Dimension-independent likelihood-informed MCMC

Tiangang Cui^a, Kody J. H. Law^b, Youssef M. Marzouk^a

^a*Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

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

Abstract

Many Bayesian inference problems require exploring the posterior distribution of high-dimensional parameters that represent the discretization of an underlying function. This work introduces a family of Markov chain Monte Carlo (MCMC) samplers that can adapt to the particular structure of a posterior distribution over functions. Two distinct lines of research intersect in the methods developed here. First, we introduce a general class of operator-weighted proposal distributions that are well defined on function space, such that the performance of the resulting MCMC samplers is independent of the discretization of the function. Second, by exploiting local Hessian information and any associated low-dimensional structure in the change from prior to posterior distributions, we develop an inhomogeneous discretization scheme for the Langevin stochastic differential equation that yields operator-weighted proposals adapted to the non-Gaussian structure of the posterior. The resulting *dimension-independent* and *likelihood-informed* (DILI) MCMC samplers may be useful for a large class of high-dimensional problems where the target probability measure has a density with respect to a Gaussian reference measure. Two nonlinear inverse problems are used to demonstrate the efficiency of these DILI samplers: an elliptic PDE coefficient inverse problem and path reconstruction in a conditioned diffusion.

Keywords: Markov chain Monte Carlo, likelihood-informed subspace, infinite-dimensional inverse problems, Langevin SDE, conditioned diffusion

1. Introduction

Many Bayesian inference problems require sampling a posterior distribution of high-dimensional parameters that represent the discretization of some underlying *function*. Examples include inverse problems governed by partial differential equations (PDEs) [1, 2, 3] and path reconstructions of stochastic differential equations [4, 5]. Yet sampling the posterior with standard Markov chain Monte Carlo (MCMC) algorithms can become arbitrarily slow as the representation of the unknown function is refined; the convergence rates of such algorithms typically degrade with increasing parameter dimension [6, 7, 8, 9]. Here we present a *dimension-independent* and *likelihood-informed* MCMC sampling framework that enables efficient exploration of posterior distributions over functions. Our framework builds on recent work in dimension-independent MCMC sampling by identifying changes in the local curvature of the log-posterior, relative to the log-prior, and using this

Email addresses: `tcui@mit.edu` (Tiangang Cui), `kody.law@kaust.edu.sa` (Kody J. H. Law), `ymarz@mit.edu` (Youssef M. Marzouk)

information to design more effective proposal distributions. We cast these proposals within a general class of operator-weighted proposals that guarantee dimension independence while allowing flexibility in the design of structure-exploiting MCMC algorithms.

Previous work [10, 11] has shown that it is possible to overcome degeneration of the convergence rates of standard MCMC algorithms by formulating them in function space. In particular, [10, 11] consider a posterior that is absolutely continuous with respect to a Gaussian reference measure, typically chosen to be the prior. By deriving proposals that yield valid Markov transition kernels on a function space, these efforts yield algorithms with performance that is *dimension independent*, i.e., robust under mesh refinement. Recent work has begun to integrate this function space perspective with other state-of-the-art sampling approaches such as sequential Monte Carlo [12, 13] and particle MCMC [14] (which were originally designed for finite-dimensional parameters) along with multi-level Monte Carlo [15]. Examples include [16], which uses function space proposals inspired by those in [10, 11] as building blocks within a sequential Monte Carlo scheme, and [17, 18], which employ function space algorithms within a multi-level MCMC scheme for variance reduction. An important takeaway from these efforts is that although proposals that are well-defined on function space provide mixing rates that are independent of discretization, they do not guarantee good mixing in absolute terms; additional analysis and exploitation of the structure of the posterior is essential.

To improve MCMC mixing in finite-dimensional settings, another important line of research has focused on using information about the local geometry of the posterior distribution to scale and guide proposals. Examples include the stochastic Newton method [19], which uses low-rank approximations of the Hessian of the log-posterior to build Gaussian proposals, and Riemannian manifold MCMC [20], which uses the Fisher information matrix (and the prior precision matrix) to construct a local metric for both Langevin proposals and Hamiltonian dynamics. Implicit in the low-rank approximations of [19] is the idea that the Hessian of the log-likelihood can be compared with the prior covariance to identify directions in parameter space along which the posterior distribution differs most strongly from the prior. This idea is formalized for linear inverse problems in [21]; at its heart lies the solution of a generalized eigenproblem involving both operators. In linear-Gaussian problems, a transformation of the resulting dominant eigendirections reveals the directions of greatest *variance reduction*, relative to the prior. Solutions of the eigenproblem can also be used to construct optimal approximations of the posterior covariance, in the form of low-rank negative semidefinite *updates* of the prior covariance [21]. The form of these approximations is crucial, as it suggests the following tie to dimension-independent algorithms: The proposals used in [10, 11] rely on sampling from the Gaussian prior measure on an infinite-dimensional space. It is natural to alter these proposals only in a *finite* number of dimensions—where the posterior differs most strongly from the prior—to achieve better mixing while preserving dimension independence.

A first attempt in this direction is the operator-weighted proposal of [22], which employs a relaxed version of the log-likelihood Hessian, evaluated at a single point near the posterior mode, to rescale the preconditioned Crank-Nicolson proposal of [11]. Of course, in inverse problems with non-linear forward operators or non-Gaussian noise, the log-likelihood Hessian varies over the parameter space. Nonetheless, if local Hessians have low rank and dominant eigenspaces that are somewhat aligned, changes from the prior to the posterior will be confined to a relatively low-dimensional subspace of the parameter space. We will call this subspace the “likelihood-informed” subspace (LIS). [23] introduced the notion of the LIS and constructed it from the dominant eigenspace of the posterior expectation of a preconditioned log-likelihood Hessian. There, the LIS was used to *approx-*

imate the posterior distribution in finite-dimensional nonlinear inverse problems, by replacing the posterior in complementary directions with the prior. In inverse problems with limited observations and smoothing forward operators, the LIS is relatively low dimensional and such approximations can be quite accurate.

We will use the notion of an LIS in the present work, but extend it to the infinite-dimensional setting and seek exact sampling; in this context, the LIS can be seen as means of convergence acceleration that captures key features of the posterior. The goal of this paper is to design MCMC algorithms that are dimension independent, that sample from the exact posterior distribution, and that exploit a posterior structure wherein departures from the prior—including non-Gaussianity—are concentrated on a finite number of directions. To realize these goals, we introduce a general class of proposal distributions that simultaneously guarantee dimension-independent MCMC sampling and allow the essential structure of the posterior (i.e., departures from the Gaussian prior) to be captured in a flexible manner, via both local gradient information and a set of global covariance-like (i.e., bounded and self-adjoint) operators. We then discuss particular instances of this class for which the operators are derived from the LIS. The proposals themselves can be obtained from an underlying Langevin SDE by applying an inhomogeneous time discretization scheme on the LIS and its complement. We then incorporate these proposals into MCMC either as Metropolis or Metropolis-within-Gibbs updates.

The rest of this paper is organized as follows. Section 2 reviews some theoretical background on the infinite-dimensional Bayesian formulation of inverse problems and function-space MCMC methods, as well as previous work aimed at improving the efficiency of these samplers. Section 3 introduces a general class of operator-weighted proposals that enables the design of dimension-independent and likelihood-informed (DILI) samplers. We then discuss the construction of a global LIS and describe several different DILI MCMC samplers based on the LIS. We also recall a few non-DILI algorithms that will subsequently be used for benchmarking. In Section 4, an elliptic PDE inverse problem is used to evaluate the performance of the various samplers introduced in Section 3. In Section 5, we apply the same samplers to the problem of path reconstruction in a conditioned diffusion and highlight key differences in performance. Section 6 offers some brief discussion and concluding remarks.

2. Theoretical background

In this section we recall the Bayesian formulation of inverse problems in the function space setting and provide some background on the relevant MCMC methods. We then review preconditioned Crank-Nicolson (pCN) proposals, which yield dimension-independent MCMC algorithms for posterior distributions over function space equipped with an appropriate Gaussian prior. We close this section with a discussion of how local geometry information can be used to inform MCMC proposals.

2.1. Bayesian inference framework

Suppose that the parameter of interest is some function $u \in \mathcal{H}$, where $\mathcal{H}$ is a separable Hilbert space. Suppose further that the prior distribution μ_0 is such that $\mu_0(\mathcal{H}) = 1$. We assume that the observation y is in $\mathbb{R}^{N_y}$. Let $\mathcal{L}(u; y)$ denote the likelihood function of y given u , which is assumed to be bounded away from zero μ_0 -almost surely. Let μ_y denote the posterior probability measure of the unknown u conditioned on the observation y . The posterior distribution on infinitesimal

volume elements du in $\mathcal{H}$ is given by

$$\mu_y(du) \propto \mathcal{L}(u; y) \mu_0(du). \quad (1)$$

Let $F : \mathcal{X} \rightarrow \mathbb{R}^{N_y}$ denote the forward operator, defined on some Banach space $\mathcal{X} \subseteq \mathcal{H}$ with $\mu_0(\mathcal{X}) = 1$, and assume the noise on the observation is additive and Gaussian with zero mean. Then the observation model has the form

$$y = F(u) + e, \quad e \sim \mathcal{N}(0, \Gamma_{\text{obs}}).$$

We define the data-misfit functional

$$\eta(u; y) = \frac{1}{2} (F(u) - y)^\top \Gamma_{\text{obs}}^{-1} (F(u) - y). \quad (2)$$

The likelihood function can then be written as

$$\mathcal{L}(u; y) \propto \exp(-\eta(u; y)). \quad (3)$$

The inner products on the data space $\mathbb{R}^{N_y}$ and the parameter space $\mathcal{H}$ are denoted $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle_{\mathcal{H}}$, respectively, with associated norms denoted by $\| \cdot \|$ and $\| \cdot \|_{\mathcal{H}}$. For brevity, where misinterpretation is not possible, we will drop the subscripts $\mathcal{H}$. Throughout this paper we adopt the following assumptions about the forward operator:

Assumption 2.1. *The forward operator $F : \mathcal{X} \rightarrow \mathbb{R}^{N_y}$ satisfies the following:*

1. *For all $\epsilon_1 > 0$, there exists a constant $K_1(\epsilon_1) > 0$ such that, for all $u \in \mathcal{X}$,*

$$\|F(u)\| \leq \exp(K_1(\epsilon_1) + \epsilon_1 \|u\|_{\mathcal{H}}^2).$$

2. *For all $\epsilon_2 > 0$ there exists a constant $K_2(\epsilon_2) > 0$ such that, for all $u_1, u_2 \in \mathcal{X}$ with $\|u_1\|_{\mathcal{H}} < \epsilon_2$ and $\|u_2\|_{\mathcal{H}} < \epsilon_2$,*

$$\|F(u_1) - F(u_2)\| \leq K_2(\epsilon_2) \|u_1 - u_2\|_{\mathcal{H}}.$$

3. *For all $u \in \mathcal{X}$, there exists a bounded linear operator $J : \mathcal{X} \rightarrow \mathbb{R}^{N_y}$ such that for all $\delta u \in \mathcal{X}$*

$$\lim_{\delta u \rightarrow 0} \frac{\|F(u + \delta u) - F(u) - J(u)\delta u\|}{\|\delta u\|_{\mathcal{H}}} = 0.$$

Given observations y such that $\|y\| < \infty$ and a forward operator that satisfies the first two conditions in Assumption 2.1, [3] shows that the resulting data-misfit function is sufficiently bounded and locally Lipschitz, and that the posterior measure is dominated by the prior measure. The third condition in Assumption 2.1 states that the forward model is first order Fréchet differentiable, and hence the Gauss-Newton approximation of the Hessian of the data-misfit functional is bounded.

We assume that the prior measure is Gaussian with mean $m_0 \in \mathcal{H}$ and a covariance operator Γ_{pr} on $\mathcal{H}$ that is self-adjoint, positive definite, and trace-class, so that the prior provides a full probability measure on $\mathcal{H}$. The covariance operator Γ_{pr} also defines the inner product $\langle \cdot, \cdot \rangle_{\Gamma_{\text{pr}}} = \langle \Gamma_{\text{pr}}^{-1/2} \cdot, \Gamma_{\text{pr}}^{-1/2} \cdot \rangle$ and the associated norm $\| \cdot \|_{\Gamma_{\text{pr}}}$ on the Hilbert space $\mathcal{H}$. It is convenient to define the covariance operator through its inverse Γ_{pr}^{-1} , referred to as the prior precision and often chosen to be a Laplace-like differential operator [3].

We note that the posterior distribution μ_y does not admit a density with respect to Lebesgue measure. Instead of a “posterior density,” we can define the limiting ratio of the probabilities of shrinking balls at two points u_1 and u_2 in $\text{Im}(\Gamma_{\text{pr}}^{1/2})$ by the exponential of the difference of the Onsager-Machlup functional (OMF),

$$\eta(u; y) + \|u - m_0\|_{\Gamma_{\text{pr}}}^2, \quad (4)$$

evaluated at u_1 and u_2 . Then the maximizer of the probability over vanishing balls, the *maximum a posteriori* (MAP) estimator, is the minimizer of this functional,¹ as shown in [24] for problems like those considered here.

Here we assume that the data y are finite dimensional. For suitable combinations of forward model and observational noise, this setting can be generalized to problems with infinite-dimensional data sets [3, 24].

2.2. MCMC on function space

MCMC methods, in particular the Metropolis-Hastings (MH) algorithm [25, 26], can be employed to sample the posterior distribution μ_y over $\mathcal{H}$. The MH algorithm defines a Markov chain by prescribing an accept-reject move based on a proposal distribution q . One step of the MH algorithm is given as follows:

Definition 2.2 (Metropolis-Hastings kernel). *Given the current state u_k , a candidate state u' is drawn from the proposal distribution $q(u_k, \cdot)$. Define the pair of measures*

$$\begin{aligned} \nu(du, du') &= q(u, du')\mu(du) \\ \nu^\perp(du, du') &= q(u', du)\mu(du'). \end{aligned} \quad (5)$$

Then the next state of the Markov chain is set to $u_{k+1} = u'$ with probability

$$\alpha(u_k, u') = \min \left\{ 1, \frac{d\nu^\perp}{d\nu}(u_k, u') \right\}; \quad (6)$$

otherwise, with probability $1 - \alpha(u_k, u')$, it is set to $u_{k+1} = u_k$.

The MH algorithm requires the absolute continuity condition $\nu^\perp \ll \nu$ to define a valid acceptance probability and hence a valid transition kernel [27]. We will refer to a MH algorithm as **well-defined** if this absolute continuity condition is satisfied. Most of the MCMC proposals used in current literature—e.g., vanilla random walks, Metropolis adjusted Langevin (MALA) [28], simplified manifold MALA [20]—violate this condition in the limit $\dim(\mathcal{H}) \rightarrow \infty$. For these proposals, the MH algorithm is defined only in finite dimensions [11] and the convergence rate degenerates under mesh refinement. Extensive investigations have been carried out to study the rate of degeneration of various proposals and optimal scaling relative to this. See [6, 29, 7] and references therein.

For target probability measures over function space, the sequence of papers [10, 11, 30, 3] provide a viable way of constructing well-defined MH algorithms using appropriate discretizations of Langevin SDEs. The key step in the construction of such algorithms is to ensure the absolute continuity condition $\nu^\perp \ll \nu$. We will revisit these constructions in the next subsection.

¹The functional is extended to $\mathcal{H}$ by prescribing it to take the value $+\infty$ outside of $\text{Im}(\Gamma_{\text{pr}}^{1/2})$.

2.3. Preconditioned semi-implicit proposal

Following [10, 11, 30], we recall the preconditioned Crank-Nicolson (pCN) proposal. Consider the Langevin SDE

$$du = -P \left(\Gamma_{\text{pr}}^{-1}(u - m_0) + \gamma \nabla_u \eta(u; y) \right) d\tau + \sqrt{2P} d\mathcal{W}(\tau), \quad (7)$$

where P is a symmetric, positive-definite, and bounded operator and $d\mathcal{W}(\tau)$ is space-time white noise over $\mathcal{H} \times \mathbb{R}_+$. For any such P , the SDE (7) has invariant distribution μ_y for $\gamma = 1$ and invariant distribution μ_0 for $\gamma = 0$; see [11] and references therein.

