_{i-1}^N) (\kappa^\nu(u_{i-1}^N) - \kappa^\nu(u_{i-1}))\|_{L^2, L^2} \\ &\quad + \|(L_N(u_{i-1}^N)^{-1}\kappa^{-\nu}(u_{i-1}^N) - L_N(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1})) \kappa^\nu(u_{i-1})\|_{L^2, L^2} \\ &= \|T_{11}\| + \|T_{12}\|. \end{aligned}$$

In the term T_{12} , we apply the resolvent identity

$$\begin{aligned} L_N(u_{i-1}^N)^{-1}\kappa^{-\nu}(u_{i-1}^N) &= L_N(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1}) \\ &+ L_N(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1})(\kappa^2(u_{i-1}) - \kappa^2(u_{i-1}^N))L_N(u_{i-1}^N)^{-1}\kappa^{-\nu}(u_{i-1}^N), \end{aligned}$$

which leads to

$$\begin{aligned} \|T_{12}\|_{L^2, L^2} &\leq \|L_N(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1})\|_{L^{3/2}, L^2} \|\kappa^2(u_{i-1}) - \kappa^2(u_{i-1}^N)\|_{L^2, L^{3/2}} \\ &\times \|L_N(u_{i-1}^N)^{-1}\kappa^{-\nu}(u_{i-1}^N)\|_{H^{-1}, H^1} \|\kappa^\nu(u_{i-1})\|_\infty. \end{aligned}$$

The space $L^{3/2}$ embeds continuously into H^{-1} , and $\|A\|_{L^{3/2}, L^2} \leq \|A\|_{H^{-1}, H^1}$ for any operator A . Hence by the Lax–Milgram theorem,

$$\|L_N(u_{i-1})^{-1}\kappa^{-\nu}(u_{i-1})\|_{L^{3/2}, L^2} \leq \frac{1}{\min(1, \inf_{\mathbf{x}} \kappa^2(u_{i-1}(\mathbf{x})))}.$$

We are now treating $\kappa^2(u_{i-1}) - \kappa^2(u_{i-1}^N)$ as a multiplication operator from L^2 to $L^{3/2}$. That is, a function g defines a multiplication operator, which takes a function f to the product of functions g and f . By Hölder's inequality, any multiplication operator $g : L^2 \rightarrow L^{3/2}$ has norm $\|g\|_{L^6}$. Similar treatment for the term T_{11} gives

$$\|T_{11}\|_{L^2, L^2} \leq \|L_N(u_{i-1}^N)^{-1}\kappa^{-\nu}(u_{i-1}^N)\|_{L^{3/2}, L^2} \|\kappa^\nu(u_{i-1}^N) - \kappa^\nu(u_{i-1})\|_{L^2, L^{3/2}}.$$

The difference of κ terms in the estimates for T_{11} and T_{12} reduces to the difference of the functions u_{i-1} through series of elementary estimates

$$\begin{aligned} \|\kappa^a(u_{i-1}^N) - \kappa^a(u_{i-1})\|_{L^6}^6 &= \int_{\mathbb{T}^d} |\kappa^a(u_{i-1}^N(\mathbf{x})) - \kappa^a(u_{i-1}(\mathbf{x}))|^6 d\mathbf{x} \\ &\leq (2 \max(\|\kappa^a(u_{i-1}^N)\|_\infty, \|\kappa^a(u_{i-1})\|_\infty))^5 \int \left| \int_{u_{i-1}(\mathbf{x})}^{u_{i-1}^N(\mathbf{x})} \kappa'(t) dt \right| d\mathbf{x} \\ &\leq C \max(\|\kappa^a(u_{i-1}^N)\|_\infty^5, \|\kappa^a(u_{i-1})\|_\infty^5) \max_{t \in B} |\kappa'(t)| \|u_{i-1}^N - u_{i-1}\|_{L^1}, \end{aligned}$$

where $a = \nu, 2$ and the set B is the interval $[\min(\inf_{\mathbf{x}} u^N(\mathbf{x}), \inf_{\mathbf{x}} u(\mathbf{x})), \max(\|u^N\|_\infty, \|u\|_\infty)]$. \square

Remark 1. When κ is a continuously differentiable function with bounds $c_1 \exp(-a_1|t|) \leq \kappa(t), |\kappa'(t)| \leq c_2 \exp(a_2|t|)$, where $c_1, c_2, a_1, a_2 > 0$, the functions G_1 and G_2 can be taken to be

$$\begin{aligned} G_1(u_{i-1}) &= C_1 \exp(C_2(\|u_{i-1}\|_\infty)) \\ G_2(u_{i-1}, u_{i-1}^N) &= C_3 \exp(C_4(\|u_{i-1}\|_\infty + \|u_{i-1}^N\|_\infty)). \end{aligned}$$

The next lemma shows that the SPDE system (6) forces the layers to be Hölder-continuous.

Lemma 3.3. *Let κ be a positive continuous function. Let $0 < \alpha < 1$ for $d = 1, 2$ and $0 < \alpha < 1/2$ for $d = 3$. The bounded weak solution u_i , $i = 0, \dots, J$, of the SPDE system (6) is α -Hölder continuous with probability 1, and has the form*

$$u_i(\mathbf{x}) = \sum_{p=-\infty}^{\infty} \hat{w}_i(p) L(u_{i-1})^{-1} \phi_p(\mathbf{x}) =: L(u_{i-1})^{-1} w_i(\mathbf{x}), \quad i = 0, \dots, J. \quad (8)$$

Proof. We will show that if u_{i-1} is continuous, then u_i is Hölder-continuous. This will prove continuity inductively, since the zeroth layer u_0 has constant u_{-1} . It is enough to verify continuity after conditioning with u_{i-1} , since $\mathbb{P}(u_i \in C^{0,\alpha}(\mathbb{T}^d)) = \mathbb{E}[\mathbb{P}(u_i \in C^{0,\alpha}(\mathbb{T}^d) \mid u_i)] = 1$ with probability 1 if and only if $\mathbb{P}(u_i \in C^{0,\alpha}(\mathbb{T}^d) \mid u_{i-1}) = 1$. After conditioning, u_i will be a Gaussian field for which we will apply Kolmogorov continuity criterium. To this end, we calculate

$$\mathbb{E}[|u_i(\mathbf{x}) - u_i(\mathbf{x}')|^{2b} \mid u_{i-1}] = C_b \sup_{\|f\|_{L^2} \leq 1} |L(u_{i-1})^{-1} f(\mathbf{x}) - L(u_{i-1})^{-1} f(\mathbf{x}')|^{2b},$$

which follows from the Itô isometry and the equivalent way to calculate $\|g\|_{L^2}$ as $\sup_{\|f\|_{L^2} \leq 1} \langle f, g \rangle$. When u_{i-1} is continuous, the function $L(u_{i-1})^{-1} f(\mathbf{x})$ belongs to H^2 by Lemma 3.1. By Sobolev's embedding theorem [47], H^2 embeds into $C^{0,\alpha}$ for $d \leq 3$. Hence,

$$\mathbb{E}[|u_i(\mathbf{x}) - u_i(\mathbf{x}')|^{2b} \mid u_{i-1}] \leq C |\mathbf{x} - \mathbf{x}'|^{2b\alpha},$$

which implies Hölder continuity with index smaller than $(2b\alpha - d)/2b = \alpha - d/2b$, where b can be arbitrarily large.

To complete the proof for existence of the solution, we insert the solution candidate (8) into the SPDE system (6) and do direct calculations. Uniqueness of the solution follows from Lax–Milgram theorem. □

We will show that u_i^N converges in probability to u_i . The proof uses uniform tightness of the distributions of $u_i^N - u_i$, which is a necessary condition for the convergence. We recall sufficient conditions for the uniform tightness (see p. 61 in [50]).

Lemma 3.4. *The random fields U^N are uniformly tight on $C(\mathbb{T}^d)$ if and only if there exists a function $K : C(\mathbb{T}^d) \rightarrow [0, \infty)$ with the following properties.*

- (1) *The set $\{g \in C(\mathbb{T}^d) : K(g) \leq C\}$ is compact for any $C > 0$,*
- (2) *$K(U^N) < \infty$ almost surely for every N , and*
- (3) *$\sup_N \mathbb{E}[K(U^N)] < \infty$.*

We will need several iterations of the logarithm so we define the iterated composition by setting $F(x) = \ln(1 + x)$, $F_0(x) = x$ and $F_{n+1}(x) = F \circ F_n(x)$.

Remark 2. The function F_i is increasing and, moreover, subadditive on non-negative numbers. That is, $F_n(x + y) \leq F_n(x) + F_n(y)$ for all $x, y \geq 0$, which follows by induction from subadditivity $\ln(1 + x + y) \leq \ln((1 + x)(1 + y)) = \ln(1 + x) + \ln(1 + y)$ of each F . Similar procedure shows that $F_n(xy) \leq F_n(x) + F_n(y)$.

Lemma 3.5. Let $d = 1, 2$ or 3 . Let κ be a continuous function with bounds $c_1 \exp(-a_1|t|) \leq \kappa(t) \leq c_2 \exp(a_2|t|)$, where $c_1, c_2, a_1, a_2 > 0$. Let u_i^N , $i = 0, \dots, J$, solve the system (7) and let u_i , $i = 0, \dots, J$ solve the system (6). Then the random fields u_i^N are uniformly tight on $C(\mathbb{T}^d)$, the random fields $u_i^N - u_i$ are uniformly tight on $C(\mathbb{T}^d)$, and the vector-valued random fields (u_i^N, u_i) are uniformly tight on $C(\mathbb{T}^d; \mathbb{R}^2)$.