Following [10, 11] we let $P = \Gamma_{\text{pr}}$. A first order semi-implicit discretization of (7), parameterized by $\vartheta \in (0, 1]$, defines a family of proposal distributions:

$$u' = \frac{1 - (1 - \vartheta)\Delta\tau}{1 + \vartheta\Delta\tau}(u - m_0) + m_0 - \frac{\gamma\Delta\tau}{1 + \vartheta\Delta\tau}\Gamma_{\text{pr}}\nabla_u \eta(u; y) + \frac{\sqrt{2\Delta\tau}}{1 + \vartheta\Delta\tau}\Gamma_{\text{pr}}^{1/2}\xi, \quad (8)$$

where $\xi \sim \mathcal{N}(0, I)$, $\Delta\tau$ is the time step, and $\gamma = \{0, 1\}$ is a tuning parameter to switch between the full Langevin proposal ($\gamma = 1$) and the random walk (Ornstein-Uhlenbeck) proposal ($\gamma = 0$). Simulating the Langevin SDE thus provides a method for approximately sampling from μ_y when $\gamma = 1$ or from μ_0 when $\gamma = 0$. Under certain technical conditions, Theorem 4.1 of [10] shows that $\vartheta = 1/2$ (a Crank-Nicolson scheme) is the unique choice that delivers well-defined MCMC algorithms when $\dim(\mathcal{H}) \rightarrow \infty$. Putting $\vartheta = 1/2$ and letting $a = (2 - \Delta\tau)/(2 + \Delta\tau)$, (8) can be rewritten as

$$u' = a(u - m_0) + m_0 - \gamma(1 - a)\Gamma_{\text{pr}}\nabla_u \eta(u; y) + \sqrt{1 - a^2}\Gamma_{\text{pr}}^{1/2}\xi, \quad (9)$$

where it is required that $a \in (-1, 1)$. If $\gamma = 0$, the invariant distribution of (9) remains μ_0 . However, if $\gamma = 1$, the invariant distribution of the discretized system is different from the invariant distribution of (7). To remove the bias, one can employ the Metropolis correction, as discussed by [28].

The pCN proposal (9) has the desired dimension independence property, in that the autocorrelation of the resulting samples does not increase as the discretization of u is refined. Since the scale of the posterior distribution necessarily tightens, relative to the prior, along directions in the parameter space that are informed by the data, maintaining a reasonable acceptance probability requires that the step size used in the proposal (9) be dictated by these likelihood-informed directions. As shown in [22], the Markov chain produced by MH with the standard pCN proposal then decorrelates more quickly in the parameter subspace that is data-informed than in the subspace that is dominated by the prior. The proposed moves in the prior-dominated directions are effectively too small or conservative, resulting in poor mixing. The underlying issue is that the scale of the proposal (9) is uniform in all directions with respect to the norm induced by the prior covariance, and hence it does not adapt to the posterior scaling induced by the likelihood function.

It is therefore desirable to design proposals that adapt to this anisotropic structure of the posterior while retaining dimension independence. A first example is the operator-weighted proposal of [22], which specifically (though without loss of generality) targets a white-noise prior, $m_0 = 0$ and $\Gamma_{\text{pr}} = I$:

$$u' = Au + (I - A^2)^{1/2}\xi. \quad (10)$$

Some intuition for this proposal can be had by comparing to (9) in the case $\gamma = 0$: the new proposal replaces the scalar a with a symmetric operator A . This operator is constructed from a regularized

approximation of the Hessian of the OMF (4) at a single posterior sample, chosen close to or at the MAP estimate. The dimension independence property of the proposal (10) can be seen simply by noticing its reversibility with respect to the prior,

$$q(u, du')\mu_0(du) = q(u', du)\mu_0(du'),$$

which implies

$$\nu(du, du') = \exp\{\eta(u; y) - \eta(u'; y)\}\nu^\perp(du, du')$$

and hence the absolute continuity condition $\nu^\perp \ll \nu$ (given that Conditions (1) and (2) of Assumption 2.1 also hold). In comparison with (9), the proposal (10) produces direction-dependent step sizes and, in particular, different step sizes for parameter directions that are prior-dominated than for directions that are influenced by the likelihood—insofar as this anisotropy is captured by A. The autocorrelation time of the resulting MH chain can thus be substantially smaller than that of a chain using the original pCN proposal [22].

Despite these performance improvements, this simple operator-weighted proposal has several limitations, especially when compared to state-of-the-art MCMC approaches for finite-dimensional problems. First, we note that A only uses information from a single point in the posterior, which may not be representative of posterior structure in a non-Gaussian problem. Also, (10) uses no gradient information and more broadly no local information, e.g., to describe the geometry of the target measure around each u . Also, the particular form of A used above does not necessarily identify directions of greatest *difference* between posterior and the prior or reference measure. (This point will be explained more fully in Section 3.2.) A more general class of operator-weighted proposals, designed in part to address these limitations, will be presented in Section 3.

2.4. Explicit proposals using local geometry

Before describing our general class of operator-weighted proposals, we briefly recall the simplified manifold MALA [20] and stochastic Newton [19] algorithms, which are related to an explicit discretization of the Langevin SDE (7). Both are well-established MCMC methods that use Hessian or Hessian-related information from the posterior distribution to achieve more efficient sampling. Simplified manifold MALA uses the proposal:

$$u' = u - P(u) (\nabla_u \eta(u; y) + \Gamma_{\text{pr}}^{-1} u) \Delta\tau + \sqrt{2\Delta\tau} P(u)^{1/2} \xi, \quad (11)$$

where the preconditioning operator $P(u)$ is derived from the expected Fisher information. Stochastic Newton prescribes $\Delta\tau = 1$ and discards the factor of $\sqrt{2}$ above, which gives the proposal

$$u' = u - P(u) (\nabla_u \eta(u; y) + \Gamma_{\text{pr}}^{-1} u) + P(u)^{1/2} \xi, \quad (12)$$

where the preconditioning operator $P(u)$ is the inverse of the posterior Hessian (or a low-rank approximation thereof; see [19, 21]) at the current sample u . We note that simplified manifold MALA and stochastic Newton are closely related for a likelihood function with additive Gaussian noise, as the expected Fisher information is equivalent to the Gauss-Newton approximation of the data-misfit Hessian in this case. Even though operations using the local preconditioner can be computationally expensive, both methods provide significant reductions in autocorrelation time per MCMC sample over standard random walks and even non-preconditioned MALA algorithms.

In a sense, the Hessian information—or more generally, the local Riemannian metric [20]—used by these proposals provides two useful insights into the structure of the posterior distribution.

First is an approximation of the local geometry, e.g., relative scaling along different directions in the parameter space, curvature, etc. Second, however, is insight into how the posterior distribution locally *differs from a base measure*, in particular, the prior distribution. When the Hessian of the data-misfit functional is compact, departures from the prior will be limited to a finite number of directions. One can explicitly identify these directions and partition the parameter space accordingly, into a finite-dimensional subspace and its complement. Of course, this particular partition is based only on local information, so one must consider methods for globalizing it—effectively incorporating the dependence of the Hessian on u . Given such an extension, this partitioning suggests a useful way of designing efficient MCMC proposals in infinite dimensions: use any effective MCMC proposal on a finite-dimensional subspace, and revert to a simpler function-space sampling framework on the infinite-dimensional complement of this subspace, thus yielding a well-defined MCMC algorithm in infinite dimensions and ensuring discretization-invariant mixing. The next section will explain how to realize this idea in the context of a more general operator-weighted MCMC.

3. Likelihood-informed MCMC

In this section, we first introduce a general class of operator-weighted MCMC proposals that guarantee dimension-independent sampling; within this class, there is considerable flexibility to capture the structure of the posterior by choosing the associated operators appropriately. We then take a detour to discuss the notion of a likelihood-informed subspace (LIS), which describes the parameter subspace where the posterior distribution departs most strongly from the prior. Based on the LIS, we derive a low-rank posterior covariance approximation scheme and then use both of these building blocks to guide the inhomogeneous discretization of a preconditioned Langevin SDE. This discretization naturally leads to the construction of various operator-weighted proposals that fall within our general class. We provide several specific examples of such proposals and then describe an adaptive posterior sampling framework that combines LIS construction, low-rank posterior covariance approximation, and automatic proposal adaptation.

3.1. Operator-weighted proposals

We consider a general class of operator-weighted proposals from u to u' given by

$$u' = \Gamma_{\text{pr}}^{1/2} A \Gamma_{\text{pr}}^{-1/2} (u - u_{\text{ref}}) + u_{\text{ref}} - \Gamma_{\text{pr}}^{1/2} G \Gamma_{\text{pr}}^{1/2} \nabla_u \eta(u; y) + \Gamma_{\text{pr}}^{1/2} B \xi, \quad (13)$$

where A , B , and G are bounded self-adjoint operators on $\text{Im}(\Gamma_{\text{pr}}^{-1/2})$, $\xi \sim \mathcal{N}(0, I)$, $u_{\text{ref}} = m_0 + m_{\text{ref}}$ with $m_{\text{ref}} \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$, and Γ_{pr} is the prior covariance operator defined in Section 2.1. Suppose further that these operators are defined through a complete orthonormal system (ψ_i) in $\text{Im}(\Gamma_{\text{pr}}^{-1/2})$ and sequences (a_i) , (b_i) , and (g_i) of positive numbers such that

$$A\psi_i = a_i\psi_i, \quad B\psi_i = b_i\psi_i, \quad G\psi_i = g_i\psi_i, \quad i \in \mathbb{N}.$$

Using matrix notation, let $\Psi = [\psi_1, \psi_2, \dots]$ and $\Psi^* = [\psi_1, \psi_2, \dots]^*$, where we define

$$\Psi^* v \equiv [\langle \psi_1, v \rangle, \langle \psi_2, v \rangle, \dots]^*.$$

Defining the diagonal operators $D_{A,ij} = \delta_{ij}a_i$, $D_{B,ij} = \delta_{ij}b_i$, and $D_{G,ij} = \delta_{ij}g_i$, we can represent the spectral decompositions of the operators above as $A = \Psi D_A \Psi^*$, $B = \Psi D_B \Psi^*$, and $G = \Psi D_G \Psi^*$, respectively. The following theorem establishes conditions on the operators A , B , and G such that the proposal (13) yields a well-defined Metropolis-Hastings algorithm for probability measures over function space.

Theorem 3.1. Suppose the self-adjoint operators A , B , and G are defined by a common set of eigenfunctions $\psi_i \in \text{Im}(\Gamma_{\text{pr}}^{-1/2})$, $i \in \mathbb{N}$, and sequences of eigenvalues (a_i) , (b_i) , and (g_i) , respectively. Suppose further that the posterior measure μ_y is equivalent to the prior measure $\mu_0 = \mathcal{N}(m_0, \Gamma_{\text{pr}})$. The proposal (13) delivers a well-defined MCMC algorithm if the following conditions hold:

1. All the eigenvalues (a_i) , (b_i) , and (g_i) are real and bounded, i.e., $a_i, b_i, g_i \in \mathbb{R}$, and there exists some $C < \infty$ such that $|a_i|, |b_i|, |g_i| < C$, $\forall i \in \mathbb{N}$.
2. For any index $i \in \mathbb{N}$ such that $b_i = 0$, we have $a_i = 1$ and $g_i = 0$.
3. For any index $i \in \mathbb{N}$ such that $b_i \neq 0$ we have $|b_i| \geq c$ for some $c > 0$.
4. The sequences of eigenvalues (a_i) and (b_i) satisfy

$$\sum_{i=1}^{\infty} (a_i^2 + b_i^2 - 1)^2 < \infty.$$

Proof. See Appendix B for the detailed proof. $\square$

We note that Condition (2) of Theorem 3.1 is designed for situations where the parameter u is partially updated, conditioned on those components corresponding to $b_i = 0$. The other conditions in Theorem 3.1 ensure that the operators A , B , and G are bounded, and that the measures ν and $\nu^\perp$ in (5) are equivalent.

To simplify subsequent derivations of acceptance probabilities and various examples of operator-weighted proposals, we now define the following useful transformations.

Definition 3.2 (Transformation that diagonalizes the prior covariance operator).

$$v = \Gamma_{\text{pr}}^{-1/2}(u - u_{\text{ref}}), \quad (14)$$

where $u_{\text{ref}} = m_0 + m_{\text{ref}}$ with $m_{\text{ref}} \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$.

Definition 3.3 (Transformation that diagonalizes operators A , B , and G).

$$w = \Psi^* v. \quad (15)$$

By applying the transformation (14) to u , which has prior measure $\mu_0 = \mathcal{N}(m_0, \Gamma_{\text{pr}})$, we obtain a prior measure on v of the form $\tilde{\mu}_0^v \equiv \mathcal{N}(-v_{\text{ref}}, \text{I})$, where $v_{\text{ref}} = \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}}$. We can then define a reference measure $\mu_0^v = \mathcal{N}(0, \text{I})$ associated with the parameter v . Since we require $m_{\text{ref}} \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$, by the Feldman-Hajek theorem (see Appendix A) we have $\tilde{\mu}_0^v \ll \mu_0^v$ with

$$\frac{d\tilde{\mu}_0^v}{d\mu_0^v}(v) = \exp\left(-\frac{1}{2}\|v_{\text{ref}}\|^2 - \langle v_{\text{ref}}, v \rangle\right).$$

The potential function η has input $u = \Gamma_{\text{pr}}^{1/2} v + u_{\text{ref}}$, which is by definition almost surely in $\mathcal{H}$. Therefore, $\mu_y^v \ll \mu_0^v$ by transitivity of absolute continuity. And thus the posterior measure defined on the parameter v has the form:

$$\frac{d\mu_y}{d\mu_0}(v) \propto \exp\left(-\eta(v; y) - \frac{1}{2}\|v_{\text{ref}}\|^2 - \langle v_{\text{ref}}, v \rangle\right), \quad (16)$$

where $\eta(v; y) = \eta(\Gamma_{\text{pr}}^{1/2} v + u_{\text{ref}}; y)$.

The proposal distribution (13) can then be rewritten as

$$v' = Av - G\nabla_v\eta(v; y) + B\xi, \quad (17)$$

where $\nabla_v\eta(v; y) = \Gamma_{\text{pr}}^{1/2}\nabla_u\eta(u; y)$. Applying the unitary transformation (15), we can simultaneously diagonalize the operators A, B, and G. The resulting proposal distribution on w is then

$$w' = D_A w - D_G \nabla_w\eta(w; y) + D_B \xi, \quad (18)$$

where we again use the degenerate notation $\nabla_w\eta(w; y) = \Psi^*\nabla_v\eta(v; y)$, with $v = \Psi w$.

Remark 3.4. Covariance operators in infinite dimensions are typically taken to be trace-class on $\mathcal{H}$, so that draws have finite $\mathcal{H}$ -norm. For the transformed parameters v and w , the associated proposals use a Gaussian measure with identity covariance operator. Even though the identity operator is not trace-class on $\mathcal{H}$ and hence draws from this Gaussian measure do not have finite $\mathcal{H}$ norm, they are still square-integrable in the weighted space $\text{Im}(\Gamma_{\text{pr}}^{-1/2})$. Furthermore, (17) or (18) still yield a well-defined function space proposal for the parameter u after the inverse transformation is applied.

For brevity of notation, the rest of this paper will define our proposals using the transformed parameter v . Dealing explicitly with the transformed parameter v also allows for a more computationally efficient implementation of the operator-weighted proposal in finite dimensions. This is because each update only requires applying the operator $\Gamma_{\text{pr}}^{1/2}$ once to transform the sample v back to u for evaluating the forward model, and another time to compute $\nabla_v\eta(v; y)$. In comparison, the operator-weighted proposal defined for parameter u requires applying the operator $\Gamma_{\text{pr}}^{1/2}$ four times and applying the operator $\Gamma_{\text{pr}}^{-1/2}$ once for each update.

The following corollary gives the acceptance probability of proposal (13).

Corollary 3.5. *Consider the proposal*

$$v' = Av - G\nabla_v\eta(v; y) + B\xi,$$

where A, B, and G satisfy the conditions of Theorem 3.1. The acceptance probability of this proposal is

$$\alpha(v, v') = \min \{1, \exp(\rho(v', v) - \rho(v, v'))\},$$

where

$$\begin{aligned} \rho(v, v') &= -\eta(v; y) - \langle v_{\text{ref}}, v \rangle - 1/2 \langle v, B_0^{-2} (A_0^2 + B_0^2 - I_0) v \rangle \\ &\quad - \langle B_0^{-1} G_0 \nabla_v\eta(v; y), B_0^{-1} (I_0 v' - A_0 v) \rangle - \frac{1}{2} \|B_0^{-1} G_0 \nabla_v\eta(v; y)\|^2, \end{aligned}$$

and where B_0, A_0, I_0, G_0 denote the respective operators projected onto $\{\psi_i\}_{b_i \neq 0}$.