Proof. We equip the Hölder space $C^{0,\alpha}(\mathbb{T}^d; \mathbb{R}^p)$ with its usual norm

$$\|g\|_\alpha = \sup_{x \neq y} \frac{|g(x) - g(y)|}{|x - y|^\alpha} + \sup_x |g(x)|.$$

Here α is chosen as in Lemma 3.3. We will use Kolmogorov–Chentsov tightness criterium (see [51]) to show the uniform tightness of the zeroth order layers, which are Gaussian. The desired estimate

$$\mathbb{E}[|u_0^N(\mathbf{x}) - u_0^N(\mathbf{x}')|^a] \leq C|\mathbf{x} - \mathbf{x}'|^{d+b}$$

follows from choosing large enough a in

$$\begin{aligned} \mathbb{E}[|u_0^N(\mathbf{x}) - u_0^N(\mathbf{x}')|^a] &\leq C \sup_{\|f\|_{L^2} \leq 1} \frac{|L_N(\ln(\kappa_0))^{-1}f(\mathbf{x}) - L_N(\ln(\kappa_0))^{-1}f(\mathbf{x}')|^a}{|\mathbf{x} - \mathbf{x}'|^{a\alpha}} |\mathbf{x} - \mathbf{x}'|^{a\alpha} \\ &\leq C \|L_N(\ln(\kappa_0))^{-1}\|_{L^2, C^{0,\alpha}}^a |\mathbf{x} - \mathbf{x}'|^{a\alpha}, \end{aligned}$$

where we applied Itô isometry and the definition of L^2 -norm as a supremum. By Lemma 3.1 and Sobolev’s embedding theorem, the operator norm of $L_N(\ln(\kappa_0))^{-1}$ is bounded for any $0 < \alpha < 1/2$. Since κ_0 is a constant, the bound is uniform. The case for $u_0^N - u_0$ and (u_0^N, u_0) follow similarly with the help of the triangle inequality.

For other layers, we use Lemma 3.4, where we choose $K(g) = F_i(g)$. For simplicity, we demonstrate Condition 1 only for $i = 3$, since the generalization is clear. The set

$$\{g : K(g) \leq C\} = \{g : \|g\|_\alpha \leq \exp(\exp(\exp(C) - 1) - 1) - 1\}$$

is clearly a closed set, which contains bounded equicontinuous functions. The set $A \subset C(\mathbb{T}^d)$ is then compact by the Arzelá–Ascoli theorem (see [52]). Moreover, the fields $u_i^N - u_i$ are almost surely α -Hölder continuous, by Lemma 3.3 and since the approximations belong to H_N . Hence, Condition 2 holds. To show Condition 3, we write u_i^N as $L_N(u_{i-1}^N)^{-1}w_i^N$ and u_i as $L(u_{i-1})^{-1}w_i$. By the triangle inequality and the subadditivity of F_i (see Remark 2), we can check the boundedness for u_i^N and u_i separately. Since the procedure is the same for both of the terms, we only show here the case for u_i^N . By Jensen’s inequality

$$\mathbb{E}[\ln(1 + \|u_i^N\|_\alpha) \mid u_{i-1}^N] \leq \ln(\mathbb{E}[1 + \|L_N(u_{i-1}^N)^{-1}w_i^N\|_\alpha \mid u_{i-1}^N]). \quad (9)$$

The conditioning with u_{i-1}^N lets us compute expectations of Gaussian variables. Instead of attacking directly the expectation in Equation (9), we will seek a Gaussian zero mean random field U with larger variance than the conditioned u_i^N . Then also certain other expectations of u_i^N will be bounded by expectations of U . Under conditioning, the variances of the random variables $\int \phi(\mathbf{x}) L_N(u_{i-1}^N)^{-1} w_i^N(\mathbf{x}) d\mathbf{x}$ are

$$\mathbb{E}[\langle w_i^N, (L_N(u_{i-1}^N)^{-1})^* \phi \rangle^2 \mid u_{i-1}^N] = \|P_N(L_N(u_{i-1}^N)^{-1})^* \phi\|_{L^2}^2 \leq \|L_N(u_{i-1}^N)^{-1}\|_{L^2, H^2}^2 \|\phi\|_{H^{-2}}^2,$$

where the last inequality follows from properties of the adjoint operator and Lemma 3.1. Set U to be a zero mean Gaussian random field U on \mathbb{T}^d whose covariance is defined by equations $\mathbb{E}[\langle U, \phi \rangle^2] = \|\phi\|_{H^{-2}}^2$ for all smooth ϕ . Then U has sample paths in Hölder space $C^{0,\alpha}(\mathbb{T}^d)$ by Sobolev's embedding theorem and Kolmogorov's continuity theorem (see [53]). By Fernique's theorem (e.g. [54]), the expectation $\mathbb{E}[\|U\|_\alpha]$ is finite. Since norms are absolutely continuous functions, also the conditional expectation of $\|u_i^N\|_\alpha$ is bounded by $\|L(u_{i-1}^N)^{-1}\|_{L^2, H^2} \mathbb{E}[\|U\|_\alpha]$ (see Corollary 3.3.7 in [54]).

Further application of Lemma 3.1 in Equation (9) gives our main estimate

$$\begin{aligned} \mathbb{E}[\ln(1 + \|u_i^N\|_\alpha) \mid u_{i-1}^N] &\leq \ln \left[1 + c \|\kappa(u_{i-1}^N(\mathbf{x}))\|_\infty^\nu \frac{\max(1, \|\kappa^2(u_{i-1}^N(\mathbf{x}))\|_\infty)}{\min(1, \inf \kappa^2(u_{i-1}^N(\mathbf{x})))^2} \mathbb{E}[\|U\|_\alpha] \right] \\ &\leq c_{\alpha,\nu} (1 + \|u_{i-1}^N\|_\infty), \end{aligned} \quad (10)$$

where we applied the bounds of κ and the elementary inequalities $\max(a, b \exp(c)) \leq \max(a, b) \exp(|c|)$, $\min(1, \inf \exp(g(x))) \geq \exp(-\|g\|_\infty)$, $(1 + ab) \leq (1 + a)(1 + b)$ and $1 + a \leq 2 \max(1, a)$. Here we can choose constants larger than 1.

When $i = 1$, the expectations (10) are uniformly bounded, since expectations of $\|u_0^N\|_\infty$ are bounded. Then Lemma 3.4 shows that u_1^N are uniformly tight. For the subsequent layers we need to operate multiple times with the logarithm and Jensen's inequality through inductive steps

$$\begin{aligned} \mathbb{E}[F_i(\|u_i^N\|_\alpha)] &\leq \mathbb{E}[F_i(\mathbb{E}[\|u_i^N\|_\alpha \mid u_{i-1}^N])] \leq \mathbb{E}[F_i(c(1 + \|u_{i-1}^N\|_\alpha))] \\ &\leq C + \mathbb{E}[F_{i-1}(\|u_{i-1}^N\|_\alpha)] \end{aligned}$$

with the help of the additivity properties of F_i from Remark 2. □

Theorem 3.6. *Let $d = 1, 2$, or 3 . Let κ be a continuously differentiable function with bounds $c_1 \exp(-a_1|t|) \leq \kappa(t)$, $|\kappa'(t)| \leq c_2 \exp(a_2|t|)$, where $c_1, c_2, a_1, a_2 > 0$. Let u_i^N , $i = 0, \dots, J$. The solution (u_0^N, \dots, u_J^N) of the Galerkin system (7) converge in probability to the weak solution (u_0, \dots, u_J) of the SPDE system (6) on $L^2(\mathbb{T}^d; \mathbb{R}^{J+1})$ as $N \rightarrow \infty$.*

Proof. It is enough to show componentwise convergence. Uniform tightness on $C(\mathbb{T}^d)$ implies uniform tightness on $L^2(\mathbb{T}^d)$, which in turn implies relative compactness in weak topology of distributions. Hence, by Lemma 3.5 each subsequence of the distributions of

$u_i^N - u_i$ has a weakly convergent subsequence, say $u_i^{N_k} - u_i$. Recall, that the convergence in probability is equivalent to the convergence of $\tau_N := \mathbb{E}[\min(1, \|u_i^N - u_i\|_{L^2})]$, where $\min(1, \|\cdot\|_{L^2})$ is now a bounded continuous function. By weak convergence of distributions, the subsequence τ_{N_k} has some limit. It remains to verify that limits of $u_i^{N_k} - u_i$ in distribution are zero. The characteristic functions of $u_i^N - u_i$ converge to 1, if $\tau_N(\phi) := \mathbb{E}[\min(1, |\langle u_i^N - u_i, \phi \rangle|)]$, converge to zero for every $\phi \in L^2$. This formulation makes calculation of expectations manageable. The essential difference to other approaches arises from the monotonicity of conditional expectations. Namely, the property $\mathbb{E}[1 - \min(1, G) | \Sigma_0] \geq 0$ for $G \geq 0$ implies that

$$\mathbb{E}[\min(1, G) | \Sigma_0] = \min(1, \mathbb{E}[\min(1, G) | \Sigma_0]) \leq \min(1, \mathbb{E}[G | \Sigma_0]). \quad (11)$$

Especially,

$$\tau_N(\phi) \leq \mathbb{E}[\min(1, \mathbb{E}[|\langle u_i^N - u_i, \phi \rangle| | w_0, \dots, w_{i-1}])], \quad (12)$$

where we applied (11) after taking a conditional expectation inside the expectation. We condition with white noises in order to handle both u_i and its approximation u_i^N easily. Moreover, the cutoff function $\min(1, \cdot)$ is nondecreasing and subadditive.

Under conditioning with w_0, \dots, w_{i-1} , the random fields u_i^N and u_i in Equation (12) become Gaussian, which enables us to compute

$$\begin{aligned} \tau_N(\phi) &= \mathbb{E}[\min(1, c\|(P_N(L_N(u_{i-1}^N)^{-1})^* - (L(u_{i-1})^{-1})^*)\phi\|_{L^2})] \\ &\leq \mathbb{E}[\min(1, c\phi\|L_N(u_{i-1}^N)^{-1} - L(u_{i-1})^{-1}\|_{L^2, L^2}) + \min(1, c\|(P_N - I)(L(u_{i-1})^{-1})^*\phi\|_{L^2})] \\ &=: \tau_N^1(\phi) + \tau_N^2(\phi) \end{aligned}$$

via the Itô isometry and the properties of adjoints. Since $L(u_{i-1})^*\phi \in L^2$, the term $\tau_N^2(\phi)$ converges to zero. We will apply Lemma 3.2 for the difference of operators $L_N(u_{i-1}^N)^{-1}$ and $L^{-1}(u_{i-1})$ in $\tau_N^1(\phi)$, which leads to a well-behaving estimate

$$\begin{aligned} \tau_N^1(\phi) &\leq \mathbb{E}\left[\min\left(1, \frac{C_\phi}{N} \exp(C_1\|u_{i-1}\|_\infty)\right)\right] + \mathbb{E}[\min(1, C_\phi \exp(C_2(\|u_{i-1}\|_\infty + \|u_{i-1}^N\|_\infty)) \\ &\quad \times \|u_{i-1}^N - u_{i-1}\|_{L^2}^{1/6})] =: \tau_N^{11}(\phi) + \tau_N^{12}(\phi). \end{aligned}$$

The first term $\tau_N^{11}(\phi)$ converges to zero by Lebesgue's dominated convergence theorem. For the term $\tau_N^{12}(\phi)$, we will use the uniform tightness of the the vector-valued random fields (u_{i-1}^N, u_{i-1}) shown in Lemma 3.5. Let $K_{i-1} = K_{i-1}(\epsilon) \subset C(\mathbb{T}^d; \mathbb{R}^2)$ be a compact set for which $P((u_{i-1}^N, u_{i-1}) \in K^C) < \epsilon$. Then

$$\begin{aligned} \tau_N^{12}(\phi) &= \mathbb{E}[\min(1, C_\phi \exp(C_2(\|u_{i-1}\|_\infty + \|u_{i-1}^N\|_\infty))\|u_{i-1}^N - u_{i-1}\|_{L^2}^{1/6})(1_{K^C} + 1_{K^C})] \\ &\leq \mathbb{E}[\min(1, C_\epsilon\|u_{i-1}^N - u_{i-1}\|_{L^2})]^{1/6} + \mathbb{E}[1_{K^C}] \end{aligned}$$

by Lemma 3.2, Remark 1 and Jensen's inequality. Thus $\tau_N^{12}(\phi)$ converges to zero if $u_{i-1}^N - u_{i-1}$ converge to zero in probability on L^2 .

When $i = 0$, the above procedure shows then that u_0^N converges to u_0 in probability, because the sublayers are then constants. By induction, u_i^N converges in probability to u_i . \square

Prohorov's theorem [51] hands us weak convergence on the space of continuous functions.

Corollary. *The distributions of u_i^N converge weakly to the distribution of u_i on $C(\mathbb{T}^d)$ and the joint distribution of (u_0^N, \dots, u_J^N) converges weakly to the joint distribution of (u_0, \dots, u_J) on $C(\mathbb{T}^d, \mathbb{R}^{J+1})$.*

Proof. The random fields u_i^N are tight on $C(\mathbb{T}^d)$ by Lemma 3.5. By Prohorov's theorem, the closure of their distributions forms a sequentially compact set, implying the existence of weak limit. By Theorem 3.6, each weakly converging subsequence of the distributions has the same limit, that is, the distribution of u_i . Indeed, since convergence in probability on $L^2(\mathbb{T}^d)$ implies weak convergence on $L^2(\mathbb{T}^d)$, the characteristic functions $\mathbb{E}[\exp(i\langle u_i^N, \phi \rangle)]$ converge to $\mathbb{E}[\exp(i\langle u_i, \phi \rangle)]$ for all $\phi \in L^2(\mathbb{T}^d)$. The set of bounded continuous functions $u \mapsto \exp(i\langle u, \phi \rangle)$, where ϕ are smooth on \mathbb{T}^d , separate the functions in $C(\mathbb{T}^d)$. Hence, the limits of these characteristic functions are enough to identify the weak limit on $C(\mathbb{T}^d)$. The joint distribution is handled similarly. \square

We recall a posterior convergence result from [55] with notation used in [17].

Theorem 3.7. *Let the posterior distribution $\mu^{\mathbf{y}}$ of an unknown u given an observation \mathbf{y} have the Radon–Nikodym density*

$$\frac{d\mu^{\mathbf{y}}}{d\mu^0}(u) \propto \exp(-\Phi(u, \mathbf{y}))$$

with respect to the prior distribution μ^0 of u , where $\Phi(\cdot, \mathbf{y}) \geq -C_{\mathbf{y}}$. Let u^N be an approximation of u and let the posterior distribution of u^N given an observation \mathbf{y}_N have also the Radon–Nikodym density

$$\frac{d\mu_N^{\mathbf{y}_N}}{d\mu_N^0}(u) \propto \exp(-\Phi(u, \mathbf{y}_N))$$

with respect to the prior distribution μ_N^0 of u^N .

If the prior distributions μ_N^0 converge weakly to μ^0 , then the posterior distributions $\mu_N^{\mathbf{y}}$ converge weakly to $\mu^{\mathbf{y}}$.

Especially, the joint prior distribution of (u_0^N, \dots, u_J^N) converge weakly on $C(\mathbb{T}^d, \mathbb{R}^{J+1})$ to the joint distribution of (u_0, \dots, u_J) . Hence, the corresponding posteriors also converge weakly.

4. Bayesian inference algorithm

We need to develop a Bayesian inference procedure to sample the Fourier coefficients from a posterior distribution where the prior is given by a multi-layered Gaussian field. Let us work directly with the Fourier coefficient \mathbf{u}_J of u_J , and denote by

$\mu^0(d\mathbf{u}_J) = \mathbb{P}(d\mathbf{u}_J)$ and $\mu^{\mathbf{y}}(d\mathbf{u}_J) = \mathbb{P}(d\mathbf{u}_J \mid \mathbf{y})$ the prior and the posterior distribution of u_J , the unknown target field, respectively, when the measurement is given by

$$\mathbf{y} = \mathbf{H}\mathbf{u}_J + \mathbf{e}. \quad (13)$$

In this equation, \mathbf{y} is a vector which contains all of the measurements. The elements of the matrix \mathbf{H} correspond to the linear mappings $\{h_k\}$ for the respective Fourier components. With the number of measurements being m , the measurement noise \mathbf{e} is an m -dimensional Gaussian random vector with zero mean and covariance \mathbf{E} . The posterior distribution $\mu^{\mathbf{y}}$ will be absolutely continuous with respect to the prior μ^0 , and the density of the posterior with respect to the prior is given by

$$\frac{d\mu^{\mathbf{y}}}{d\mu^0}(\mathbf{u}_J) = \frac{1}{Z} \exp(-\Phi(\mathbf{u}_J, \mathbf{y})), \quad (14)$$

where $\Phi(\mathbf{u}_J, \mathbf{y}) := \frac{1}{2} \|\mathbf{E}^{-1/2}(\mathbf{y} - \mathbf{H}\mathbf{u}_J)\|^2$ is the potential function and the normalization constant $Z := \int \exp(-\Phi(\mathbf{u}_J, \mathbf{y}))\mu^0(d\mathbf{u}_J)$. The posterior distribution $\mu^{\mathbf{y}}(d\mathbf{u}_J)$ can be written in the following form

$$\mu^{\mathbf{y}}(d\mathbf{u}_J) \propto \exp(-\Phi(\mathbf{u}_J, \mathbf{y}))\mu^0(d\mathbf{u}_J). \quad (15)$$

In the next section, we describe the MCMC algorithm to sample from the posterior distribution.