Proof. The detailed proof is given in Appendix C. □

Theorem 3.1 and Corollary 3.5 provide a wide range of possibilities for constructing valid operator-weighted proposals for infinite-dimensional problems. As described in Section 2, one important advantage of operator-weighted proposals over simple pCN proposals is the ability to prescribe different step sizes in different directions of the parameter space. By utilizing information from the Hessian of the data-misfit functional and the prior covariance, one can construct operators A, B, and G that adapt to the structure of the posterior.

3.2. Likelihood-informed subspace

The first step in constructing the operators appearing in proposal distribution (13) is to identify a suitable set of basis functions Ψ . We will choose these basis functions in order to expose a posterior structure wherein departures from the prior may be limited to a low-dimensional subspace of the parameter space. This structure is a consequence of several factors typical of inverse problems and other distributed-parameter inference configurations: the smoothing action of the forward operator, the limited accuracy or number of observations, and smoothing from the prior. The parameter subspace capturing this prior-to-posterior update will be called the *likelihood-informed subspace* (LIS). In designing function-space MCMC proposals, identification of the LIS will provide a useful way to adapt to posterior structure through the construction of the operators A, B, and G in (13).

For finite-dimensional inverse problems, the LIS can be derived from the dominant eigenvectors of the prior-preconditioned Hessian of the data misfit function, which is a finite dimensional version of the potential $\eta(u; y)$ in (2). This construction was used in [31] to build low-rank update approximations to the posterior covariance matrix for large-scale linear inverse problems. It has also been generalized to inverse problems with infinite-dimensional parameters, for the purpose of constructing a Laplace approximation of the posterior measure [32, 33]. [21] proved the optimality of this approximation in the linear-Gaussian case—in the sense of minimizing the Förstner-Moonen [34] distance from the *exact* posterior covariance matrix over the class of positive definite matrices that are rank- r negative semidefinite updates of the prior covariance, for any given r . These optimality results also extend to optimality statements between Gaussian distributions [21].

The investigations mentioned above compute the LIS *locally*, through the eigendecomposition of the prior-preconditioned Hessian at a given parameter value u . When the forward model is nonlinear, however, the Hessian of the potential varies over the parameter space and hence a local LIS may not be sufficient to capture the entire prior-to-posterior update. One approach for dealing with a non-constant Hessian, introduced in [23], is to “globalize” the LIS by combining local LIS information from many points in the posterior. [23] uses this global LIS to approximate a non-Gaussian posterior in finite dimensions, where the posterior distribution in the complement of the global LIS is taken to be equal to the prior. In other words, the approximation decomposes the posterior into a finite subspace of likelihood-informed directions and a *complementary subspace* (CS) of prior-dominated directions, with mutually independent distributions in each subspace. Here we will extend the notion of a global LIS to the infinite-dimensional setting, where the CS becomes infinite-dimensional. Then we will incorporate the global LIS into the operator-weighted MCMC presented in Section 3.1, in order to achieve exact posterior sampling rather than posterior approximation.

In the remainder of Section 3.2, we review the computation of the local LIS and then discuss the globalization of the LIS through the estimation of posterior expectations. Based on the global LIS, we will also present a structure-exploiting posterior covariance approximation for non-Gaussian posteriors. This covariance approximation naturally integrates the global LIS with the second-order geometry of the posterior, and will be used in later subsections to design operator-weighted proposals. We conclude this subsection by discussing iterative construction procedures for the global LIS and the posterior covariance approximation.

3.2.1. Local likelihood-informed subspace

Let the forward model be Fréchet differentiable. Then the linearization of the forward model at a given parameter u , $J(u) = \nabla_u F(u)$, yields the local sensitivity of the parameter-to-data map.

The Gauss-Newton approximation of the Hessian of the data-misfit functional (2) at u (hereafter referred to as the GNH) then becomes

$$\mathbf{H}(u) = \mathbf{J}^*(u)\Gamma_{\text{obs}}^{-1}\mathbf{J}(u), \quad (19)$$

and can be used to construct a local Gaussian approximation $\mathcal{N}(m(u), \Gamma_{\text{pos}}(u))$ of the posterior measure, where

$$\Gamma_{\text{pos}}^{-1}(u) = \mathbf{H}(u) + \Gamma_{\text{pr}}^{-1} \quad \text{and} \quad m(u) = u - \Gamma_{\text{pos}}(u)\nabla_u \eta(u; y). \quad (20)$$

Along a function $\varphi \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$, consider the local Rayleigh ratio

$$\mathcal{R}(\varphi; u) = \frac{\langle \varphi, \mathbf{H}(u)\varphi \rangle}{\langle \varphi, \Gamma_{\text{pr}}^{-1}\varphi \rangle}. \quad (21)$$

This ratio quantifies how strongly the likelihood constrains variation in the φ direction relative to the prior. When this ratio is large, the likelihood limits variation in the φ direction more strongly than the prior, and vice-versa. We note that similarly to the Onsager-Machlup functional, (21) can be extended to $\mathcal{H}$ by setting $R(\varphi; u) = 0$ outside $\text{Im}(\Gamma_{\text{pr}}^{1/2})$. We will in particular be concerned with φ such that $\|\Gamma_{\text{pr}}^{-1/2}\varphi\| = 1$.

Applying the change of variables $\phi = \Gamma_{\text{pr}}^{-1/2}\varphi \in \mathcal{H}$, the Rayleigh ratio (21) becomes

$$\tilde{\mathcal{R}}(\varphi; u) = \mathcal{R}(\phi; u) = \frac{\langle \Gamma_{\text{pr}}^{1/2}\phi, \mathbf{H}(u)\Gamma_{\text{pr}}^{1/2}\phi \rangle}{\langle \phi, \phi \rangle}. \quad (22)$$

This naturally leads to considering the eigendecomposition of the prior-preconditioned GNH (ppGNH)

$$\left(\Gamma_{\text{pr}}^{1/2}\mathbf{H}(u)\Gamma_{\text{pr}}^{1/2}\right)\phi_i = \lambda_i\phi_i, \quad (23)$$

which is positive semidefinite by construction. The dominant eigenfunctions of the ppGNH are local maximizers of the Rayleigh quotient (22). Indeed, for the basis function $\varphi_i = \Gamma_{\text{pr}}^{1/2}\phi_i \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$, the Rayleigh quotient has value λ_i , i.e.,

$$\mathcal{R}(\varphi_i; u) = \lambda_i.$$

Thus, the ppGNH captures the (local) balance of likelihood and prior information described above: the largest eigenvalues of (23) correspond to directions along which the likelihood dominates the prior, and the smallest eigenvalues correspond to directions along which the posterior is determined almost entirely by the prior. The basis functions $\{\varphi_1, \dots, \varphi_l\}$ corresponding to the l leading eigenvalues of (23), such that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_l \geq \varrho_l > 0$, span the *local* likelihood-informed subspace (LIS). The resulting approximation of the local posterior covariance $\Gamma_{\text{pos}}(u)$ has the form

$$\Gamma_{\text{pos}}(u) \approx \Gamma_{\text{pos}}^{(l)} = \Gamma_{\text{pr}} - \sum_{i=1}^l \frac{\lambda_i}{\lambda_i + 1} \varphi_i \varphi_i^*, \quad (24)$$

where φ_i and λ_i depend on the point u . Note that a Gaussian measure with this covariance and with mean $m(u)$ (20) still endows the Hilbert space $\mathcal{H}$ with full measure. In a finite-dimensional setting,

this covariance corresponds precisely to the approximation shown to be optimal in [21]. Note also that the number of nonzero eigenvalues in (23) is most N_y , and thus the local LIS is indeed low-dimensional compared to the infinite-dimensional parameter space. If the nonzero eigenvalues λ_i decay quickly, then one can reduce the dimension of the local LIS even further, with an appropriate choice of truncation threshold $\varrho_1 > 0$.

Remark 3.6. The change of variables $\phi = \Gamma_{\text{pr}}^{-1/2}\varphi$ has the same role as the whitening transformation in Definition 3.2. For the parameter $v = \Gamma_{\text{pr}}^{1/2}(u - u_{\text{ref}})$, the GNH w.r.t. v is

$$H(v) = \Gamma_{\text{pr}}^{1/2}H(u)\Gamma_{\text{pr}}^{1/2},$$

which is identical to the ppGNH in (23). The associated local Rayleigh quotient for parameter v has the form

$$\tilde{\mathcal{R}}(\phi; v) = \frac{\langle \phi, H(v)\phi \rangle}{\langle \phi, \phi \rangle}. \quad (25)$$

In this case, the orthonormal basis $\{\phi_1, \dots, \phi_l\}$ forms the local LIS for the parameter v . Henceforth, for brevity and to maintain notation consistent with our exposition of operator-weighted proposals above, we will discuss the global LIS and the associated posterior covariance approximation in terms of the transformed parameter v .

3.2.2. Global likelihood-informed subspace

To extend the pointwise criterion (25) into a global criterion for identifying likelihood-informed directions for nonlinear problems, we consider the posterior expectation of the local Rayleigh quotient,

$$\mathbb{E} [\tilde{\mathcal{R}}(\phi; v)] = \frac{\langle \phi, S\phi \rangle}{\langle \phi, \phi \rangle},$$

where S is the expectation of the local GNH (working in whitened coordinates):

$$S = \int_{\text{Im}(\Gamma_{\text{pr}}^{-1/2})} H(v)\mu_y(dv). \quad (26)$$

As in the local case, it is natural to then construct a global LIS from the eigendecomposition

$$S\theta_j = \rho_j\theta_j. \quad (27)$$

The eigenfunctions corresponding to the r leading eigenvalues $\rho_1 \geq \rho_2 \geq \dots \geq \rho_r \geq \varrho_g > 0$, for some threshold $\varrho_g > 0$, form a basis $\Theta_r = [\theta_1, \dots, \theta_r]$ of the global LIS (again in the whitened coordinates). We typically choose a truncation threshold $\varrho_g \approx 0.1$, thus retaining parameter directions where the “impact” of likelihood is at least one tenth that of prior. The expected GNH can be approximated using a Monte Carlo estimator based on samples adaptively drawn from the posterior. This adaptive procedure is described in Section 3.2.4.

Remark 3.7. Since we assume that the Fréchet derivative of the forward operator exists, the GNH operator is defined on function space. We take the approach of differentiating and then discretizing, so that dimension independence is conserved assuming a convergent discretization of the linearization.

3.2.3. Low-rank posterior covariance approximation

The posterior covariance describes important aspects of the global structure of the posterior (e.g., its scale and orientation) and thus is often used in constructing efficient MCMC samplers for the finite dimensional setting—for example, the adaptive Metropolis [35] and adaptive Metropolis-adjusted Langevin [36] methods. When the posterior covariance matrix arises from the discretization of an infinite-dimensional posterior, however, computing, storing, and factorizing this matrix—all standard operations required by most MCMC methods—are not feasible. Suppose that the discretized parameter space has dimension N . Then estimating the posterior covariance from samples (without assuming any special structure) requires at least $O(N)$ posterior samples; the storage cost is $O(N^2)$; and the computational cost of a covariance factorization is $O(N^3)$. Here, we will instead use the global LIS to derive a covariance approximation scheme for which the costs of estimating, storing, and factorizing the posterior covariance all scale linearly with the discretization dimension N .

Since the global LIS captures those parameter directions where the posterior differs most strongly from the prior, we can approximate the posterior measure by projecting the argument of the likelihood function onto the finite-dimensional global LIS spanned by Θ_r . Given the rank- r orthogonal projection operator $\Pi_r = \Theta_r \Theta_r^*$, the approximation takes the form

$$\mu_y(dv) \approx \tilde{\mu}_y(dv) \propto \mathcal{L}(\Pi_r v; y) \mu_0(dv). \quad (28)$$

The orthogonal projector Π_r decomposes the parameter as

$$v = \Pi_r v + (\mathbf{I} - \Pi_r) v,$$

where $\Pi_r v$ and $(\mathbf{I} - \Pi_r) v$ are projections of the parameter v onto the low-dimensional LIS and the infinite-dimensional CS, respectively. Since the prior measure $\mu_0(dv)$ is Gaussian with zero mean and identity covariance, $\Pi_r v$ and $(\mathbf{I} - \Pi_r) v$ are independent under the prior. Furthermore, the argument of the likelihood function in (28) depends only on the LIS-projected parameter $\Pi_r v$. Therefore, $\Pi_r v$ and $(\mathbf{I} - \Pi_r) v$ are also independent under the approximated posterior $\tilde{\mu}_y(dv)$. This independence and the LIS projection lead to an approximate posterior covariance (for the parameter v) that takes the form

$$\text{Cov}_{\mu_y}[v] \approx \Sigma := \text{Cov}_{\mu_y}[\Pi_r v] + \text{Cov}_{\mu_0}[(\mathbf{I} - \Pi_r) v].$$

This covariance Σ is the sum of the marginal posterior covariance within the LIS,

$$\text{Cov}_{\mu_y}[\Pi_r v] = \Theta_r \Sigma_r \Theta_r^*,$$

where the $r \times r$ matrix Σ_r is the posterior covariance projected onto the LIS basis Θ_r , and the marginal prior covariance within the CS,

$$\text{Cov}_{\mu_0}[(\mathbf{I} - \Pi_r) v] = \mathbf{I} - \Theta_r \Theta_r^*.$$

Introducing the eigendecomposition $\Sigma_r = W_r D_r W_r^*$ and reweighing the basis Θ_r using the orthogonal matrix W_r ,

$$\Psi_r = \Theta_r W_r, \quad (29)$$

we can diagonalize the covariance approximation and write it as a low-rank update of the identity:

$$\begin{aligned} \Sigma &= \Theta_r \Sigma_r \Theta_r^* + (\mathbf{I} - \Theta_r \Theta_r^*) \\ &= \Psi_r (D_r - \mathbf{I}_r) \Psi_r^* + \mathbf{I}. \end{aligned} \quad (30)$$

where $\mathbf{I}_r$ denotes the $r \times r$ identity matrix. The corresponding approximate posterior covariance for the parameter u is $\Gamma_{\text{pr}}^{1/2} \Sigma \Gamma_{\text{pr}}^{1/2}$. Note that $\Pi_r = \Theta_r \Theta_r^* = \Psi_r \Psi_r^*$.

Given a basis for the LIS, the number of samples required to estimate the low-rank approximation (30) now depends on the rank of the LIS rather than on the discretized parameter dimension N . Furthermore, the cost of storing this approximated covariance is dominated by the storage of the LIS basis, which is proportional to the discretization dimension N . The cost of computing the eigendecomposition of Σ_r is of order r^3 , which is asymptotically negligible compared with the discretization dimension. Similarly, the factorization of the diagonalized posterior covariance (30) only involves the $r \times r$ diagonal matrix $\mathbf{D}_r$, and hence these operations are also asymptotically negligible. The dominant computational cost here is reweighing the LIS basis as in (29), which is proportional to the discretization dimension N .

3.2.4. Iterative construction

The low-rank approximation of the posterior covariance (30) will be used to design operator-weighted proposals. Yet building this covariance approximation requires posterior samples, both for computing the global LIS and for constructing the projected posterior covariance Σ_r . These computations must therefore be carried out on-the-fly during posterior sampling. Here we discuss two building block algorithms—for estimating the global LIS and the projected posterior covariance from a given set of posterior samples—which are later used in the adaptive sampling algorithms of Section 3.4.

We first consider the global LIS. Suppose that we have a set of m posterior samples $v_k \sim \mu_y$ for $k = 1 \dots m$. The expected GNH (26) could be directly approximated using the Monte Carlo estimator

$$\mathbf{S}_m = \frac{1}{m} \sum_{k=1}^m \mathbf{H}(v_k). \quad (31)$$

If the discretized parameter has dimension N , the cost of storing $\mathbf{S}_m$ and the computational cost of updating it are both proportional to N^2 . Furthermore, the cost of a direct eigendecomposition of $\mathbf{S}_m$ for computing the global LIS basis is of order N^3 . Thus, it is not feasible to directly apply the Monte Carlo estimate and eigendecomposition here. Instead, we again use the low-rank structure of the problem to design an iterative updating scheme for computing and updating the eigendecomposition of $\mathbf{S}_m$ as more posterior samples are drawn.