4.1. Non-centered algorithm

Consider the case of hierarchical prior distribution $\mu(d\mathbf{u}_J)$ with J hyperprior layers which we construct as

$$\mathbf{L}_0\mathbf{u}_0 = \mathbf{w}_0, \quad (16a)$$

$$\mathbf{L}(\mathbf{u}_{j-1})\mathbf{u}_j = \mathbf{w}_j, \quad j = 1, \dots, J, \quad (16b)$$

where \mathbf{u}_j contains the Fourier coefficients of the field u_j , for $j = 0, \dots, J$. We fix $g(x) = \exp(-x)$. Since $u_{-1} = \ln(\kappa_0)$ is a constant, we can write, $\mathbf{L}_0 = \mathbf{L}(\mathbf{u}_{-1})$, where k -th element of \mathbf{u}_{-1} is equal to $\ln(\kappa_0)\delta_{k,0}$, where $\delta_{k,0}$ is the Kronecker delta. By (16), we can define a linear transformation from \mathbf{w}_j to \mathbf{u}_j for $j > 0$ as follows:

$$\mathbf{u}_j = \tilde{U}(\mathbf{w}_j, \mathbf{u}_{j-1}) := \mathbf{L}(\mathbf{u}_{j-1})^{-1}\mathbf{w}_j. \quad (17)$$

Using (17) we can define a transformation from \mathbf{w} to \mathbf{u} as follows

$$\mathbf{u} = U(\mathbf{w}) = (\tilde{U}(\mathbf{w}_0, \mathbf{u}_{-1}), \tilde{U}(\mathbf{w}_1, \cdot) \circ \tilde{U}(\mathbf{w}_0, \mathbf{u}_{-1}), \dots, \tilde{U}(\mathbf{w}_J, \cdot) \circ \dots \circ \tilde{U}(\mathbf{w}_0, \mathbf{u}_{-1})). \quad (18)$$

The dependence of \mathbf{u}_j on \mathbf{u}_{j-1} , $j = 1, \dots, J$, leads us to

$$\mu^0(d\mathbf{u}_J) = \mathbb{P}(d\mathbf{u}_0) \prod_{j=1}^J \mathbb{P}(d\mathbf{u}_j \mid \mathbf{u}_{j-1}). \quad (19)$$

When evaluating the posterior function (15), it is necessary to compute the log determinant of $\mathbf{L}(\mathbf{u}_{j-1})$ in $\mathbb{P}(d\mathbf{u}_j|\mathbf{u}_{j-1})$ for each $1 < j \leq J$, respectively. These calculations are expensive in general [56]. There is also a singularity issue if we sample directly from $\mathbb{P}(d\mathbf{u}|\mathbf{y})$ if N approaches infinity [16]. We can avoid these issues by using the reparametrization (17), where instead of sampling the Fourier coefficients \mathbf{u} , we sample the Fourier coefficient of the noises \mathbf{w} , which is then called *non-centered* algorithm [25, 57]. That is, we can write:

$$\frac{d\tilde{\mu}^{\mathbf{y}}}{d\tilde{\mu}^0}(\mathbf{w}) = \frac{1}{\tilde{Z}} \exp\left(-\tilde{\Phi}(\mathbf{w}, \mathbf{y})\right) := \frac{1}{\tilde{Z}} \exp(-\Phi(U(\mathbf{w}), \mathbf{y})). \quad (20)$$

In this equation, $\tilde{\mu}^0 := \mathbb{P}(d\mathbf{w})$ and $\tilde{\mu}^{\mathbf{y}} := \mathbb{P}(d\mathbf{w}|\mathbf{y})$ are the prior and posterior of the \mathbf{w} , respectively. The preconditioned Crank–Nicolson (pCN) algorithm [19] can be used to sample from $\mathbb{P}(d\mathbf{w}|\mathbf{y})$, and it is well defined even for the case of N goes to infinity, using the fact that the prior for \mathbf{w} is a standard Gaussian distribution. One implementation of the pCN algorithm is given by Algorithm 1. Due to linearity assumption of the forward model (13), we can leverage the standard Gaussian regression in addition to the non-centered reparametrization. This procedure has been proposed in [16] for general deep Gaussian fields. Here, we adapt this algorithm for our Galerkin method. The resulting algorithm is a Metropolis within Gibbs type [58] where the Fourier coefficients $\{\mathbf{u}_j\}, j = 0, \dots, J-1$ are sampled using the pCN algorithm via reparametrization (17), and the Fourier coefficients for the last layer \mathbf{u}_J are sampled directly. The detail is given as follows. By marginalization of \mathbf{u}_{J-1} , we can write $\mathbb{P}(d\mathbf{u}_J|\mathbf{y}) = \int \mathbb{P}(d\mathbf{u}_J | \mathbf{u}_{J-1}, \mathbf{y}) \mathbb{P}(d\mathbf{u}_{J-1} | \mathbf{y})$. Furthermore, from the standard Gaussian regression, we can sample directly from $\mathbb{P}(d\mathbf{u}_J|\mathbf{u}_{J-1}, \mathbf{y})$ by using:

$$\begin{aligned} V(\mathbf{u}_{J-1}, \mathbf{y}) &:= \begin{pmatrix} \mathbf{E}^{-1/2}\mathbf{H} \\ \mathbf{L}(\mathbf{u}_{J-1}) \end{pmatrix}^\dagger \left(\begin{pmatrix} \mathbf{E}^{-1/2}\mathbf{y} \\ 0 \end{pmatrix} + \tilde{\mathbf{v}} \right), \\ \mathbf{u}_{J-1} &= \tilde{U}(\mathbf{w}_{J-1}, \cdot) \circ \dots \circ \tilde{U}(\mathbf{w}_0, \mathbf{u}_{-1}), \\ \tilde{\mathbf{v}} &\sim N(0, \mathbf{I}). \end{aligned} \quad (21)$$

Writing $\mathbf{y} = \mathbf{HL}(\mathbf{u}_{J-1})^{-1}\mathbf{w}_J + \mathbf{e}$, the conditional probability density of \mathbf{y} given \mathbf{u}_{J-1} is given by $p(\mathbf{y}|\mathbf{u}_{J-1}) = N(\mathbf{y}|0, \mathbf{HL}(\mathbf{u}_{J-1})^{-1}\mathbf{L}(\mathbf{u}_{J-1})^{-\top}\mathbf{H}^\top + \mathbf{E})$. Let $\bar{\mathbf{w}} = (\mathbf{w}_0, \dots, \mathbf{w}_{J-1})$. To obtain samples from $\mathbb{P}(d\mathbf{u}_{J-1}|\mathbf{y})$ we can use reparametrization (17) and Algorithm 1 to sample from $\mathbb{P}(d\bar{\mathbf{w}}|\mathbf{y})$. The probability distribution $\mathbb{P}(d\bar{\mathbf{w}}|\mathbf{y})$ is given as follows:

$$\mathbb{P}(d\bar{\mathbf{w}}|\mathbf{y}) \propto \exp\left(-\tilde{\Psi}(\bar{\mathbf{w}}, \mathbf{y})\right) \mathbb{P}(d\bar{\mathbf{w}}), \quad (22a)$$

$$\tilde{\Psi}(\bar{\mathbf{w}}, \mathbf{y}) := \Psi(U(\bar{\mathbf{w}}), \mathbf{y}) = \frac{1}{2} \|\mathbf{y}\|_{\mathbf{Q}}^2 + \frac{1}{2} \log \det(\mathbf{Q}), \quad (22b)$$

$$\mathbf{Q} = \mathbf{HL}(\mathbf{u}_{J-1})^{-1}\mathbf{L}(\mathbf{u}_{J-1})^{-\top}\mathbf{H}^\top + \mathbf{E}. \quad (22c)$$

To sample from $\mathbb{P}(d\bar{\mathbf{w}}|\mathbf{y})$ using Algorithm 1, we use $J-1$, and $\tilde{\Psi}$ for J and $\tilde{\Phi}$, respectively.

Algorithm 1: preconditioned Crank–Nicolson algorithm.

Input : $J, \mathbf{w}, \mathbf{y}, \tilde{\Phi}(\mathbf{w}, \mathbf{y})$
Output: accepted, \mathbf{w}
 accepted = 0
 draw $\mathbf{w}' \sim N(0, \mathbf{I})$
 $\tilde{\mathbf{w}} = \sqrt{1 - s^2}\mathbf{w} + s\mathbf{w}'$
 logRatio = $\tilde{\Phi}(\mathbf{w}, \mathbf{y}) - \tilde{\Phi}(\tilde{\mathbf{w}}, \mathbf{y})$
 draw $\omega \sim \text{Uniform}[0, 1]$
if logRatio > $\ln(\omega)$ **then**
 $\mathbf{w} = \tilde{\mathbf{w}}$
 accepted = 1

It is worth noting that the proposed approach could also be extended to nonlinear measurement models. However, a few changes need to be made. First, for a nonlinear forward problem $\mathbf{h} : L^2(\Omega) \rightarrow \mathbb{R}^m$, we form a composition $\mathbf{h} \circ \mathcal{F}^{-1}(\mathbf{u}_J)$, where \mathcal{F}^{-1} is the inverse Fourier transform. Secondly, the potential function $\Phi(\mathbf{u}_J, \mathbf{y})$ in (14) needs to be replaced with $\Phi(\mathbf{u}_J, \mathbf{y}) := \frac{1}{2} \|\mathbf{E}^{-1/2}(\mathbf{y} - \mathbf{h} \circ \mathcal{F}^{-1}(\mathbf{u}_J))\|^2$. Using this new potential, we could come up with equivalent expressions to ones given in (21) and (22).

4.2. Computational Cost

We need to estimate the computational cost in generating one sample by the proposed MCMC algorithm. For this purpose, we fix the following numbers: J, n, d , and m . Recall that we have the following relation $\bar{N} = 2N + 1 = (2n + 1)^d$, where N is the total number of the Fourier basis functions. Suppose that we would like to produce an estimate of the unknown in $\bar{N}_o = (2n_o + 1)^d$ points, for some $n_o > n$. We will estimate the computational complexity from a naive implementation of this algorithm without considering possible optimizations (e.g., parallelization) as it is beyond the scope of this article.

First, we need to generate $\mathbf{w}_0, \dots, \mathbf{w}_J$, which has a negligible cost compared to the rest of linear algebraic evaluations. Then we need to evaluate Fourier coefficients \mathbf{u}_j for each layer. The computational cost to obtain \mathbf{u}_0 in (16a) is cheap since \mathbf{L}_0 is a fixed diagonal matrix. For the remaining layers, let us now examine the cost to calculate $\mathbf{L}(\mathbf{u}_{j-1})$ for $j = 1, \dots, J$. Based on (4) we need to evaluate the following steps:

- (i) Inverse Fourier transform of \mathbf{u}_{j-1} , which costs $\mathcal{O}(\bar{N}_o \log(\bar{N}_o))$ by IFFT.
- (ii) Two element-wise evaluations of nonlinear function $\kappa(u_{j-1})^\gamma$.
- (iii) Two Fourier transforms of $\kappa(u_{j-1})^\gamma$, which cost $\mathcal{O}(\bar{N}_o \log(\bar{N}_o))$ by FFT.
- (iv) Forming of M_N in (3), which is only a matrix element assignment with two for loops.

- (v) One matrix addition $\mathcal{O}(\bar{N}^2)$ as well as one dense matrix and diagonal matrix multiplication, which totally cost $\mathcal{O}(\bar{N}^2)$, and one scalar and matrix multiplication of cost $\mathcal{O}(\bar{N}^2)$.

Therefore, the computational cost to evaluate $\mathbf{L}(\mathbf{u}_{j-1})$ for $j = 1, \dots, J$ is $\mathcal{O}(J\bar{N}^2)$. To obtain the Fourier coefficients \mathbf{u}_j , we use a linear algebraic solver to evaluate $\mathbf{L}(\mathbf{u}_{j-1})^{-1}\mathbf{w}_j$. This computation costs $\mathcal{O}(\bar{N}^3)$ [59]. Performing this computation for $j = 0$ until $J - 1$ translates to $\mathcal{O}(J\bar{N}^3)$. For the last layer, we need to solve the least squares problem (21), which has the complexity of $\mathcal{O}((m + \bar{N})\bar{N}^2)$ [59], assuming that the $\mathbf{E}^{-1}\mathbf{H}$ and $\mathbf{E}^{-1}\mathbf{y}$ are precomputed, and the least squares problem is solved via QR factorization and triangular solver. The last part is to compute the equation (22). To compute Q in (22c), we need $\mathcal{O}(\bar{N}^2m) + \mathcal{O}(\bar{N}^3)$ to evaluate $\mathbf{Z} = \mathbf{L}(\mathbf{u}_{J-1})^{-\top}\mathbf{H}^\top$, and then we need $\mathcal{O}(m^2\bar{N})$ to evaluate $\mathbf{Z}^\top\mathbf{Z} + \mathbf{E}$. Then we use QR factorization to compute the log determinant and the weighted norm terms, which cost $\mathcal{O}(m^3)$. Hence, using the assumption that the ratio $\varrho = \bar{N}_o/\bar{N}$ is not too big so that $\mathcal{O}(\bar{N}_o \log(\bar{N}_o)) \approx \mathcal{O}(\varrho\bar{N} \log(\bar{N}))$ and $m \leq \bar{N}$, the total computational cost for generating one sample is $\mathcal{O}(J\bar{N}^3)$.

5. Numerical results

Now, we present examples of Bayesian inversion using a multi-layer Gaussian prior presented in the previous section. Our main focus in this section is to show the effectiveness of the proposed finite-dimensional approximation method for selected examples. Therefore, we will not discuss the properties of the MCMC algorithm used to generate the samples as they are based on the MCMC algorithms described in [60, 16]. We aim at acceptance ratio between 25–50%, which is obtained by tuning the pCN step size s in Algorithm 1. Our experience in the numerical implementations below indicates that the MCMC algorithm based on the pCN and non-centered algorithm is quite robust. For one and two hyperprior layers implementation, the step size s does not need to be extremely small. The step sizes s in the first and second examples are varying around 10^{-1} to 10^{-3} .

Depending upon the characteristic of the unknown, such as the presence of discontinuities and whether the length scales of the unknown are varied locally, we advise to choose the number of the hyperprior layers as low as possible and only opt to a higher number if a lower number gives an unsatisfactory result. The reason is that with the addition of the hyperprior layers, the number of the Fourier coefficients increases linearly. This will greatly increase the mixing time for the MCMC algorithm.

5.1. Continuous-time random processes with finite-time discrete measurements

Let us consider the application of the proposed technique to address the non-parametric denoising of two piecewise smooth signals. The first test signal is a rectangular shape signal where the value is zero except on interval $[0.2, 0.8]$. The second test signal is a

	Rectangle		Bell-Rectangle	
J	L^2 error	PSNR	L^2 error	PSNR
1	1.044	24.325	1.527	21.603
2	0.922	25.605	1.475	21.801

Table 1. Quantitative performance comparison of a shallow Gaussian fields prior inversion for one dimensional signal in Section 5.1.

combination between a smooth bell shaped signal and a rectangular signal [12]:

$$v_{rect}(t) = \begin{cases} 1, & t \in [0.2, 0.8], \\ 0, & \text{otherwise.} \end{cases}, \quad v_{bell,rect}(t) = \begin{cases} \exp\left(4 - \frac{1}{2t-4t^2}\right), & t \in (0, 0.5), \\ 1, & t \in [0.7, 0.8], \\ -1, & t \in (0.8, 0.9], \\ 0, & \text{otherwise.} \end{cases} \quad (23)$$

Previously, in [28], for $J = 1$ and with the unknown signal $v_{bell,rect}(t)$, we have demonstrated that upon increasing the number of the Fourier basis functions, the L^2 error between the ground truth and the posterior sample mean decreased significantly.

In this section, we will compare the estimation results using one hyperprior and two hyperprior layers respectively. To allow high variation near points of discontinuities, the length-scale of v is expected to be smaller around the discontinuities than the rest of the domain. We take measurement of on one dimensional grid of 2^8 points and set the standard deviation of the measurement noise to be 0.1. The proposed algorithm is tested with $N = 2^6 - 1$. After estimation, we reconstruct the signal using inverse Fourier transform with a finer grid with 2^8 points equally spaced between zero and one. We take ten million samples for each MCMC run.

To have a fair comparison, we use the same measurement record for each run with different J . Figures 2 and 3 show the reconstructed signals and their respective length-scale estimations. It can be clearly seen that the addition of another hyper prior layer improves the reconstruction result for the unknown signals. For $J = 2$, although there are no sudden drops near the points of discontinuities, the length-scale value is sufficiently high in the a smooth part of v signal, and substantially low near the points of discontinuities. This variation translates to a better smoothness detection, as can be seen in Figures 2a, 3a, 2c, and 3c. In contrast, Figures 2b and 3b show that when $J = 1$, the posterior sample means are considerably overfitting the data on the smooth part of the signals. Although the length-scale drops suddenly near the points of discontinuities of the ground truth v , the variation of ℓ is limited (see Figure 2d and 3d). This contributes to a decreased smoothness in the smooth region of the sample mean. The quantitative performance is given in Table 1.

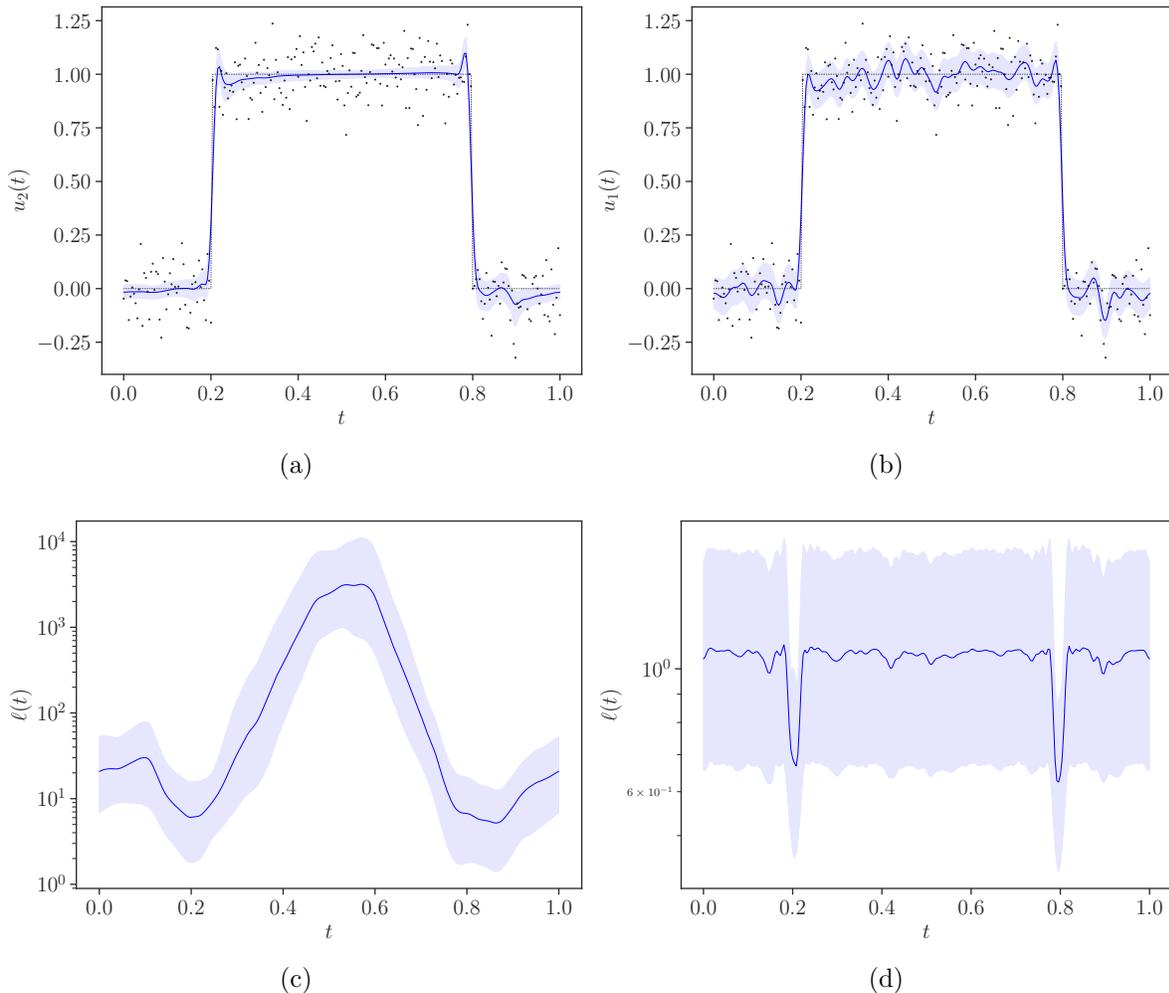


Figure 2. Simulation results of example in Section 5.1. Figures 2a, 2c, 2b, and 2d describe the lowest random fields from u for $J = 2$ and $J = 1$, and their length-scales. In Figure 2a and 2b, the blue and the black lines are the mean Fourier inverse of the samples in the 95 % confidence shades of the lowest layer and the original unknown signals respectively. The blue line and the shades in the remaining figures are sample means and 95 % confidence interval of the estimated length-scales.