As discussed earlier for the local LIS, each local GNH $\mathbf{H}(v_k)$ is well captured by a low-rank decomposition

$$\mathbf{H}(v_k) \approx \sum_{i=1}^{l(v_k)} \lambda_i(v_k) \phi_i(v_k) \phi_i^*(v_k), \quad (32)$$

where $\lambda_i(v_k)$ and $\phi_i(v_k)$ are eigenvalues and eigenvectors depending on the parameter v_k . The eigenspectrum is truncated at some local threshold ϱ_1 , i.e., $\lambda_1(v_k) \geq \lambda_2(v_k) \geq \dots \geq \lambda_{l(v_k)}(v_k) \geq \varrho_1 > 0$. Similar to the global LIS truncation threshold, we typically choose the local truncation threshold $\varrho_1 = \varrho_g = 0.1$. Now let $\Lambda(v_k)$ be the $l(v_k) \times l(v_k)$ diagonal matrix such that $\Lambda_{ij}(v_k) = \delta_{ij} \lambda_i(v_k)$, and let $\Phi(v_k) = [\phi_1(v_k), \dots, \phi_{l(v_k)}(v_k)]$ be the rank- $l(v_k)$ local LIS basis. Using matrix notation, the decomposition of the local GNH can be rewritten as

$$\mathbf{H}(v_k) \approx \Phi(v_k) \Lambda(v_k) \Phi^*(v_k).$$

By computing the action of the local GNH on functions—each action requires one forward model evaluation and one adjoint model evaluation—the eigendecomposition can be computed using either Krylov subspace methods [37] or randomized algorithms [38, 39]. This way, the number of forward and adjoint evaluations used to obtain (32) is proportional to the target rank $l(v_k)$.

For the current sample set $\{v_k\}_{k=1}^m$, suppose that the expected GNH also admits a rank- M approximation for some $M(m)$, in the form of

$$S_m \approx \Theta_m \Xi_m \Theta_m^*,$$

where $\Xi_m \in \mathbb{R}^{M \times M}$, with $[\Xi_m]_{ij} = \delta_{ij} \rho_i$ and $\Theta_m = [\theta_1, \dots, \theta_M]$ containing the eigenvalues and eigenfunctions of S_m , respectively. To maintain an accurate representation of the expected GNH and its eigendecomposition during the adaptive construction, we truncate this eigendecomposition conservatively rather than applying the $\rho_g = 0.1$ threshold used to truncate the global LIS. In this paper, we set the truncation threshold to $\rho_s = 10^{-4}$. Given a new sample v_{m+1} , the eigendecomposition can be updated using the incremental SVD method [40]; this procedure is described in Steps 6–9 of Algorithm 1. In the discretization of the infinite-dimensional problem, the bases Θ_m and $\Phi(v_{m+1})$ are represented by “tall and skinny” matrices, and thus the computational cost of this iterative update is proportional to the discretization dimension N . The major computational cost in this procedure is the eigendecomposition of the local GNH—which involves a number of linearized forward model and adjoint model solves—in Step 2 of Algorithm 1. To initialize the global LIS, we choose the MAP estimate (defined in Section 2.1) as the initial sample at which to decompose the local GNH.

Algorithm 1 Incremental update of the expected GNH and global LIS.

Input: Θ_m, Ξ_m, v_{m+1}

Output: $\Theta_{m+1}, \Xi_{m+1}, \Theta_r, \Xi_r$

- 1: **procedure** UPDATELIS($\Theta_m, \Xi_m, v_{m+1}, \Theta_{m+1}, \Xi_{m+1}, \Theta_r, \Xi_r$)
 - 2: Compute the local LIS basis through $H(v_{m+1}) \approx \Phi(v_{m+1}) \Lambda(v_{m+1}) \Phi^*(v_{m+1})$
 - 3: **if** $m = 0$ **then**
 - 4: Initialize S_1 by setting $\Theta_1 = \Phi(v_1)$ and $\Xi_1 = \Lambda(v_1)$
 - 5: **else**
 - 6: Compute the QR decomposition $[\Theta_m, \Phi(v_{m+1})] = QR$
 - 7: Compute the new eigenvalues through the eigendecomposition

$$\frac{1}{m+1} R \begin{bmatrix} m \Xi_m & 0 \\ 0 & \Lambda(v_{m+1}) \end{bmatrix} R^* = W \Xi_{m+1} W^*$$
 - 8: Compute the new basis $\Theta_{m+1} = QW$
 - 9: Return new decomposition Θ_{m+1} and Ξ_{m+1} such that all $\rho_i \geq \rho_s$
 - 10: **end if**
 - 11: Return LIS basis $\Theta_r = \Theta_{m+1}(:, 1:r)$ and LIS eigenvalues $\Xi_r = \Xi_{m+1}(1:r, 1:r)$,
 such that $\rho_1, \dots, \rho_r \geq \rho_g$ ▷ MATLAB notation is used here
 - 12: **end procedure**
-

Algorithm 1 provides an efficient way of updating the expected GNH given a new posterior sample v_{m+1} . In Section 3.4 we will describe how to embed Algorithm 1 into an adaptive framework for posterior exploration, where new samples for each update are generated by operator-weighted

MCMC. This approach thus *automatically* explores the landscape of the likelihood; depending on this landscape, the number of samples needed to give a reasonable estimate of the expected GNH (and in particular, its dominant eigendirections) may vary significantly. (For instance, if the log-likelihood is quadratic, only one sample is needed!) In general, to limit the total number of expensive eigendecompositions of the local GNH, we would like to terminate this adaptive construction process as soon as the estimated LIS stabilizes.

To this end, let $\{\Theta_r, \Xi_r\}$ and $\{\Theta_{r'}, \Xi_{r'}\}$ denote the global likelihood informed subspaces and associated eigenvalues evaluated at successive iterations of the LIS construction process. And let

$$S_r = \Theta_r \Xi_r \Theta_r^*, \quad \text{and} \quad S_{r'} = \Theta_{r'} \Xi_{r'} \Theta_{r'}^*,$$

denote the (truncated) expected GNHs at the same two iterations. We use the Förstner distance [34] between $(I + S_r)$ and $(I + S_{r'})$,

$$d_{\mathcal{F}}(I + S_r, I + S_{r'}) = \left(\sum_{i=1}^{\infty} \ln^2(\lambda_i(I + S_r, I + S_{r'})) \right)^{\frac{1}{2}}, \quad (33)$$

to monitor the converge of the LIS. Here $\lambda_i(A, B)$ is the i th generalized eigenvalue of the pencil (A, B) . The Förstner distance was introduced in [34] as a metric for positive definite covariance matrices, and is also related to the Kullback-Leibler divergence or Hellinger distance between Gaussian distributions, as described in [21]. In the present setting, the infinite-dimensional generalized eigenproblem in (33) can be reformulated as a finite-dimensional generalized eigenproblem, so that the Förstner metric becomes directly applicable. Since the expected GNHs S_r and $S_{r'}$ are symmetric and finite-rank operators, the corresponding operators $(I + S_r)$ and $(I + S_{r'})$ differ only on the finite dimensional subspace $\mathcal{V}$ that is the sum of the ranges of S_r and $S_{r'}$. In other words, there are only a finite number of nonzero terms in the summation (33). Letting $\mathcal{V} = \text{span}(\Theta_r) + \text{span}(\Theta_{r'})$, $d_{\mathcal{F}}$ can be equivalently and efficiently computed by projecting the operators $(I + S_r)$ and $(I + S_{r'})$ onto $\mathcal{V}$. Recall that in the discretized space both Θ_r and $\Theta_{r'}$ are “tall and skinny” matrices, and therefore the computational cost of evaluating $d_{\mathcal{F}}$ scales linearly with the discretization dimension N .

An alternative criterion for monitoring the convergence of the global LIS is the weighted subspace distance from [41], which is used in [23]. In practice, we find that both criteria produce consistent results.

Algorithm 2 Update of the low-rank posterior covariance approximation.

Input: Θ_r, Σ_r

Output: Ψ_r, D_r

- 1: **procedure** UPDATECOV($\Theta_r, \Sigma_r, \Psi_r, D_r$)
 - 2: Compute eigendecomposition $W_r D_r W_r^* = \Sigma_r$
 - 3: Return the reweighted LIS basis $\Psi_r = \Theta_r W_r$ and the diagonalized covariance D_r
 - 4: **end procedure**
-

Given a global LIS basis Θ_r and an empirical estimate of the posterior covariance projected onto the LIS, Σ_r , the low-rank approximation (30) of the posterior covariance can be constructed using the procedure described in Algorithm 2. When the LIS is fixed, estimation of Σ_r from posterior samples only requires storage of order r^2 for the covariance matrix itself and an empirical

mean. Given a new posterior sample v , Σ_r can be updated by a rank-one modification using the r -dimensional vector $\Theta_r^* v$. Since the factorization and inversion of the posterior covariance (30) only involve the $r \times r$ diagonal matrix D_r , their computational costs are negligible compared with the discretization dimension N . The computational costs of projecting the sample v , reweighing the LIS basis as in (29), and computing the eigendecomposition of Σ_r scale as Nr , Nr^2 , and r^3 , respectively. The sample projection $\Theta_r^* v$ has negligible cost compared with other operations such as PDE solves, and hence does not require special treatment.

Algorithm 3 Posterior covariance re-projection for each LIS update.

Input: $\Theta_r, \Sigma_r, \Theta_{r'}$

Output: $\Sigma_{r'}$

- 1: **procedure** PROJECTCOV($\Theta_r, \Sigma_r, \Theta_{r'}, \Sigma_{r'}$)
 - 2: $T = \Theta_{r'}^* \Theta_r$
 - 3: $\Sigma_{r'} = T(\Sigma_r - I_r)T^* + I_{r'}$
 - 4: **end procedure**
-

Finally, we note that each time the global LIS is updated, the estimate Σ_r of the posterior covariance in the LIS must itself be updated. Suppose that the current LIS basis is Θ_r and that the associated projected posterior covariance is Σ_r . Given a new LIS basis $\Theta_{r'}$, the new projected posterior covariance $\Sigma_{r'}$ can be obtained by projecting the current posterior covariance estimate $\Sigma = \Theta_r(\Sigma_r - I_r)\Theta_r^* + I$ onto the new LIS, as shown in Algorithm 3. The computational cost of this procedure is dominated by the multiplication of the LIS bases in step 2, which is proportional to Nrr' .

3.3. DILI proposals

Now we use the global LIS and the low-rank posterior covariance approximation (30) to develop operator-weighted MCMC proposals that are efficient and that satisfy the conditions of Theorem 3.1, thus ensuring discretization-invariant mixing in a function space setting.

The low-rank posterior covariance $\Sigma = \Psi_r(D_r - I_r)\Psi_r^* + I$ is entirely specified by the global LIS basis Ψ_r and the r -dimensional diagonal matrix D_r . Given these objects, the transformed parameter space $\mathcal{H}_v \equiv \text{Im}(\Gamma_{\text{pr}}^{-1/2})$ can be partitioned as $\mathcal{H}_v = \mathcal{H}_r \oplus \mathcal{H}_\perp$, where the LIS is $\mathcal{H}_r = \text{span}(\Psi_r)$ and has finite Hilbert dimension, i.e., $\text{card}(\Psi_r) = r < \infty$. We can associate the infinite-dimensional CS² $\mathcal{H}_\perp \equiv (\mathcal{H}_v)^\perp$ with a set of eigenfunctions $\Psi_\perp = [\psi_{r+1}, \psi_{r+2}, \dots]$ such that $\mathcal{H}_\perp = \text{span}(\Psi_\perp)$. $\Psi = [\Psi_r, \Psi_\perp]$ thus forms a complete orthonormal system in $\mathcal{H}_v$. Recall that the LIS captures parameter directions where the posterior differs most strongly from the prior. This partition of the parameter space, along with the approximate posterior covariance Σ , can then be used to assign eigenvalues to the basis Ψ in order to construct the operators A, B, and G in the proposal (13).

An appropriate discretization of the Langevin SDE (7) naturally leads to an eigenvalue assignment that satisfies Theorem 3.1. Following the sequence of transformations in Definitions 3.2 and 3.3, the Langevin SDE (7) preconditioned by the low-rank approximation of the posterior covariance $\Gamma_{\text{pr}}^{1/2} \Sigma \Gamma_{\text{pr}}^{1/2}$ (for the parameter u) takes the form

$$\begin{pmatrix} dw_r \\ dw_\perp \end{pmatrix} = - \begin{pmatrix} D_r w_r d\tau \\ w_\perp d\tau \end{pmatrix} - \begin{pmatrix} \gamma_r D_r (\Psi_r^* \nabla_w \eta(w; y)) d\tau \\ \gamma_\perp \Psi_\perp^* \nabla_w \eta(w; y) d\tau \end{pmatrix} + \begin{pmatrix} \sqrt{2D_r} d\mathcal{W}_r(\tau) \\ \sqrt{2} d\mathcal{W}_\perp(\tau) \end{pmatrix}, \quad (34)$$

²In whitened coordinates, i.e., in terms of the transformed parameter v , the CS is the *orthogonal* complement of the LIS.

where $w_r = \Psi_r^* v$ and $w_\perp = \Psi_\perp^* v$ are coefficients of the projection of v onto the LIS and CS, respectively. The two tuning parameters γ_r and $\gamma_\perp$ take values $\{0, 1\}$ to switch between a Langevin proposal and random walk proposal in the LIS and CS. Now, discretizing this SDE is equivalent to assigning eigenvalues to the operators A, B, and G of the operator-weighted proposal (13). We will employ an inhomogeneous time discretization strategy: a time discretization scheme tailored to the non-Gaussian structure—that is not necessarily dimension-independent—is used in the finite-dimensional LIS, while a standard Crank-Nicolson scheme—to ensure dimension-independence—is used in the infinite-dimensional CS.

In the rest of this subsection, we will provide specific examples of these operator-weighted proposals. We will prescribe two different time steps to discretize (34), denoted by $\Delta\tau_r$ and $\Delta\tau_\perp$. All the proposals we present depend only on the global LIS basis Ψ_r , the $r \times r$ diagonalized covariance D_r , and the time steps $\Delta\tau_r$ and $\Delta\tau_\perp$. Given the partition defined by the LIS, we will also present extensions of those proposals to Metropolis-within-Gibbs schemes.

3.3.1. Basic proposals

Here we consider all-at-once proposals that update all the parameter in a single iteration. Two examples will be provided, one which does not use gradient information on the LIS ($\gamma_r = 0$) and one which does ($\gamma_r = 1$). On the CS, both proposals use the pCN update with $\gamma_\perp = 0$. Thus the proposal in the CS is given by

$$w'_\perp = a_\perp w_\perp + b_\perp \xi_\perp, \quad (35)$$

where

$$\xi_\perp \sim \mathcal{N}(0, I_\perp), \quad a_\perp = \frac{2 - \Delta\tau_\perp}{2 + \Delta\tau_\perp}, \quad \text{and} \quad b_\perp = \sqrt{1 - a_\perp^2}.$$

Here we use $I_\perp$ to denote the identity operators associated with $\mathcal{H}_\perp$, whereas the identity operator associated with $\mathcal{H}_r$ is an $r \times r$ identity matrix I_r . For the corresponding eigenfunctions $\Psi_\perp$, the operators A, B, and G have uniform eigenvalues $a_\perp$, $b_\perp$, and 0, respectively.

Various discretization schemes can be used for the Langevin SDE in the LIS. Let the eigenfunctions Ψ_r be associated with sequences of eigenvalues (a_i) , (b_i) , and (g_i) , $i = 1, \dots, r$, giving rise to the $r \times r$ diagonal matrices D_{A_r} , D_{B_r} , and D_{G_r} . The operators A, B, and G can then be written as

$$\begin{aligned} A &= \Psi_r (D_{A_r} - a_\perp I_r) \Psi_r^* + a_\perp I, \\ B &= \Psi_r (D_{B_r} - b_\perp I_r) \Psi_r^* + b_\perp I, \\ G &= \Psi_r D_{G_r} \Psi_r^*. \end{aligned} \quad (36)$$

Theorem 3.1 and Corollary 3.5 can be used to show when the operator-weighted proposals given by (36) yield well-defined MCMC algorithms, and to derive the associated acceptance probability.