5.2. X-ray tomography

Now we apply the methods developed to the 2D X-ray tomography reconstruction problem. For the tomography problem, the linear functional is given by a line integration known as the Radon transform [61, 62, 39, 63].

Assume that the field of interest has support in the a circle with center at $(1/2, 1/2)$ and radius equal to half in the two dimensional Euclidean space. Also recall that at a distance r with detection angle θ , we can write the Radon transform of the Fourier basis $\phi_{\mathbf{k}} = \exp(i2\pi\mathbf{k}^\top \mathbf{x})$ as below

$$\langle \phi_{\mathbf{k}}, H_{r,\theta} \rangle = \int_{\Omega} \chi\left(x - \frac{1}{2}, y - \frac{1}{2}\right) \exp(i2\pi\mathbf{k}^\top \mathbf{x}) \delta\left(r - \left(\left(x - \frac{1}{2}\right) \cos \theta + \left(y - \frac{1}{2}\right) \sin \theta\right)\right) d\mathbf{x},$$

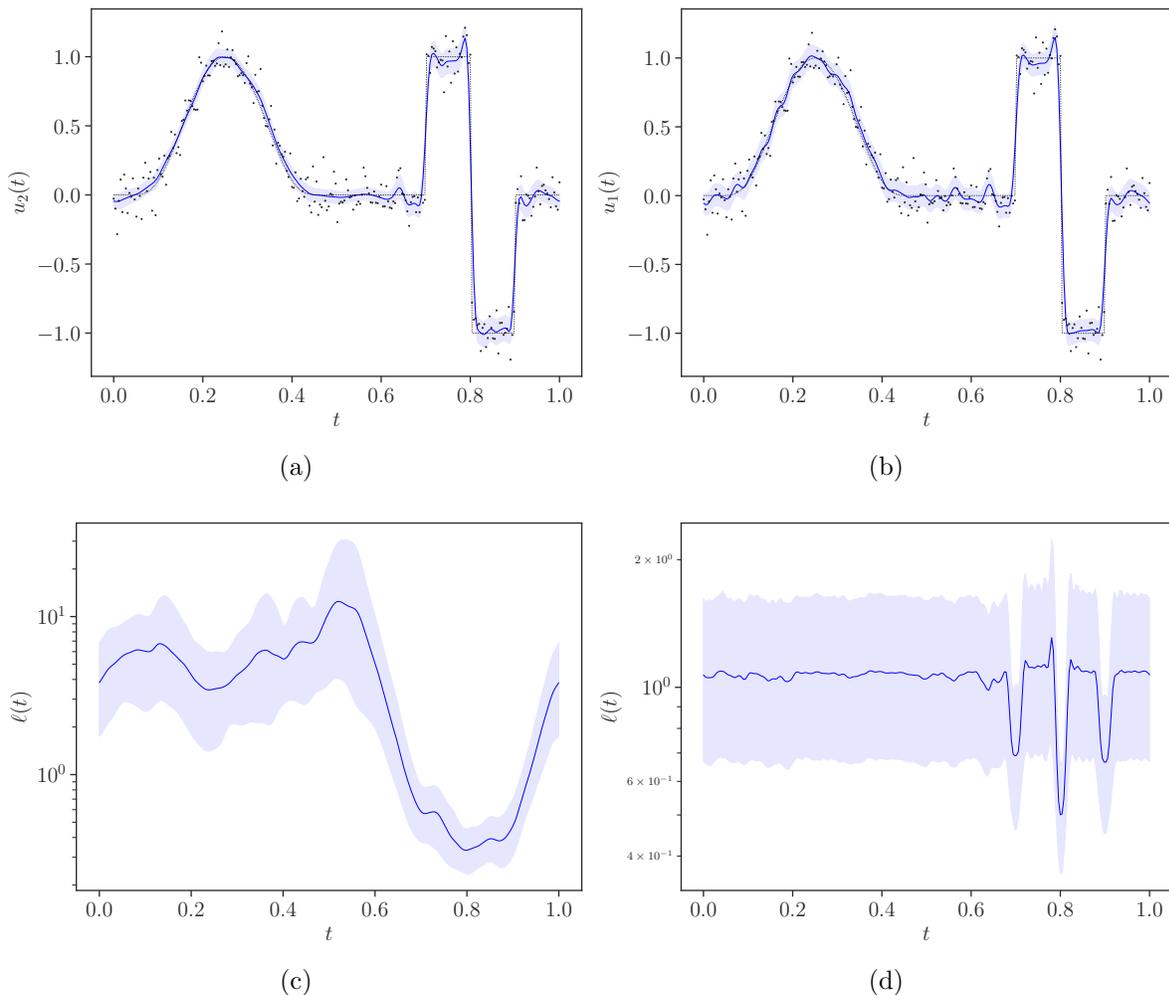


Figure 3. Similar to Figure 2, with the unknown is $v_{bell,rect}$.

where $\chi(x, y)$ is an indicator function with support in $(x^2 + y^2 < \frac{1}{4})$. Introducing a rotation matrix \mathbf{R}_θ , and $\mathbf{p} = [p \ q]^\top$, and $x' = x - \frac{1}{2}, y' = y - \frac{1}{2}$, $\mathbf{p} = \mathbf{R}_\theta \mathbf{x}$ we can write

$$\langle \phi_{\mathbf{k}}, H_{r,\theta} \rangle = \exp(i\pi(k_x + k_y)) \int_{\Omega} \chi(x', y') \exp(i2\pi(\mathbf{R}_\theta \mathbf{k})^\top \mathbf{p}) \delta(r - p) d\mathbf{p}.$$

Using the assumption we have mentioned and $\tilde{\mathbf{k}} = [\tilde{k}_x \ \tilde{k}_y]^\top = \mathbf{R}_\theta \mathbf{k}$, we end up with $\langle \phi_{\mathbf{k}}, H_{r,\theta} \rangle = \exp(i\pi(k_x + k_y)) \exp(i2\pi\tilde{k}_x r) \frac{1}{\pi\tilde{k}_y} [\sin(2\pi\tilde{k}_y \sqrt{\frac{1}{4} - r^2})]$. Notice when $\tilde{k}_y = 0$ we replace the above equation with its limit, that is, $\lim_{\tilde{k}_y \rightarrow 0} \langle \phi_{\mathbf{k}}, H_{r,\theta} \rangle = 2\sqrt{\frac{1}{4} - r^2} \exp(i\pi(k_x + k_y)) \exp(i2\pi\tilde{k}_x r)$.

We modify the MCMC implementation used for the previous one-dimensional example to suit to a GPU architecture. For our test comparison in X-ray tomography, the Shepp–Logan phantom with 511×511 resolution is used (see Figure 4a). We will use this phantom to evaluate the proposed method.

We take 45 sparsely full projections out of 180 degrees using parallel beams. The measurement is corrupted by a white Gaussian noise with standard deviation 0.2. Due

	Shallow-GP	FBP	Tikhonov
PSNR	22.246	23.399	21.673
L^2 Error	44.795	47.156	42.145

Table 2. Quantitative performance comparison of a shallow Gaussian field Bayesian inversion for the tomography application in Section 5.2.

to the restriction in GPU memory, J is set to one. To excel the speed Fourier transform and inverse Fourier transform computations, we use FFT and IFFT routine from CUPY [64]. For a performance comparison, we perform the filtered back projection (FBP) on sinogram using iradon routine from Skimage [65]. A Tikhonov regularization is also used to reconstruct the Fourier coefficients of the unknown field \mathbf{u}_J [66]. The Tikhonov regularization parameter λ is selected to be 5×10^{-2} . We set the Fourier basis number n to 31. The total number of parameter for each layer is 1985, which makes the total number of parameters for all layers 3970. The total number of parameters in this example is greatly reduced compared to [67] where each pixel in the target image count as a parameter, that is, for our example it translates to 261121 parameters.

Figure 4 shows that compared to the FBP and Tikhonov regularization reconstructions, the posterior sample mean of our MCMC method resulted in an image with less streak artifact and noise. The features of the phantom appear much more clear compared to those on the FBP and Tikhonov reconstructions. As examples, examine the mouth parts, dark circles between eyes and at the forehead, and the two eyes are both relatively much more clear than the other two. Nonetheless, since we set n only 31, the edge of phantom face which has very high values is not fully recovered as much as the FBP reconstruction. The posterior mean produces a lower L^2 error (44.795) compared to the FBP reconstruction (47.156), but higher than Tikhonov regularization method, (42.145). However, it has a slightly higher PSNR, 22.246 compared to Tikhonov regularization method, 21.673. Ideally we could double n in our proposed Bayesian method to get a much better reconstruction. However, it is not possible to accomplish this within our current setup due to restrictions in the GPU memory. We can fairly conclude that the use of a shallow Gaussian field prior with $n = 31$ resulted in a highly reduced amount of artifact at the expense of light blur at the edge. As in the one dimensional example, adding another Gaussian field layer might help to increase the sharpness of the edges in the X-ray tomography application.