Our first operator-weighted proposal employs a Crank-Nicolson scheme in the LIS, with $\gamma_r = 0$. The resulting discretized Langevin SDE in the LIS has the form

$$w'_r - w_r = -\frac{1}{2} D_r (w'_r + w_r) \Delta\tau_r + \sqrt{2 D_r \Delta\tau_r} \xi_r,$$

where $\xi_r \sim \mathcal{N}(0, I_r)$. Together with the CS proposal (35), we have the following complete proposal:

Proposal 3.1 (LI-Prior: Likelihood-informed proposal that leaves the shifted prior distribution invariant).

$$v' = A v + B \xi,$$

where A , B , and G are given by (36), with

$$\begin{aligned} D_{A_r} &= (2I_r + \Delta\tau_r D_r)^{-1} (2I_r - \Delta\tau_r D_r), \\ D_{B_r} &= (I_r - D_{A_r}^2)^{1/2}, \\ D_{G_r} &= 0. \end{aligned} \quad (37)$$

Since this proposal has $D_{A_r}^2 + D_{B_r}^2 = I_r$, the acceptance probability can be simplified to

$$\alpha(v, v') = \min \{1, \exp(\eta(v) - \eta(v') + \langle v_{\text{ref}}, v - v' \rangle)\}, \quad (38)$$

where $v_{\text{ref}} = \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}}$.

We can instead consider a mixed discretization scheme, which uses an explicit discretization of the Langevin proposal ($\gamma_r = 1$) in the LIS. The discretized Langevin SDE in the LIS has the form

$$w'_r - w_r = -\Delta\tau_r D_r (w_r + \Psi_r^* \nabla_w \eta(w; y)) + \sqrt{2D_r \Delta\tau_r} \xi_r,$$

where $\xi_r \sim \mathcal{N}(0, I_r)$. Combined with the CS proposal (35), this leads to the proposal:

Proposal 3.2 (LI-Langevin: Likelihood-informed proposal that uses the Langevin proposal in the LIS).

$$v' = Av - G \nabla_v \eta(v; y) + B\xi,$$

where A , B , and G are given by (36), with

$$\begin{aligned} D_{A_r} &= I_r - \Delta\tau_r D_r, \\ D_{B_r} &= \sqrt{2\Delta\tau_r D_r}, \\ D_{G_r} &= \Delta\tau_r D_r. \end{aligned} \quad (39)$$

The acceptance probability of this proposal has the form

$$\alpha(v, v') = \min \{1, \exp(\rho(v', v) - \rho(v, v'))\}, \quad (40)$$

where $\rho(v, v')$ is given by

$$\begin{aligned} \rho(v, v') &= -\eta(v; y) - \langle v_{\text{ref}}, v \rangle - \frac{1}{2} \|\Psi_r^* v\|^2 \\ &\quad - \frac{1}{2} \left\| D_{B_r}^{-1} (\Psi_r^* v' - D_{A_r} \Psi_r^* v + D_{G_r} \Psi_r^* \nabla_v \eta(v; y)) \right\|^2, \end{aligned}$$

where $v_{\text{ref}} = \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}}$.

3.3.2. Metropolis-within-Gibbs scheme

Given a basis Ψ_r for the LIS, the parameter v can be split as

$$v = v_r + v_{\perp}, \quad v_r = \Psi_r w_r, \quad v_{\perp} = \Psi_{\perp} w_{\perp},$$

and hence a Metropolis-within-Gibbs update can be quite naturally introduced for v_r and $v_{\perp}$. This update consists of a pair of proposals:

$$\text{Update } v_r \text{ conditioned on } v_{\perp} : v' = A_r v - G_r \nabla_v \eta(v; y) + B_r \xi, \quad (41a)$$

$$\text{Update } v_{\perp} \text{ conditioned on } v_r : v' = A_{\perp} v + B_{\perp} \xi, \quad (41b)$$

where $\xi \sim \mathcal{N}(0, \mathbf{I})$. By setting $a_\perp = 1$ and $b_\perp = 0$ in (36), we can derive the proposal (41a), which updates v_r conditional on $v_\perp$. This leads to the operators

$$\mathbf{A}_r = \Psi_r (\mathbf{D}_{\mathbf{A}_r} - \mathbf{I}_r) \Psi_r^* + \mathbf{I}, \quad \mathbf{B}_r = \Psi_r \mathbf{D}_{\mathbf{B}_r} \Psi_r^*, \quad \text{and} \quad \mathbf{G}_r = \Psi_r \mathbf{D}_{\mathbf{G}_r} \Psi_r^*. \quad (42)$$

Similarly, by setting $\mathbf{D}_{\mathbf{A}_r} = \mathbf{I}_r$, $\mathbf{D}_{\mathbf{B}_r} = 0$, and $\mathbf{D}_{\mathbf{G}_r} = 0$, we have the proposal (41b) which updates $v_\perp$ conditional on v_r . This leads to the operators

$$\mathbf{A}_\perp = \Psi_r \Psi_r^* + a_\perp (\mathbf{I} - \Psi_r \Psi_r^*), \quad \text{and} \quad \mathbf{B}_\perp = b_\perp (\mathbf{I} - \Psi_r \Psi_r^*). \quad (43)$$

These proposals both satisfy the conditions of Theorem 3.1, and therefore lead to a Metropolis-Hastings scheme that is well-defined on function space. The following two Metropolis-within-Gibbs schemes are derived from Proposals 3.1 and 3.2.

Proposal 3.3 (MGLI-Prior: Metropolis-within-Gibbs update using a likelihood-informed proposal that is invariant w.r.t. the shifted prior distribution). *Let $\mathbf{D}_{\mathbf{A}_r}$, $\mathbf{D}_{\mathbf{B}_r}$, and $\mathbf{D}_{\mathbf{G}_r}$ be given by (37). Define $\mathbf{A}_r$, $\mathbf{B}_r$, and $\mathbf{G}_r$ as in (42) and $\mathbf{A}_\perp$, $\mathbf{B}_\perp$ as in (43). The updates are given by (41a) and (41b) and both have acceptance probability (38).*

Proposal 3.4 (MGLI-Langevin: Metropolis-within-Gibbs update using the Langevin proposal in the LIS). *Let $\mathbf{D}_{\mathbf{A}_r}$, $\mathbf{D}_{\mathbf{B}_r}$, and $\mathbf{D}_{\mathbf{G}_r}$ be given by (37). Define $\mathbf{A}_r$, $\mathbf{B}_r$, and $\mathbf{G}_r$ as in (42) and $\mathbf{A}_\perp$ and $\mathbf{B}_\perp$ as in (43). The updates are given by (41a) and (41b), with acceptance probabilities given by (40) and (38), respectively.*

3.4. Adaptive posterior sampling framework

Now we describe how the construction of the global LIS, the low-rank posterior covariance approximation, and the operator-weighted proposals can be integrated into a single adaptive posterior sampling framework. The building blocks have already been given; in particular, Section 3.2.4 describes how to construct and iteratively update the global LIS basis Θ_r (or equivalently Ψ_r (29)) and the posterior covariance Σ from posterior samples, while Section 3.3 shows how to construct several different operator-weighted proposals from the LIS and covariance information. All that remains is to embed the first two operations within an MCMC approach. The overall framework is described in Algorithm 4. At specified intervals (controlled by the parameter n_{lag}) and until a final tolerance (controlled by Δ_{LIS} and n_{max}) is reached, the LIS is updated from previous posterior samples. Every n_b samples and/or when the LIS basis is updated, and for the duration of the algorithm, the posterior covariance Σ and the operators in the DILI proposal are updated from previous samples. Thus the algorithm adaptively captures the global structure of the posterior from local Hessian information and from samples, and tunes the resulting operator-weighted proposals accordingly.

Generally speaking, the major computational cost of Algorithm 4 is the evaluation of the acceptance probability at Step 2, which involves a computationally expensive likelihood evaluation and—when the Langevin proposals are used—one extra adjoint solve for the gradient evaluation. The proposal evaluation in Step 1 only incurs a minor computational cost, as the cost of applying the low-rank decompositions of the operators $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{G}$ is proportional to the discretization dimension N . The computational cost of the LI-Prior proposal is therefore comparable to that of a pCN random walk proposal—or other types of random walk proposals used in the finite dimensional setting—and the computational cost of the LI-Langevin proposal is comparable to standard

Algorithm 4 Adaptive function space MCMC with operator-weighted proposals.

Require: During the LIS construction, we retain (1) $\{\Theta_m, \Xi_m\}$ to store the expected GNH evaluated from m samples; and (2) the value of $d_{\mathcal{F}}$ (33) between the most recent two updates of the expected GNH, for LIS convergence monitoring.

Require: At step n , given the state v_n , LIS basis Θ_r , projected empirical posterior covariance Σ_r , and operators $\{A, B, G\}$ induced by $\{\Psi_r, D_r, \Delta\tau_r, \Delta\tau_{\perp}\}$, one step of the algorithm is:

- 1: Compute a candidate $v' = q(v_n, \cdot; A, B, G)$ using either LI-prior or LI-Langevin
- 2: Compute the acceptance probability $\alpha(v_n, v')$
- 3: **if** $\text{Uniform}(0, 1] < \alpha(v_n, v')$ **then** $v_{n+1} = v'$ **else** $v_{n+1} = v_n$ $\triangleright$ Accept/reject v'
- 4: **UPDATEFLAG** = **false**
- 5: **if** $\text{rem}(n+1, n_{\text{lag}}) = 0$ & $m < n_{\text{max}}$ & $d_{\mathcal{F}} \geq \Delta_{\text{LIS}}$ **then** $\triangleright$ Update LIS using selected samples
- 6: $\text{UPDATELIS}(\Theta_m, \Xi_m, v_{n+1}, \Theta_{m+1}, \Xi_{m+1}, \Theta_{r'}, \Xi_{r'})$
- 7: $\text{PROJECTCOV}(\Theta_r, \Sigma_r, \Theta_{r'}, \Sigma_{r'})$
- 8: Update the LIS convergence diagnostic $d_{\mathcal{F}}$ as in (33)
- 9: $\Theta_r = \Theta_{r'}, \Xi_r = \Xi_{r'}, \Sigma_r = \Sigma_{r'}, m = m + 1$
- 10: **UPDATEFLAG** = **true**
- 11: **else** $\triangleright$ The LIS is fixed
- 12: Perform rank-1 update of the empirical covariance Σ_r given $\Theta_r^* v_{n+1}$
- 13: **if** $\text{rem}(n+1, n_b) = 0$ **then** **UPDATEFLAG** = **true**
- 14: **end if**
- 15: **if** **UPDATEFLAG** **then**
- 16: $\text{UPDATECOV}(\Theta_r, \Sigma_r, \Psi_r, D_r)$
- 17: Update the operators $\{A, B, G\}$
- 18: **end if**

Metropolis-adjusted Langevin proposals. Purely local methods, such as stochastic Newton [19] and manifold MCMC [20], are much more computationally expensive than both LI-Prior and LI-Langevin, as they require high-order derivatives of the posterior at every step.

Another non-negligible computational burden comes from the adaptive construction of LIS (Steps 5–10), where each iteration involves a number of linearized forward model and adjoint model evaluations for computing the eigendecomposition of the local GNH, as well as the incremental update of the factorized expected GNH in Algorithm 1. Here we employ a fixed-size adaptation strategy that limits the total number of local GNH eigendecompositions to be no more than n_{max} , and we terminate the iterative updates once the subspace convergence diagnostic (33) drops below a specified threshold Δ_{LIS} . Furthermore, to ensure that essentially uncorrelated posterior samples are used to estimate the expected GNH, sub-sampling with lag n_{lag} is used in Step 7. Since the LIS basis is low-rank, the computational costs of updating the posterior covariance (Step 16) and the operators (Step 17) are only proportional to the discretization dimension N , and thus do not dominate the cost of the forward solve. These costs are further suppressed by only renewing these two terms every n_b iterations during MCMC sampling. In the numerical experiments below, we choose $n_b = 50$.

We also note that once the additive construction of LIS terminates, we simulate the Markov chain using the operator-weighted proposal built from a *fixed* LIS, where the empirical covariance is henceforth updated only within this fixed r -dimensional subspace. In this situation, existing adaptive MCMC theory [35, 42, 36, 43, 44] can be used to establish ergodicity of the chain.

Algorithm 5 Adaptive function space MCMC with operator-weighted proposals, using Metropolis-within-Gibbs updates.

Require: During the LIS construction, we have $\{\Theta_m, \Xi_m\}$ and $d_{\mathcal{F}}$ as in Algorithm 4.

Require: At step n , given the state v_n , LIS basis Θ_r , projected empirical posterior covariance Σ_r , and operators $\{A_r, B_r, G_r\}$ and $\{A_{\perp}, B_{\perp}\}$ induced by $\{\Psi_r, D_r, \Delta\tau_r, \Delta\tau_{\perp}\}$, one step of the algorithm is:

- 1: Compute a LIS candidate $v' = q(v_n, \cdot; A_r, B_r, G_r)$ using either MGLI-Prior or MGLI-Langevin
 - 2: Compute the acceptance probability $\alpha(v_n, v')$
 - 3: **if** $\text{Uniform}(0, 1] < \alpha(v_n, v')$ **then** $\tilde{v} = v'$ **else** $\tilde{v} = v_n$ **end if** $\triangleright$ Accept/reject v'
 - 4: Compute a CS candidate $v'' = q(\tilde{v}, \cdot; A_{\perp}, B_{\perp})$ using (41b)
 - 5: Compute the acceptance probability $\alpha(\tilde{v}, v'')$
 - 6: **if** $\text{Uniform}(0, 1] < \alpha(\tilde{v}, v'')$ **then** $v_{n+1} = v''$ **else** $v_{n+1} = \tilde{v}$ $\triangleright$ Accept/reject v''
 - 7: Follow Steps 4–14 of Algorithm 4 to update the LIS and sample covariance within the LIS
 - 8: **if** UPDATEFLAG **then**
 - 9: UPDATECOV($\Theta_r, \Sigma_r, \Psi_r, D_r$)
 - 10: Update the operators $\{A_r, B_r, G_r\}$ and $\{A_{\perp}, B_{\perp}\}$
 - 11: **end if**
-

Just as the all-at-once proposals of Section 3.3.1 are integrated into the adaptive framework of Algorithm 4, the Metropolis-within-Gibbs proposals given in Section 3.3.2 can be integrated into Algorithm 5. Here the operators $\{A_r, B_r, G_r\}$ and $\{A_{\perp}, B_{\perp}\}$ —for the LIS update and the CS update, respectively—are determined by $\{\Psi_r, D_r, \Delta\tau_r, \Delta\tau_{\perp}\}$. Two proposal and accept/reject steps are involved in each iteration of Algorithm 5, and thus it doubles the number of likelihood evaluations per iteration relative to Algorithm 4. The computational cost of all the other components remains the same.

3.5. Benchmark algorithms

In the numerical examples of Sections 4 and Section 5, we will benchmark our operator-weighted proposals by comparing their performance with two representative proposals from the literature, one that is dimension independent and another that exploits posterior Hessian information.

Proposal 3.5 (pCN-RW: Prior-preconditioned Crank-Nicolson random walk proposal [10]). *By setting $\gamma = 0$ in (9) and applying the transformation in Definition 3.2, we obtain the proposal:*

$$v' = av + \sqrt{1 - a^2}\xi.$$

where $a \in (-1, 1)$.

Note that the gradient term could have been included by setting $\gamma = 1$ in (9). In the numerical examples described below, however, the performance of such a pCN Langevin proposal [11] is very close to that of the pCN-RW proposal. For brevity, we will therefore only report on pCN-RW in our numerical comparisons.

Proposal 3.6 (H-Langevin: Hessian-preconditioned explicit Langevin proposal). *Let H_M denote the Hessian of the data-misfit functional evaluated at the MAP estimate of the parameter v . Precondition the Langevin SDE (7) with the inverse Hessian of the OMF, $(H_M + I)^{-1}$, and discretize*

this SDE explicitly, arriving at a proposal of the form:

$$\begin{aligned} v' &= (\mathbf{I} - \Delta\tau(\mathbf{H}_M + \mathbf{I})^{-1}) v - \Delta\tau(\mathbf{H}_M + \mathbf{I})^{-1} D_v \Phi(v; y) + \sqrt{2\Delta\tau(\mathbf{H}_M + \mathbf{I})^{-1}} \xi \\ &= \mathbf{A}_H v - \mathbf{G}_H D_v \Phi(v; y) + \mathbf{B}_H \xi. \end{aligned}$$

Now take a low rank approximation of $\mathbf{H}_M$,

$$\mathbf{H}_M \approx \Psi_M \mathbf{D}_M \Psi_M^*.$$

to obtain

$$\begin{aligned} \mathbf{A}_H &= \mathbf{I} - \Delta\tau (\Psi_M (-\mathbf{D}_M (\mathbf{D}_M + \mathbf{I}_r)^{-1}) \Psi_M^* + \mathbf{I}), \\ \mathbf{B}_H &= \sqrt{2\Delta\tau} \left(\Psi_M \left((\mathbf{D}_M + \mathbf{I}_r)^{-1/2} - \mathbf{I}_r \right) \Psi_M^* + \mathbf{I} \right), \\ \mathbf{G}_H &= \Delta\tau (\Psi_M (-\mathbf{D}_M (\mathbf{D}_M + \mathbf{I}_r)^{-1}) \Psi_M^* + \mathbf{I}). \end{aligned}$$

This H-Langevin proposal essentially combines preconditioned MALA [45] with the low-rank Hessian approximation used in the stochastic Newton algorithm [19]. But instead of using a location-dependent preconditioner as in simplified manifold MALA [20] and stochastic Newton, it uses a constant preconditioner $(\mathbf{H}_M + \mathbf{I})^{-1}$, thus avoiding the computational burden of evaluating the local Hessian at each iteration. The same principle has recently been used in [46] to modify the stochastic Newton algorithm. The main difference between the H-Langevin proposal and the variants of stochastic Newton proposed in [46] is that H-Langevin employs an adjustable discretization step $\Delta\tau$, whereas [46] uses $\Delta\tau = 1$.

4. Example 1: Elliptic PDE

Our first numerical example is an inverse problem of an elliptic PDE, i.e., inferring the transmissivity field of the Poisson equation. We use this test case to compare the efficiency of the new operator-weighted proposals described in Section 3.3 with the benchmark proposals of Section 3.5. We also demonstrate the advantages of using multiple posterior samples to construct the LIS and the invariance of the LIS under grid refinement.

4.1. Problem setup

Consider the spatial domain $\Omega = [0, 1]^2$ with boundary $\partial\Omega$. We denote the spatial coordinate by $s \in \Omega$. Let $\kappa(s)$ be the transmissivity field, $p(s)$ be the potential function, and $f(s)$ the forcing term. The potential function for a given realization of the transmissivity field is governed by

$$\begin{cases} -\nabla \cdot (\kappa(s) \nabla p(s)) &= f(s), & s \in \Omega \\ \langle \kappa(s) \nabla p(s), \vec{n}(s) \rangle &= 0, & s \in \partial\Omega \end{cases} \quad (44)$$

where $\vec{n}(s)$ is the outward normal to the boundary. To make a well-posed boundary value problem, a further condition

$$\int_{\partial\Omega} p(s) dl(s) = 0, \quad (45)$$

is imposed. The source/sink term $f(s)$ is defined by the superposition of four weighted Gaussian plumes with standard deviation 0.05, centered at $[0.3, 0.3]$, $[0.7, 0.3]$, $[0.7, 0.7]$, $[0.3, 0.7]$, with weights

$\{2, -3, -2, 3\}$. The system of equations (44) is solved by the finite element method with bilinear elements on a uniform 40×40 grid.

The transmissivity field is endowed with a log-normal prior distribution, i.e.,

$$\kappa(s) = \exp(u(s)), \text{ and } u(s) \sim \mathcal{N}(0, C), \quad (46)$$

where the covariance operator C is defined through an exponential kernel function

$$c(s, s') = \sigma_u^2 \exp\left(-\frac{\|s - s'\|}{2s_0}\right), \text{ for } s, s' \in \Omega. \quad (47)$$

In this example, we set the prior standard deviation $\sigma_u = 1.25$ and the correlation length $s_0 = 0.0625$. To make the inverse problem more challenging, we use a “true” transmissivity field that is not directly drawn from the prior distribution. The true transmissivity field—i.e., the field used to generate noisy observations—and the corresponding potential are shown in Figures 1(a) and (b), respectively.

Figure 1: Setup of the elliptic example. (a) The true transmissivity field. (b) The potential $p(s)$ computed from the transmissivity field defined in (a), with measurement locations indicated by circles. (c) Data sets: the dashed line represents the noise-free potential at each measurement location, while crosses and circles are observations of the potential perturbed with Gaussian noise of signal-to-noise ratios 10 and 50, respectively.

Partial observations of the potential function are collected at 25 measurement sensors as shown by the circles in Figure 1(b). The corresponding observation operator M yields the outputs

$$y = Mp(s) + e, \quad e \sim \mathcal{N}(0, \sigma^2 I_{25}),$$

where the observed data is $y \in \mathbb{R}^{25}$. Since amount of information carried in the likelihood function affects the structure of the LIS, we use two cases with different noise magnitudes to illustrate this relationship. Let the signal-to-noise ratio (SNR) be defined as $\max_s \{u(s)\} / \sigma$. The first data set is generated with signal to noise ratio 10 (SNR10), and the second set is generated with signal to noise ratio 50 (SNR50). Both data sets are shown in Figure 1(c). *We note that changing the signal-to-noise ratio will have an impact on the LIS dimension similar to changing the number of measurements; both affect the amount of information carried in the data.*

4.2. Sampling efficiency

We now compare the posterior sampling performance of our operator-weighted proposals (Proposals 3.1–3.4) with that of the pCN-RW and H-Langevin proposals (Proposals 3.5 and 3.6). To

build the LIS, the adaptive sampling procedure in Section 3.2.4 is run for 2×10^5 MCMC iterations, from which 1000 posterior samples are selected, one every $n_{\text{lag}} = 200$ iterations, to compute the expected Hessian. Then we sample the posterior distribution for 10^6 MCMC iterations using the operator-weighted proposals and the benchmark proposals. Samples from the second half of these iterations are used to estimate the autocorrelation of the Onsager-Machlup functional (4) and of the parameter u projected onto selected eigenfunctions of the prior covariance operator. As the metrics of sampling efficiency below will indicate, these sample sizes are significantly larger than one might use in practice, particularly with the DILI proposals. Our goal here, however, is to characterize the mixing properties of each chain without any potential bias resulting from burn-in and finite sample size. Also, our comparisons of mixing along particular directions in the parameter space use projections onto prior covariance eigenfunctions because these projections provide a *common* basis for comparison—in contrast to projections w onto modes of the LIS, which can be different for different proposals. Moreover, evaluating performance on LIS modes could conceivably yield results that are artificially favorable to the DILI proposals.

For the SNR10 case, Figure 2 shows the autocorrelation functions of selected components of the v parameter. Note that we chose the square root $\Gamma_{\text{pr}}^{1/2}$ in Definition 3.2 such that the selected components of v correspond to projections of u onto the 1st, 10th, 100th, and 1000th eigenfunctions of the prior covariance (i.e., Karhunen-Loève modes). Overall, MGLI-Langevin outperforms the other proposals, in some cases quite dramatically. On the low-index eigenfunctions (1 and 10), the operator-weighted proposals and the H-Langevin proposal yield similar mixing, as indicated by the decay of the autocorrelation. Yet the operator-weighted proposals outperform H-Langevin on the higher-index eigenfunctions, as these directions in the parameter space tend to be less constrained by the data; indeed, they are more aligned with the CS. This effect can already be observed for 100th eigenfunction and becomes more significant for the 1000th. Chains produced by pCN-RW decorrelate very slowly compared to the other samplers.

The autocorrelation of the OMF, shown in Figure 3 along with trace plots for two of the chains, provides an alternative summary of sampling performance. Again, we observe that the operator-weighted proposals yield shorter autocorrelation times than H-Langevin and pCN-RW.

Figure 4 compares the MGLI-Langevin, H-Langevin, and pCN-RW samplers for the SNR50 case. In parameter directions corresponding to the smoother prior covariance eigenfunctions (e.g., components 1, 10, and 100 of v), the MGLI-Langevin and H-Langevin proposals produce very similar autocorrelations. Yet, as in the SNR10 case, MGLI-Langevin significantly outperforms H-Langevin for parameter directions corresponding to the high-index eigenfunctions, e.g., v_{1000} . The pCN-RW chain again decorrelates very slowly compared to the other proposals. Looking at the OMF in Figure 5, we observe relative performances similar to the SNR10 case: the MGLI-Langevin proposal outperforms both pCN-RW and H-Langevin.

In both of these test cases, H-Langevin and the operator-weighted proposals have comparable sampling efficiency along parameter directions that are dominated by the likelihood; this is not surprising, as both use Hessian information from the log-likelihood to scale proposed moves. For directions along which the likelihood has a relatively weaker influence than the prior, however, the operator-weighted proposals have better sampling efficiency than H-Langevin. To further explore this pattern, Figure 6 shows the lag-1 autocorrelation of all 1600 components of v , for both the SNR10 and SNR50 cases (subfigures 6(a) and 6(b), respectively). The star symbols represent the lag-1 autocorrelation of MCMC samples produced by H-Langevin while the square symbols represent results from MGLI-Langevin. We observe that the lag-1 autocorrelations produced by

Figure 2: Autocorrelation functions of selected components of the v parameter produced by various proposals, SNR10 case. (a) Projection onto the first eigenfunction of the prior covariance; (b) Projection onto the second eigenfunction; (c) Projection onto the 100th eigenfunction; (d) Projection onto the 1000th eigenfunction.

Figure 3: Summaries of MCMC mixing via the Onsager-Machlup functional, SNR10 case. (a) Trace plot of OMF values produced by MCMC with the MGLI-Langevin proposal. (b) Trace plot of OMF values produced by MCMC with the pCN-RW proposal. (c) Autocorrelation of OMF values produced by various proposals.

H-Langevin stay almost constant across components of v (i.e., across the prior eigenfunctions). In comparison, the lag-1 autocorrelations produced by MGLI-Langevin decrease significantly for the components of v corresponding to the high-index eigenfunctions, and are lower on average than the autocorrelations from H-Langevin. The other operator-weighted proposals produce similar results, though the lag-1 autocorrelations of the lowest-index eigenfunctions are not always smaller with

Figure 4: Autocorrelation functions of selected components of the v parameter produced by various proposals, SNR50 case.

Figure 5: Summaries of MCMC mixing via the OMF, SNR50 case.

the operator-weighted proposals than with H-Langevin. For brevity, these results are not shown.

In the SNR10 case, the CPU times spent, after adaptive LIS construction, on 10^6 iterations of MCMC with the LI-Prior, LI-Langevin, pCN-RW, and H-Langevin proposal are 0.88×10^5 seconds, 1.12×10^5 seconds, 0.85×10^5 seconds and 1.25×10^5 seconds, respectively. Here the CPU time of LI-Prior is comparable to that of pCN-RW, and the CPU time of LI-Langevin is comparable to that of H-Langevin. We also note that the Langevin proposals do not spend much additional CPU time on gradient evaluations compared to LI-Prior and pCN-RW, since in the present example

Figure 6: The lag-1 autocorrelation of each component of the v parameter. Star and square symbols are the results of H-Langevin and MGLI-Langevin, respectively. (a) SNR10 case. (b) SNR50 case.

the underlying PDE is self-adjoint and hence the forward solution can be recycled for the adjoint evaluation. The CPU times spent on 10^6 MCMC steps with MGLI-Prior and MGLI-Langevin are 1.70×10^5 seconds and 2.10×10^5 seconds, respectively. These essentially double the CPU times of LI-Prior and LI-Langevin, as expected. The CPU time spent on the adaptive construction of the LIS using 1000 posterior samples is about 0.23×10^5 seconds, which is not significant relative to the time spent on MCMC sampling. We note that using 1000 samples here is a rather conservative choice. From our experience in [23] and the convergence study shown next (see Figure 9), about 500 samples are sufficient to estimate an LIS with good accuracy. For the SNR50 case, the CPU times spent on each of the proposals and on the LIS construction are similar to those of the SNR10 case, and thus are not reported for brevity.

4.3. Global versus local LIS

It is useful to quantify the impact of “globalizing” the LIS on sampling efficiency. In particular, we wish to contrast the performance of operator-weighted proposals constructed from a global LIS, based on the posterior expected Hessian (as in Section 3.2.2), with that of operator-weighted proposals constructed from a single local LIS, based on the Hessian at the MAP. For brevity, we limit our comparisons to Proposal 3.4 (MGLI-Langevin). The global LIS is constructed using the adaptive sampling strategy detailed at the start of Section 4.2. Results produced with the global LIS are labeled ‘Adapt-LIS’ in the figures below. Results obtained with the MGLI-Langevin proposal employing a local LIS at the MAP are denoted ‘MAP-LIS.’ Note that the local LIS in the current problem has dimension at most 25, since the inverse problem has 25 observations. The global LIS can be much larger, since it accounts for posterior variation in the dominant eigenspace of the Hessian. In the current setup, the global LIS for the SNR10 case has dimension $r = 66$ while for the SNR50 case it has dimension $r = 193$; both of these are obtained with truncation thresholds $\tau_{\text{loc}} = \tau_{\text{g}} = 0.01$.

Figures 7 and 8 summarize relative sampling performance for both the SNR10 and SNR50 cases. We show the autocorrelation of the OMF for each algorithm and the lag-1 autocorrelations of each component of v . In both test cases, using a global LIS produces better mixing than a single local LIS. The improvement in mixing for the OMF seems slightly more pronounced for the larger-noise case (SNR10) than the smaller-noise case (SNR50). This may be due to the broader posterior of the former; since the nonlinear forward model is the same in both cases, the variation of the Hessian becomes more significant when wider ranges of the parameter space are explored. On the other hand, the lag-1 autocorrelations show greater improvements in the SNR50 case. Here, the higher

dimension of the global LIS relative to the local LIS (193 to 25 in the SNR50 case versus 66 to 25 in the SNR10 case) may play a role. The global LIS is larger in the SNR50 case because the data carry more information.

Figure 7: Sampling performance of Adapt-LIS (solid line) and MAP-LIS (dashed line), SNR10 case. (a) Autocorrelation of the OMF. (b) Lag-1 autocorrelation of each component of the v parameter.

Figure 8: Sampling performance of Adapt-LIS (solid line) and MAP-LIS (dashed line), SNR50 case.

4.4. LIS under grid refinement

We now explore how grid refinement affects the dimensionality and structure of the global LIS, as well as the convergence of the adaptive procedure for constructing it. In the examples below, the transmissivity field $\kappa(s)$ and the associated potential function $p(s)$ are discretized on 40×40 , 80×80 , and 120×120 grids.

Figure 9 shows the dimension of the LIS and the convergence diagnostic (33) versus the number of samples used in the adaptive construction process, for the SNR50 case only. For all three discretizations, the distance (33) between likelihood-informed subspaces at adjacent iterations drops by several orders of magnitude over the course of the adaptive sampling procedure, as shown in Figure 9(b). The rates of convergence of this distance are comparable for all three discretizations. After a few hundred samples, the dimensionality of LIS also converges for all three discretizations. Note that the 40×40 grid yields a slightly higher-dimensional LIS than the two more refined grids: at the end of the adaptive procedure, the LIS of the 40×40 grid has dimension $r = 193$, while the 80×80 and 120×120 grids yield $r = 178$ and $r = 176$, respectively. This effect can be ascribed to discretization errors on the 40×40 grid; since the forward model converges under grid refinement, we expect the dimension of the associated LIS also to converge. The first five basis vectors of

Figure 9: The dimension of the LIS and the subspace convergence diagnostic (33) versus the number of samples used in the adaptive algorithm; SNR50 case. Dotted, dashed, and solid lines represent the 40×40 , 80×80 , and 120×120 grids, respectively. (a) The dimension of the LIS. (b) The convergence diagnostic.

the LIS for the various discretizations are plotted in Figure 10. Similar structures are observed at different levels of grid refinement.

Figure 10: The first five basis functions of the LIS for three different levels of grid refinement; SNR50 case. Top row: 40×40 grid, middle row: 80×80 grid, and bottom row: 120×120 grid.

5. Example 2: Conditioned diffusion

In this section, we use noisy point-wise observations of the path of a particle, diffusing in a double-well potential, to infer the driving force on the particle and hence its pushforward to the path itself. This example is motivated by applications in molecular dynamics.