6. Conclusion

We have presented a multi-layered Gaussian-field Bayesian inversion using a Galerkin method. We rigorously showed that our approach enjoys a nice convergence-in-probability property to the weak solution of the forward model. We achieved this result by establishing an upper bound for the square root of the precision operator of the Gaussian field $L(u_{i-1})$. Then, we showed for the dimension of field $d \leq 3$ that there

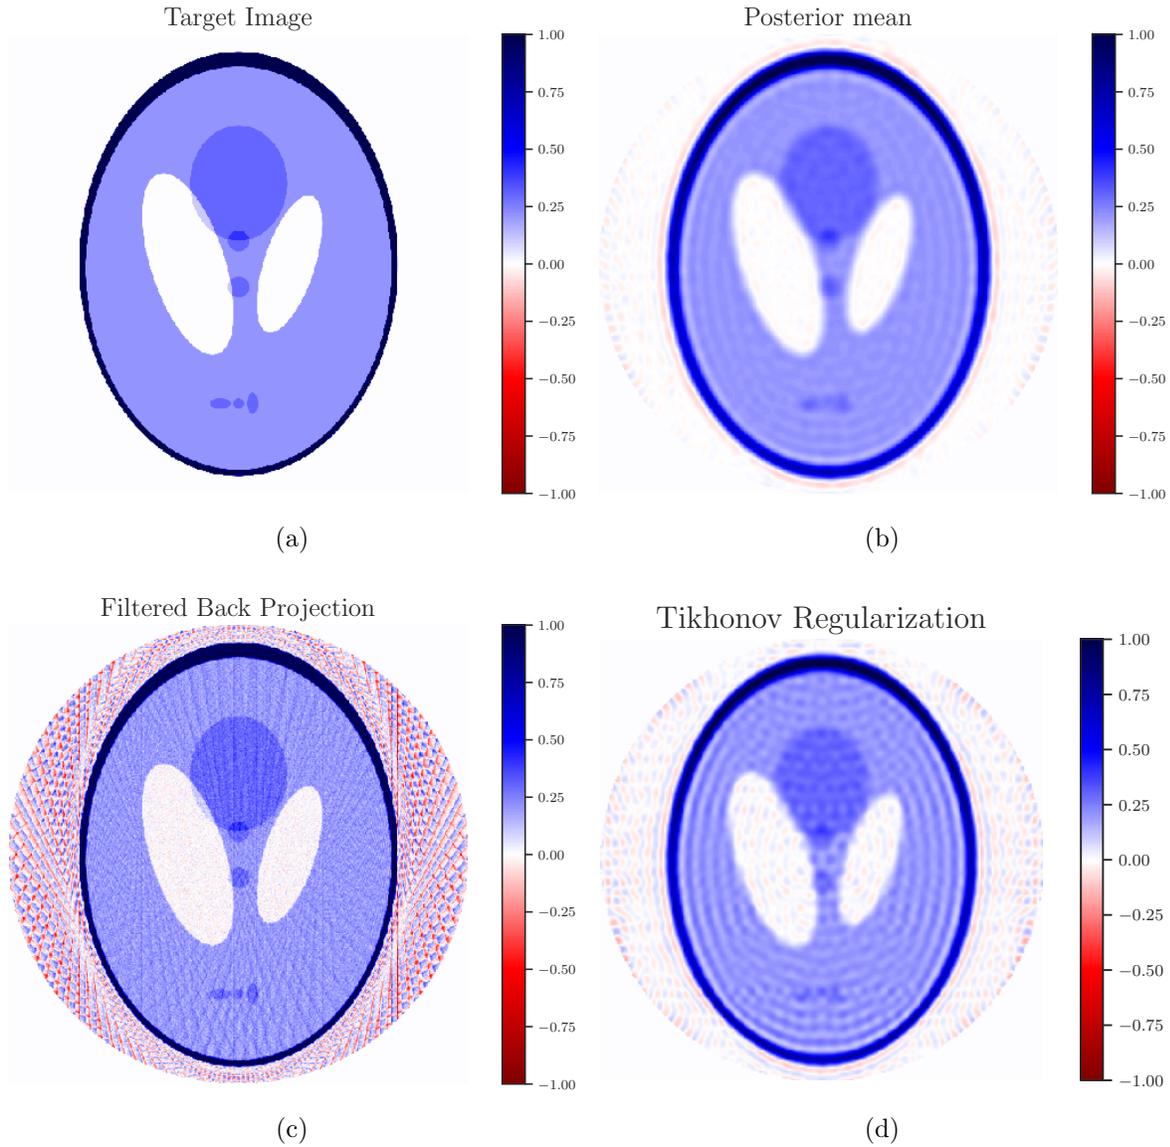


Figure 4. Simulation result of a shallow SPDE with $J = 1$ and $n = 31$, where the number of projections is 45. Figures 4b and 4d show posterior means of field v and the Tikhonov regularization result using $\lambda = 5 \times 10^{-2}$, while Figure 4c shows the FBP reconstruction.

is an upper bound for the error between $L(u_{i-1})$ and $L(u_{i-1}^N)$. The desired result was then shown via a uniform tightness result and Hölder continuity of the original weak solution. As a consequence, we can guarantee that the approximated prior and posterior probability distribution converges weakly to the original prior and posterior, respectively. This implies weak convergence of the joint posterior distribution of the Gaussian field, and hence gives an assurance that our proposed method is well defined and robust upon increasing the number of Fourier basis functions. Generalization of this result to the case $d > 3$ requires further technical care. To achieve a reduced computational time without sacrificing too much accuracy, in the future, it will also interesting to use the

proposed Bayesian inversion method in combination with variational Bayesian methods.

We showed, numerically, for one-dimensional denoising problem that we can achieve joint smoothing and edge-preserving inversion at the same time by using two hyperprior layers and non-centered version of the preconditioned Crank–Nicolson algorithm. For the X-ray tomography problem, with a single hyperprior layer and a very small number of Fourier basis ($n = 31$), the posterior sample mean of our proposed approach gives an image with less streak artifact and noise compared to the FBP and Tikhonov regularization reconstructions. Although traces of streak artefact and edge blurring still present, the L^2 error and PSNR of our proposed method sit in the middle of those from the FBP and Tikhonov regularization. Furthermore, adding another Gaussian field layer might help to increase the sharpness of the edge in the X-ray tomography application. One future outlook is to apply the method on a real data.

Acknowledgments

The authors would like to thank Academy of Finland for financial support (application numbers: 326240, 326341, 334816, 321891, 321900, and 314474).

References

- [1] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*. Springer, 2004.
- [2] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press Ltd, 2005.
- [3] M. J. Heaton, A. Datta, A. O. Finley, R. Furrer, J. Guinness, R. Guhaniyogi, F. Gerber, R. B. Gramacy, D. Hammerling, M. Katzfuss, F. Lindgren, D. W. Nychka, F. Sun, and A. Zammit-Mangion, “A case study competition among methods for analyzing large spatial data,” *J. Agr. Biol. Envir. St.*, vol. 24, pp. 398–425, dec 2018.
- [4] C. J. Paciorek and M. J. Schervish, “Nonstationary covariance functions for Gaussian process regression,” in *Adv. Neur. Inf. Proc. Sys.*, pp. 273–280, 2004.
- [5] G.-A. Fuglstad, D. Simpson, F. Lindgren, and H. Rue, “Does non-stationary spatial data always require non-stationary random fields?,” *Spat. Stat.*, vol. 14, pp. 505–531, nov 2015.
- [6] C. J. Paciorek, *Non-Stationary Gaussian processes for regression and spatial modelling*. PhD thesis, 2003.
- [7] E. Snelson, Z. Ghahramani, and C. E. Rasmussen, “Warped Gaussian processes,” in *Adv. Neur. Inf. Proc. Sys.*, pp. 337–344, 2004.
- [8] P. D. Sampson and P. Guttorp, “Nonparametric estimation of nonstationary spatial covariance structure,” *Journal of the American Statistical Association*, vol. 87, pp. 108–119, mar 1992.
- [9] P. J. Green, “Reversible jump Markov chain Monte Carlo computation and Bayesian model determination,” *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.
- [10] A. Ray and D. Myer, “Bayesian geophysical inversion with trans-dimensional Gaussian process machine learning,” *Geophysical Journal International*, vol. 217, pp. 1706–1726, feb 2019.
- [11] 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. Royal Stat. Soc. B*, vol. 73, no. 4, pp. 423–498, 2011.
- [12] L. Roininen, M. Girolami, S. Lasanen, and M. Markkanen, “Hyperpriors for Matérn fields with applications in Bayesian inversion,” *Inverse Probl. & Imaging*, vol. 13, no. 1, pp. 1–29, 2019.