5.1. Problem setup

Consider a stochastic process $p : [0, T] \rightarrow \mathbb{R}$ governed by the following Langevin SDE:

$$dp_t = f(p_t)dt + du_t, \quad p_0 = 0, \quad (48)$$

where $f : \mathbb{R} \rightarrow \mathbb{R}$ is globally Lipschitz and du_t is an increment of the Brownian motion $u \sim \mu_0 = \mathcal{N}(0, C)$, where $C(t, t') = \min(t, t')$. Let the function f have the following form

$$f(p) := \beta p(1 - p^2)/(1 + p^2),$$

with $\beta > 0$. The corresponding potential is $E(p) = -\int_{-\infty}^p f(s)ds$. From any initial condition p_0 , the state p_t approaches the invariant measure of the SDE, which has a density $Z^{-1} \exp(-2E(p))$ with respect to Lebesgue measure, where $Z = \int_{\mathbb{R}} \exp(-2E(s)) ds$. Since the potential E has a double-well shape, the invariant measure is bimodal and paths of (48) will transition from one well to the other, with a probability dependent on the magnitude of the stochastic forcing and on the scale β of the potential. We will use $\beta = 10$ in the examples below. This model is ubiquitous in the sciences, perhaps most notably in molecular dynamics where it represents the motion of a particle with negligible mass trapped in an energy potential E with thermal fluctuations represented by the Brownian forcing.

We refer the reader to [5] for a proof that the map $u \mapsto p$ is continuous and differentiable from $C([0, T], \mathbb{R}) \rightarrow C([0, T], \mathbb{R})$. The fact that $\mu_0(C([0, T], \mathbb{R})) = 1$ follows from the well-known property that continuous functions have probability one under Wiener measure. Therefore, this Bayesian inverse problem is well-defined [5], and it fits into the framework developed in Section 3. The observation operator is defined by $Mp := [p_{t_1}, p_{t_2}, \dots, p_{t_{20}}]^T$, and we let

$$y = Mp + e, \quad e \sim \mathcal{N}(0, \sigma^2 I_{20}),$$

where $\sigma = 0.1$ and the observation times t_i are equispaced within the interval $[0, 10]$. An Euler-Maruyama scheme is used for integration with $\Delta t = 10^{-2}$, and hence the dimension of our approximation of this infinite-dimensional pathspace is $N = 1000$.

5.2. Numerical results

We now evaluate the efficiency of the MCMC proposals from Section 3.3 and Section 3.5 on the conditioned diffusion example. Figure 11 illustrates the solution to the forward problem taken as the truth, the noisy observations of this particular path, and the results of subsequent conditioning—i.e., aspects of the posterior. In particular, we show the posterior mean, marginals of the posterior density at each time step, and the 0.05 and 0.95 quantiles of the local marginals. Note that the time interval $t \in [0, 10]$ considered here is long enough to capture two transitions between the potential wells. Our numerical experiments use the same chain lengths and burn-in intervals as in the elliptic example, detailed in Section 4.2. We also use the same number of iterations as before to adaptively construct the global LIS.

The numerical results in Figures 12–14 highlight the strong performance of the DILI samplers developed in this paper. Figure 12 summarizes the sampling efficiency of each MCMC scheme via traces and autocorrelations of the OMF. Not only do the DILI algorithms outperform pCN, but they also show dramatically improved mixing over H-Langevin; here, H-Langevin performs almost as poorly as pCN. This contrasts with the elliptic PDE cases examined in Section 4, where the performance of H-Langevin was reasonably close to that of the DILI proposals. Plots of lag-1 autocorrelation in Figure 13 reinforce this notion, showing much faster decay of lag-1 autocorrelation with MGLI-Langevin than with H-Langevin across the first 1000 eigenfunctions of the prior.

Figure 14 shows how mixing performance is improved by using a global LIS rather than a local LIS based on the Hessian at the MAP. Both traces in this figure use the MGLI-Langevin algorithm, but those marked ‘Adapt-LIS’ use a global LIS constructed via the adaptive algorithm

Figure 11: Setup of the conditioned diffusion example. Top: the true trajectory, the noisy observations, the posterior mean of the trajectory, the 0.05 and 0.95 quantiles of the trajectory marginalized at each time. The posterior density of the trajectory marginalized at each time is shown as the grayscale background image. Bottom: the posterior density of the trajectory marginalized at selected times.

of Section 3.2.4. The global LIS clearly yields a much faster decay of the OMF autocorrelation and better mixing, as measured by lag-1 autocorrelations, across all components of v .

6. Conclusions

This paper has introduced a general class of operator-weighted MCMC proposals that are well defined on function space. In Bayesian inference problems that require exploring a posterior distribution over functions, these proposals yield MCMC sampling performance that is invariant under refinement of the discretization of these functions, and hence dimension-independent. While this class includes earlier dimension-independent algorithms [10, 11, 22] as particular cases, we use the versatility offered by our new family of proposals to design samplers that exploit posterior structure wherein departures from the prior—including non-Gaussianity—are concentrated on a finite number of directions, captured by a global likelihood-informed subspace (LIS). The global LIS is identified by approximating the posterior expectation of the Hessian of a preconditioned data-misfit functional. A further adaptive strategy based on this decomposition yields computationally affordable

Figure 12: Summaries of MCMC mixing for the conditioned diffusion example. (a) Trace plot of the OMF produced by MGLI-Langevin. (b) Trace plot of the OMF produced by pCN-RW. (c) Autocorrelation of the OMF, as produced by various proposals.

Figure 13: The lag-1 autocorrelation of each component of the v parameter for the conditioned diffusion example. Star and square symbols are produced by H-Langevin and MGLI-Langevin proposals, respectively.

Figure 14: Sampling performance of MGLI-Langevin with a global LIS (Adapt-LIS, solid line) versus a local LIS (MAP-LIS, dashed line) for the conditioned diffusion example. (a) Autocorrelation of the OMF. (b) Lag-1 autocorrelation of each component of the v parameter.

approximations of the posterior covariance. All of this geometric information is used to construct *dimension-independent* and *likelihood-informed* (DILI) proposal distributions; the four variations studied here (Proposals 3.1–3.4) also make use of local gradient information and Metropolis-within-Gibbs updates. Numerical experiments for two nonlinear inverse problems suggest that these DILI

samplers offer significant gains in sampling efficiency over current state-of-the-art algorithms.

This work can be extended in many directions. First, one can certainly consider alternative constructions of the LIS. Of particular interest are parallel methods for constructing the global LIS, strategies for infinite batch-updating of the LIS that are compatible with adaptive MCMC, and even directed strategies for optimally exploring the variation of the preconditioned data-misfit Hessian. If only gradients of the data-misfit function are available, one might resort to so-called active subspace methods [47, 48] as an alternative way of building up LIS-like information. If gradients are also not available, then a different approach might be taken, for example adaptively building covariance information from samples and using regularized estimates to build an approximate LIS.

Second, many other operator-weighted proposal constructions are possible. For example, given the approximation of the empirical posterior covariance $\Sigma = \Psi_r (D_r - I_r) \Psi_r^* + I$ defined in (30), we can define the operators

$$\begin{aligned} A &= (a_r - a_\perp) \Psi_r \Psi_r^* + a_\perp I \\ B &= \Psi_r (b_r D_r^{1/2} - b_\perp I_r) \Psi_r^* + b_\perp I \\ G &= 0, \end{aligned}$$

where $a_r, a_\perp \in (-1, 1)$, and b_r and $b_\perp$ take values $\sqrt{1 - a_r^2}$ and $\sqrt{1 - a_\perp^2}$, respectively. Given a set of posterior samples $\{u_1, \dots, u_n\}$, by setting the reference variable m_{ref} to be the difference of the empirical posterior mean and the prior mean, i.e., $m_{\text{ref}} = \frac{1}{n} \sum_{k=1}^n u_k - m_0$, the resulting proposal (13) has a Gaussian approximation of the posterior as its invariant distribution.

Finally, while this work has focused on the design of global operators (e.g., A, B, and G that do not depend on the local parameter value), there is much room to combine global and local operators. As an example of the latter, [49] uses a semi-implicit discretization of a locally-preconditioned Langevin equation to derive a proposal that, under appropriate assumptions, yields a dimension-independent MCMC algorithm. The present work may provide a rather general way of constructing local preconditioners that satisfy these assumptions. In particular, the partition of parameter space induced by the LIS provides an opportunity to introduce a variety of efficient proposals on the LIS (e.g., locally-preconditioned manifold MALA, even RMHMC) while preserving dimension independence through a suitable discretization of a Langevin equation on the CS. Combining proposals in this manner extends beyond the notion of local preconditioning. We leave a fuller development of these ideas to future work.

Acknowledgments

T. Cui and Y. M. Marzouk acknowledge financial support from the US Department of Energy, Office of Advanced Scientific Computing Research (ASCR), under grant number DE-SC0009297, as part of the DiaMonD Multifaceted Mathematics Integrated Capability Center. K. J. H. Law was supported by the King Abdullah University of Science and Technology (KAUST), and is a member of the KAUST SRI-UQ Center at KAUST.

References

- [1] A. Tarantola, Inverse Problem Theory and Methods for Model Parameter Estimation, Society for Industrial Mathematics, Philadelphia, 2005.

- [2] J. P. Kaipio, E. Somersalo, *Statistical and Computational Inverse Problems*, volume 160, Springer, New York, 2004.
- [3] A. M. Stuart, Inverse problems: a Bayesian perspective, *Acta Numerica* 19 (2010) 451–559.
- [4] A. Beskos, O. Papaspiliopoulos, G. O. Roberts, P. Fearnhead, Exact and computationally efficient likelihood based estimation for discretely observed diffusion processes (with discussion), *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68 (2006) 333–382.
- [5] M. Hairer, A. M. Stuart, J. Voss, Signal processing problems on function space: Bayesian formulation, stochastic PDEs and effective MCMC methods, in: D. Crisan, B. Rozovsky (Eds.), *The Oxford Handbook of Nonlinear Filtering*, Oxford University Press, 2011.
- [6] G. O. Roberts, A. Gelman, W. R. Gilks, Weak convergence and optimal scaling of random walk Metropolis algorithms, *Annals of Applied Probability* 7 (1997) 110–120.
- [7] G. O. Roberts, J. S. Rosenthal, Optimal scaling of various Metropolis-Hastings algorithms, *Statistical Science* 16 (2001) 351–367.
- [8] J. C. Mattingly, N. Pillai, A. M. Stuart, Diffusion limits of the random walk Metropolis algorithm in high dimensions, *Annals of Applied Probability* 22 (2012) 881–930.
- [9] N. S. Pillai, A. M. Stuart, A. H. Thiery, Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions, *Annals of Applied Probability* 22 (2012) 2320–2356.
- [10] A. Beskos, G. O. Roberts, A. M. Stuart, J. Voss, MCMC methods for diffusion bridges, *Stochastic Dynamics* 8 (2008) 319–350.
- [11] S. L. Cotter, G. O. Roberts, A. M. Stuart, D. White, MCMC methods for functions: modifying old algorithms to make them faster, *Statistical Science* 28 (2013) 424–446.
- [12] N. Chopin, A sequential particle filter method for static models, *Biometrika* 89 (2002) 539–552.
- [13] P. D. Moral, A. Doucet, A. Jasra, Sequential Monte Carlo samplers, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68 (2006) 411–436.
- [14] C. Andrieu, A. Doucet, R. Holenstein, Particle Markov chain Monte Carlo methods, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72 (2010) 269–342.
- [15] M. B. Giles, Multi-level Monte Carlo path simulation, *Operations Research* 56 (2008) 607–617.
- [16] A. Beskos, A. Jasra, N. Kantas, Sequential Monte Carlo methods for high-dimensional inverse problems: A case study for the Navier-Stokes equations, preprint (2013) arXiv:1307.6127.
- [17] V. H. Hoang, C. Schwab, A. M. Stuart, Complexity analysis of accelerated MCMC methods for Bayesian inversion, *Inverse Problems* 29 (2013) 085010.
- [18] C. Ketelsen, R. Scheichl, A. L. Teckentrup, A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow, preprint (2013) arXiv:1303.7343.
- [19] J. Martin, L. C. Wilcox, C. Burstedde, O. Ghattas, A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion, *SIAM Journal on Scientific Computing* 34 (2012) A1460–A1487.
- [20] M. Girolami, B. Calderhead, Riemann manifold Langevin and Hamiltonian Monte Carlo methods, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73 (2011) 123–214.
- [21] A. Spantini, A. Solonen, T. Cui, J. Martin, L. Tenorio, Y. Marzouk, Optimal low-rank approximation of linear Bayesian inverse problems, preprint (2014). arXiv:1407.3463.
- [22] K. J. Law, Proposals which speed up function-space MCMC, *Journal of Computational and Applied Mathematics* 262 (2014) 127–138.
- [23] T. Cui, J. Martin, Y. Marzouk, A. Solonen, A. Spantini, Likelihood-informed dimension reduction for nonlinear inverse problems, *Inverse Problems* 30 (2014) 114015.
- [24] M. Dashti, K. J. H. Law, A. M. Stuart, J. Voss, MAP estimators and posterior consistency in Bayesian nonparametric inverse problems, *Inverse Problems* 29 (2013) 095017.
- [25] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, Equation of state calculations by fast computing machines, *Journal of Chemical Physics* 21 (1953) 1087–1092.
- [26] W. Hastings, Monte Carlo sampling using Markov chains and their applications, *Biometrika* 57 (1970) 97–109.
- [27] L. Tierney, A note on Metropolis-Hastings kernels for general state spaces, *Annals of Applied Probability* 8 (1998) 1–9.
- [28] J. E. Besag, Discussion of “representations of knowledge in complex systems” by U. Grenander and M. I. Miller, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 56 (1994) 549–603.
- [29] G. O. Roberts, J. S. Rosenthal, Optimal scaling of discrete approximations to Langevin diffusions, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 60 (1998) 255–268.
- [30] M. Hairer, A. M. Stuart, S. Vollmer, Spectral gaps for a Metropolis-Hastings algorithm in infinite dimensions,

- Annals of Applied Probability 24 (2014) 2455–2490.
- [31] H. P. Flath, L. C. Wilcox, V. Akcelik, J. Hill, B. van Bloemen Waanders, O. Ghattas, Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations, *SIAM Journal on Scientific Computing* 33 (2011) 407–432.
 - [32] T. Bui-Thanh, C. Burstedde, O. Ghattas, J. Martin, G. Stadler, L. C. Wilcox, Extreme-scale UQ for Bayesian inverse problems governed by PDEs, in: *ACM/IEEE Supercomputing SC12*, Salt Lake City, Utah, 2012.
 - [33] T. Bui-Thanh, O. Ghattas, J. Martin, G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems. Part I: The linearized case, with application to global seismic inversion, *SIAM Journal on Scientific Computing* 35 (2013) A2494–A2523.
 - [34] W. Förstner, M. Boudewijn, A metric for covariance matrices, in: *Geodesy: The Challenge of the 3rd Millennium*, Springer, Berlin Heidelberg, 2003, pp. 299–309.
 - [35] H. Haario, E. Saksman, J. Tamminen, An adaptive Metropolis algorithm, *Bernoulli* 7 (2001) 223–242.
 - [36] Y. F. Atchadé, An adaptive version for the Metropolis adjusted Langevin algorithm with a truncated drift, *Methodology and Computing in Applied Probability* 8 (2006) 235–254.
 - [37] G. H. Golub, C. F. V. Loan, *Matrix Computations*, JHU Press, 2012.
 - [38] N. Halko, P. Martinsson, J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review* 53 (2011) 217–288.
 - [39] E. Liberty, F. Woolfe, P. G. Martinsson, V. Rokhlin, M. Tygert, Randomized algorithms for the low-rank approximation of matrices, *Proceedings of the National Academy of Sciences* 104 (2007) 20167–20172.
 - [40] M. Brand, Incremental singular value decomposition of uncertain data with missing values, in: *Computer Vision – ECCV 2002*, Springer, Berlin Heidelberg, 2002, pp. 707–720.
 - [41] F. Li, Q. Dai, W. Xu, G. Er, Weighted subspace distance and its applications to object recognition and retrieval with image sets., *Signal Processing Letters, IEEE* 16 (2009) 227–230.
 - [42] C. Andrieu, E. Moulines, On the ergodicity properties of some adaptive MCMC algorithms, *The Annals of Applied Probability* 16 (2006) 1462–1505.
 - [43] G. O. Roberts, J. S. Rosenthal, Coupling and ergodicity of adaptive Markov chain Monte Carlo algorithms, *Journal of Applied Probability* 44 (2007) 458–475.
 - [44] E. Saksman, M. Vihola, et al., On the ergodicity of the adaptive metropolis algorithm on unbounded domains, *The Annals of applied probability* 20 (2010) 2178–2203.
 - [45] G. O. Roberts, S. Stramer, Langevin diffusions and Metropolis-Hastings algorithms, *Methodology and Computing in Applied Probability* 4 (2002) 337–357.
 - [46] N. Petra, J. Martin, G. Stadler, O. Ghattas, A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet flow inverse problems, *SIAM Journal on Scientific Computing* (2014). To appear.
 - [47] T. M. Russi, Uncertainty quantification with experimental data and complex system models, Ph.D. thesis, University of California, Berkeley, 2010.
 - [48] P. G. Constantine, E. Dow, Q. Wang, Active subspace methods in theory and practice: Applications to kriging surfaces, *SIAM Journal on Scientific Computing* 36 (2014) A1500–A1524.
 - [49] A. Beskos, A stable manifold MCMC method for high dimensions, *Statistics & Probability Letters* 90 (2014) 46–52.
 - [50] G. D. Prato, J. Zabczyk, *Stochastic equations in infinite dimensions*, Cambridge University Press, 1992.