- [13] K. Montterrubio-Gómez, L. Roininen, S. Wade, T. Damoulas, and M. Girolami, “Posterior inference for sparse hierarchical non-stationary models,” *Comp. Stat. & Data Anal.*, p. 106954, mar 2020.
- [14] A. C. Damianou and N. D. Lawrence, “Deep Gaussian processes,” *Artif. Intelli. & Stat.*, 2013.
- [15] D. Duvenaud, O. Rippel, R. P. Adams, and Z. Ghahramani, “Avoiding pathologies in very deep networks,” *Artif. Intelli. & Stat.*, 2014.
- [16] M. M. Dunlop, M. A. Girolami, A. M. Stuart, and A. L. Teckentrup, “How deep are deep Gaussian processes?,” *J. Mach. Learn. Res.*, vol. 19, no. 54, pp. 1–46, 2018.
- [17] A. M. Stuart, “Inverse problems: A Bayesian perspective,” *Acta Numerica*, vol. 19, pp. 451–559, 2010.
- [18] M. Dashti and A. M. Stuart, “The Bayesian approach to inverse problems,” in *Handbook of Uncertainty Quantification*, pp. 311–428, Springer International Publishing, 2017.
- [19] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, “MCMC methods for functions: Modifying old algorithms to make them faster,” *Statist. Sci.*, vol. 28, no. 3, pp. 424–446, 2013.
- [20] K. J. H. Law, “Proposals which speed up function-space MCMC,” *J. Comp. & Appl. Math.*, vol. 262, pp. 127–138, 2014.
- [21] A. Beskos, M. Girolami, S. Lan, P. E. Farrell, and A. M. Stuart, “Geometric MCMC for infinite-dimensional inverse problems,” *J. Comp. Phys.*, vol. 335, pp. 327–351, 2017.
- [22] D. Rudolf and B. Sprungk, “On generalization of the preconditioned Crank-Nicolson Metropolis algorithm,” *Found. Comp. Math.*, vol. 18, pp. 309–343, 2018.
- [23] M. Hairer, A. M. Stuart, and S. J. Vollmer, “Spectral gaps for a Metropolis-Hastings algorithm in infinite dimensions,” *Ann. Appl. Probab.*, vol. 24, pp. 2455–2490, dec 2014.
- [24] V. Chen, M. M. Dunlop, O. Papaspiliopoulos, and A. M. Stuart, “Dimension-robust MCMC in Bayesian inverse problems,” *arXiv*, 2018.
- [25] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld, “A general framework for the parametrization of hierarchical models,” *Stat. Sci.*, vol. 22, pp. 59–73, feb 2007.
- [26] Z. Hu, Z. Yao, and J. Li, “On an adaptive preconditioned Crank-Nicolson MCMC algorithm for infinite dimensional bayesian inference,” *J. Comp. Phys.*, vol. 332, pp. 492–503, mar 2017.
- [27] J. M. Bernardo and A. F. M. Smith, *Bayesian Theory*. John Wiley & Sons, 2000.
- [28] M. Emzir, S. Lasanen, Z. Purisha, and S. Särkkä, “Hilbert-space reduced-rank methods for deep Gaussian processes,” in *Proceeding of IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2019.
- [29] A. Solin and S. Särkkä, “Hilbert space methods for reduced-rank Gaussian process regression,” *Stat. Comput.*, vol. 30, pp. 419–446, aug 2019.
- [30] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *Journal of the American Statistical Association*, vol. 112, pp. 859–877, feb 2017.
- [31] M. J. Beal, *Variational algorithms for approximate Bayesian inference*. PhD thesis, university of London London, 2003.
- [32] Y. Wang and D. M. Blei, “Frequentist consistency of variational Bayes,” *Journal of the American Statistical Association*, vol. 114, pp. 1147–1161, aug 2018.
- [33] F. J. Pinski, G. Simpson, A. M. Stuart, and H. Weber, “Algorithms for Kullback–Leibler approximation of probability measures in infinite dimensions,” *SIAM Journal on Scientific Computing*, vol. 37, pp. A2733–A2757, jan 2015.
- [34] F. J. Pinski, G. Simpson, A. M. Stuart, and H. Weber, “Kullback–Leibler approximation for probability measures on infinite dimensional spaces,” *SIAM Journal on Mathematical Analysis*, vol. 47, pp. 4091–4122, jan 2015.
- [35] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, “Solving inverse problems using data-driven models,” *Acta Numerica*, vol. 28, pp. 1–174, may 2019.
- [36] M. A. Nawaz, *Efficient probabilistic inversion of geophysical data*. PhD thesis, 2019.
- [37] M. Lassas and S. Siltanen, “Can one use total variation prior for edge-preserving Bayesian inversion?,” *Inverse Problems*, vol. 20, pp. 1537–1563, aug 2004.
- [38] J. Kaipio and E. Somersalo, “Statistical inverse problems: Discretization, model reduction and

- inverse crimes,” *J. Comp. & Appl. Math.*, vol. 198, pp. 493–504, jan 2007.
- [39] F. Natterer, *The Mathematics of Computerized Tomography (Classics in Applied Mathematics)*. SIAM: Society for Industrial and Applied Mathematics, 2001.
- [40] A. Tarantola, *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics, 2005.
- [41] D. Li, J. Svensson, H. Thomsen, F. Medina, A. Werner, and R. Wolf, “Bayesian soft X-ray tomography using non-stationary Gaussian processes,” *Rev. Sci. Inst.*, vol. 84, p. 083506, aug 2013.
- [42] C. Plagemann, K. Kersting, and W. Burgard, “Nonstationary Gaussian process regression using point estimates of local smoothness,” in *Machine Learning and Knowledge Discovery in Databases European Conference, ECML PKDD 2008 Antwerp, Belgium, September 15-19, 2008 Proceedings, Part II*, 2008.
- [43] Z. Purisha, C. Jidling, N. Wahlström, T. B. Schön, and S. Särkkä, “Probabilistic approach to limited-data computed tomography reconstruction,” *Inverse Probl.*, vol. 35, p. 105004, sep 2019.
- [44] T. Frese, C. Bouman, and K. Sauer, “Adaptive wavelet graph model for Bayesian tomographic reconstruction,” *IEEE Trans. Imag. Proc.*, vol. 11, pp. 756–770, jul 2002.
- [45] Y. Chen, Y. Li, W. Yu, L. Luo, W. Chen, and C. Toumoulin, “Joint-MAP tomographic reconstruction with patch similarity based mixture prior model,” *Multiscale Modeling & Simulation*, vol. 9, pp. 1399–1419, oct 2011.
- [46] V. Antun, F. Renna, C. Poon, B. Adcock, and A. C. Hansen, “On instabilities of deep learning in image reconstruction and the potential costs of AI,” *PNAS*, p. 201907377, may 2020.
- [47] L. C. Evans, *An Introduction to Stochastic Differential Equations*. American Mathematical Society, 2014.
- [48] S. Lasanen, L. Roininen, and J. M. Huttunen, “Elliptic boundary value problems with Gaussian white noise loads,” *Stoch. Proc. & Appl.*, vol. 128, no. 11, pp. 3607–3627, 2018.
- [49] S. C. Brenner and L. R. Scott, *The mathematical theory of finite element methods*, vol. 15 of *Texts in Applied Mathematics*. Springer, New York, third ed., 2008.
- [50] V. I. Bogachev, *Weak convergence of measures*, vol. 234 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2018.
- [51] O. Kallenberg, *Foundations of modern probability*. Probability and its Applications (New York), Springer-Verlag, New York, second ed., 2002.
- [52] W. Rudin, *Principles of Mathematical Analysis*, vol. 39 of *International series in pure and applied mathematics*. McGraw-Hill, 1976.
- [53] M. B. Marcus and J. Rosen, *Markov processes, Gaussian processes, and local times*, vol. 100 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2006.
- [54] V. I. Bogachev, *Gaussian measures*, vol. 62 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1998.
- [55] S. Lasanen, “Non-Gaussian statistical inverse problems. Part II: Posterior convergence for approximated unknowns,” *Inverse Probl. Imaging*, vol. 6, no. 2, pp. 267–287, 2012.
- [56] K. Dong, D. Eriksson, H. Nickisch, D. Bindel, and A. G. Wilson, “Scalable log determinants for Gaussian process kernel learning,” in *Advances in Neural Information Processing Systems*, pp. 6327–6337, 2017.
- [57] Y. Yu and X.-L. Meng, “To center or not to center: That is not the question- an ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC efficiency,” *J. Comp. & Graph. Stat.*, vol. 20, pp. 531–570, jan 2011.
- [58] H. F. L. Dani Gamerman, *Markov Chain Monte Carlo*. Taylor & Francis Inc, 2006.
- [59] H. Wendland, *Numerical Linear Algebra*. Cambridge University Press, nov 2017.
- [60] T. Cui, K. J. Law, and Y. M. Marzouk, “Dimension-independent likelihood-informed MCMC,” *J. Comp. Phys.*, vol. 304, pp. 109–137, jan 2016.
- [61] T. Buzug, *Computed Tomography*. Springer-Verlag GmbH, 2008.
- [62] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Society for Industrial

- and Applied Mathematics, jan 2001.
- [63] S. R. Deans, *The Radon transform and some of its applications*. Wiley, 1983.
 - [64] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, “CuPy: A NumPy-compatible library for NVIDIA GPU calculations,” in *Proceeding MM '17 Proceedings of the 25th ACM international conference on Multimedia*, pp. 1217–1220, 2017.
 - [65] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, “Scikit-image: image processing in Python,” *PeerJ*, vol. 2, p. e453, jun 2014.
 - [66] J. Mueller and S. Siltanen, *Linear and nonlinear inverse problems with practical applications*. Philadelphia: Society for Industrial and Applied Mathematics, 2012.
 - [67] J. Suuronen, M. Emzir, S. Lasanen, S. Särkkä, and L. Roininen, “Enhancing industrial X-ray tomography by data-centric statistical methods,” *Data-Centric Engineering*, 2020. (to appear).