Appendix A. Feldman-Hajek Theorem

By the Feldman-Hajek Theorem (see Theorem 2.23 of [50]), a pair of Gaussian measures $\mathcal{N}(m_1, \Gamma_1)$ and $\mathcal{N}(m_2, \Gamma_2)$ are equivalent measures if and only if the following conditions are satisfied:

1. The two measures have the same Cameron-Martin space, i.e., $\text{Im}(\Gamma_1^{1/2}) = \text{Im}(\Gamma_2^{1/2})$.
2. The difference between the mean functions is in the Cameron-Martin space, i.e., $m_1 - m_2 \in \text{Im}(\Gamma_1^{1/2})$.
3. The operator $T = \left(\Gamma_1^{-1/2}\Gamma_2^{1/2}\right)\left(\Gamma_1^{-1/2}\Gamma_2^{1/2}\right)^* - I$ is Hilbert-Schmidt.

Appendix B. Proof of Theorem 3.1

Let $q(u, du')$ denote the proposal distribution defined by (13), which has the form

$$\begin{aligned} u' &= \Gamma_{\text{pr}}^{1/2} \mathbf{A} \Gamma_{\text{pr}}^{-1/2} (u - u_{\text{ref}}) + u_{\text{ref}} + \Gamma_{\text{pr}}^{1/2} \mathbf{B} \xi - \Gamma_{\text{pr}}^{1/2} \mathbf{G} \Gamma_{\text{pr}}^{1/2} \nabla_u \eta(u; y) \\ &= \Gamma_{\text{pr}}^{1/2} \mathbf{A} \Gamma_{\text{pr}}^{-1/2} (u - u_{\text{ref}}) + u_{\text{ref}} + \Gamma_{\text{pr}}^{1/2} \Psi \left(\mathbf{D}_B \xi' - \mathbf{D}_G \Psi^* \Gamma_{\text{pr}}^{1/2} \nabla_u \eta(u; y) \right) \end{aligned} \quad (\text{B.1})$$

where $\xi \sim \mathcal{N}(0, \mathbf{I})$, $\xi' = \Psi^* \xi$, and $u_{\text{ref}} = m_0 + m_{\text{ref}}$ with $m_{\text{ref}} \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$. We note that $\xi' \sim \mathcal{N}(0, \mathbf{I})$ because Ψ is an unitary operator.

The term $\mathbf{D}_B \xi' - \mathbf{D}_G \Psi^* \Gamma_{\text{pr}}^{1/2} \nabla_u \eta(u; y)$ can be decomposed as

$$\begin{cases} b_i \xi'_i - g_i u_i^* \left(\Gamma_{\text{pr}}^{1/2} \nabla_u \eta(u; y) \right) & \text{if } b_i \neq 0 \\ 0 & \text{if } b_i = 0 \end{cases}$$

Therefore the Gaussian laws for $\Gamma_{\text{pr}}^{1/2} \Psi \mathbf{D}_B \xi'$ and $\Gamma_{\text{pr}}^{1/2} \Psi \left(\mathbf{D}_B \xi' - \mathbf{D}_G \Psi^* \Gamma_{\text{pr}}^{1/2} \nabla_u \eta(u; y) \right)$ are equivalent since, by assumption, $\Gamma_{\text{pr}} \nabla_u \eta(u; y) \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$.

Following the argument in the proof of Theorem 4.1 of [10], we can define a simplified proposal distribution $\tilde{q}(u, du')$ by

$$\begin{aligned} u' &= \Gamma_{\text{pr}}^{1/2} \mathbf{A} \Gamma_{\text{pr}}^{-1/2} (u - u_{\text{ref}}) + u_{\text{ref}} + \Gamma_{\text{pr}}^{1/2} \Psi \mathbf{D}_B \xi' \\ &= \Gamma_{\text{pr}}^{1/2} \mathbf{A} \Gamma_{\text{pr}}^{-1/2} (u - u_{\text{ref}}) + u_{\text{ref}} + \Gamma_{\text{pr}}^{1/2} \mathbf{B} \xi, \end{aligned} \quad (\text{B.2})$$

where $\tilde{q}(u, du')$ is equivalent to $q(u, du')$. Now we consider the pair of measures

$$\begin{aligned} \tilde{\nu}(du, du') &= \tilde{q}(u, du') \mu_0(du), \\ \tilde{\nu}^\perp(du, du') &= \tilde{q}(u', du) \mu_0(du'), \end{aligned} \quad (\text{B.3})$$

which are joint Gaussian measures on (u, u') . Under the assumption that μ_y and μ_0 are equivalent, $\tilde{\nu}$ and $\tilde{\nu}^\perp$ are equivalent to the measures $\nu(du, du') = q(u, du') \mu_y(du)$ and $\nu^\perp(du, du') = q(u', du) \mu_y(du')$, respectively. To provide a well-defined MH algorithm for the infinite dimensional posterior measure μ_y , one requires that $\tilde{\nu}$ and $\tilde{\nu}^\perp$ are equivalent measures, i.e., that

$$\frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(u, u')$$

is positive and bounded $\tilde{\nu}$ -almost surely. In the rest of this proof, we will establish the equivalence of $\tilde{\nu}$ and $\tilde{\nu}^\perp$, given the conditions provided above.

The pair of Gaussian measures $\tilde{\nu}$ and $\tilde{\nu}^\perp$ have the form

$$\tilde{\nu}(du, du') = \mathcal{N}(m_1, \mathbf{C}_1) \text{ and } \tilde{\nu}^\perp(du, du') = \mathcal{N}(m_2, \mathbf{C}_2). \quad (\text{B.4})$$

Recall that $u_{\text{ref}} = m_0 + m_{\text{ref}}$, where $m_{\text{ref}} \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$. The mean functions of $\tilde{\nu}$ and $\tilde{\nu}^\perp$ have the form

$$m_1 = \begin{pmatrix} m_0 \\ m_0 + m_{\text{ref}} - \Gamma_{\text{pr}}^{1/2} \mathbf{A} \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}} \end{pmatrix} \text{ and } m_2 = \begin{pmatrix} m_0 + m_{\text{ref}} - \Gamma_{\text{pr}}^{1/2} \mathbf{A} \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}} \\ m_0 \end{pmatrix}, \quad (\text{B.5})$$

respectively. The covariance operators of $\tilde{\nu}$ and $\tilde{\nu}^\perp$ are

$$C_1 = LQ_1L^* \text{ and } C_2 = LQ_2L^*, \quad (\text{B.6})$$

respectively, where

$$Q_1 = \begin{pmatrix} I & D_A \\ D_A & D_A^2 + D_B^2 \end{pmatrix}, \quad Q_2 = \begin{pmatrix} D_A^2 + D_B^2 & D_A \\ D_A & I \end{pmatrix}, \text{ and } L = \begin{pmatrix} \Gamma_{\text{pr}}^{1/2}\Psi & 0 \\ 0 & \Gamma_{\text{pr}}^{1/2}\Psi \end{pmatrix}.$$

We now apply the Feldman-Hajek (FH) Theorem (see [Appendix A](#)) to show that $\tilde{\nu}$ and $\tilde{\nu}^\perp$ are equivalent measures, given that the conditions in Theorem 3.1 hold. First we consider the mean (condition (2) of the FH theorem). The difference between the mean functions m_1 and m_2 is

$$m_1 - m_2 = \begin{pmatrix} \Gamma_{\text{pr}}^{1/2} A \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}} - m_{\text{ref}} \\ m_{\text{ref}} - \Gamma_{\text{pr}}^{1/2} A \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}} \end{pmatrix},$$

by (B.5). Since $m_{\text{ref}} \in \text{Im}(\Gamma_{\text{pr}}^{1/2})$ and A is a bounded operator, we have $m_1 - m_2 \in \text{Im}(C_1^{1/2})$.

To show the covariance conditions of the FH theorem, first we decompose the space into $\mathcal{H}' \oplus \mathcal{H}_0$ where $\mathcal{H}' = \text{span}\{\psi_i\}_{b_i=0}$ and $\mathcal{H}_0 = \text{span}\{\psi_i\}_{b_i \neq 0}$. Furthermore, note that due to the block structure of diagonal operators, we decompose the product space into $\mathcal{H}'^2 \oplus \mathcal{H}_0^2$, where $C_i = C'_i \oplus C_{0,i}$ and C'_i and $C_{0,i}$ act only on $\mathcal{H}'^2$ and $\mathcal{H}_0^2$ respectively. Note that $C'_1 = C'_2$ by Assumption 2.1 (2), so conditions (1) and (3) of the FH theorem are automatically satisfied. We will consider the measures restricted to the $\mathcal{H}_0^2$ space. Indeed, without loss of generality, we will assume $|b_i| \geq c > 0$ for all $i \in \mathbb{N}$ for the remainder of the proof, so that $\mathcal{H}^2 = \mathcal{H}_0^2$.

For condition (1) of the FH theorem, note that given conditions (1)–(3) of Theorem 3.1, there exist constants $c_1, c_2 > 0$ such that

$$c_1 \|x\| < \|Q_1 x\| < c_2 \|x\| \text{ and } c_1 \|x\| < \|Q_2 x\| < c_2 \|x\|, \quad \forall x \in \mathcal{H}^2.$$

Thus the operators $Q_1^{1/2}$ and $Q_2^{1/2}$ have the same image, i.e., $\text{Im}(Q_1^{1/2}) = \text{Im}(Q_2^{1/2})$. Since $C_i^{1/2} = LQ_i^{1/2}$ for $i = 1, 2$, we have immediately $\text{Im}(C_1^{1/2}) = \text{Im}(C_2^{1/2})$.

For condition (3) of the FH theorem, note that the operators Q_1 and Q_2 can be decomposed as $Q_1 = R_1 R_1^*$ and $Q_2 = R_2 R_2^*$, where

$$R_1 = \begin{pmatrix} I & 0 \\ D_A & D_B \end{pmatrix} \text{ and } R_2 = \begin{pmatrix} D_B & D_A \\ 0 & I \end{pmatrix}. \quad (\text{B.7})$$

Thus C_1 and C_2 can be decomposed as

$$C_1 = (LR_1)(LR_1)^* \text{ and } C_2 = (LR_2)(LR_2)^*,$$

respectively. Recollecting terms we have

$$T = \begin{pmatrix} D_A^2 + D_B^2 - I & -D_A D_B^{-1} (D_A^2 + D_B^2 - I) \\ -D_A D_B^{-1} (D_A^2 + D_B^2 - I) & D_A^2 + D_B^{-2} (I - D_A^2)^2 - I \end{pmatrix}. \quad (\text{B.8})$$

So

$$\text{trace}(T^* T) = \sum_{i=1}^{\infty} \frac{(a_i^2 + b_i^2 - 1)^2 \left((a_i^2 + b_i^2 - 1)^2 + 2b_i^2 \right)}{b_i^4} < K \sum_{i=1}^{\infty} (a_i^2 + b_i^2 - 1)^2,$$

for some $K < \infty$, and hence the operator T is Hilbert-Schmidt given that condition (4) of Theorem 3.1 holds.

Therefore for operators A , B , and G that satisfy the conditions of Theorem 3.1, the proposal (13) yields a well-defined MH algorithm for the infinite dimensional posterior measure μ_y that is equivalent to the prior measure μ_0 .

Appendix C. Proof of Corollary 3.5

Condition (2) of Theorem 3.1 defines a conditional update, in which the components corresponding to $b_i = 0$ are fixed. These fixed components do not affect the acceptance probability, so we will assume $|b_i| \geq c > 0$ for all $i \in \mathbb{N}$ in this proof.

As in Appendix B, we first define the pair of proposals

$$\begin{aligned} q(v, dv') &: v' = Av + B\xi - G\nabla_v \eta(v; y), \\ \tilde{q}(v, dv') &: v' = Av + B\xi. \end{aligned}$$

Then we define the pair of measures $\nu(dv, dv')$ and $\tilde{\nu}(dv, dv')$ associated with this pair of proposals as $\nu(dv, dv') = q(v, dv')\mu_y(dv)$, and $\tilde{\nu}(dv, dv') = \tilde{q}(v, dv')\mu_0(dv)$, respectively. Here $\mu_0 = \mathcal{N}(0, I)$. We can define $\nu^\perp(dv, dv')$ and $\tilde{\nu}^\perp(dv, dv')$ in a similar way. The acceptance probability has the form

$$\alpha(v, v') = \min \left\{ 1, \frac{d\nu^\perp}{d\tilde{\nu}^\perp}(v, v') \right\}.$$

Letting

$$\beta(v, v') = \log \left(\frac{d\nu}{d\tilde{\nu}}(v, v') \right) = \log \left(\frac{dq(v, \cdot)}{d\tilde{q}(v, \cdot)}(v') \right) \log \left(\frac{d\mu_y}{d\mu_0}(v) \right),$$

we have

$$\frac{d\nu^\perp}{d\tilde{\nu}^\perp}(v, v') = \frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(v, v') \exp(\beta(v', v) - \beta(v, v')). \quad (\text{C.1})$$

Recall equation (16),

$$\frac{d\mu_y}{d\mu_0}(v) = \exp \left(-\eta(v; y) - \frac{1}{2}\|v_{\text{ref}}\|^2 - \langle v_{\text{ref}}, v \rangle \right),$$

where $v_{\text{ref}} = \Gamma_{\text{pr}}^{-1/2} m_{\text{ref}}$. The term $\frac{1}{2}\|v_{\text{ref}}\|^2$ is cancelled out in $\beta(v', v)$ and $\beta(v, v')$, and hence we have

$$\beta(v, v') = -\eta(v; y) - \langle v_{\text{ref}}, v \rangle + \log \left(\frac{dq(v, \cdot)}{d\tilde{q}(v, \cdot)}(v') \right). \quad (\text{C.2})$$

The first two terms in (C.2) are μ_0 -a.s. finite by the discussion preceding (16).

To simplify the derivation of the right-hand side of (C.1), we apply the transformation (15). Then the pair of proposals $q(v, dv')$ and $\tilde{q}(v, dv')$ defined on the transformed parameters (w, w') have the form

$$\begin{aligned} q(w, dw') &: w' = D_A w + D_B \xi - D_G \nabla_w \eta(w; y), \\ \tilde{q}(w, dw') &: w' = D_A w + D_B \xi. \end{aligned}$$

The pair of measures $\tilde{\nu}^\perp(v, dv')$ and $\tilde{\nu}(v, dv')$ transformed onto (w, w') have the form

$$\begin{aligned}\tilde{\nu}^\perp(dw, dw') &= \tilde{q}(w, dw')\mu_0^w(dw), \\ \tilde{\nu}(dw, dw') &= \tilde{q}(w', dw)\mu_0^w(dw').\end{aligned}$$

where $\mu_0^w = \mathcal{N}(0, \mathbf{I})$ as the transformation (15) is unitary. We want to derive the Radon-Nikodym derivatives

$$\frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(w, w') = \frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(v, v'), \text{ and } \frac{dq(w, \cdot)}{d\tilde{q}(w, \cdot)}(w') = \frac{dq(v, \cdot)}{d\tilde{q}(v, \cdot)}(v').$$

Following the same derivation as in [Appendix B](#), we can show that $\tilde{\nu}^\perp(dw, dw')$ and $\tilde{\nu}(dw, dw')$ follow the Gaussian measures

$$\tilde{\nu}(dw, dw') = \mathcal{N}(0, \mathbf{Q}_1), \text{ and } \tilde{\nu}^\perp(dw, dw') = \mathcal{N}(0, \mathbf{Q}_2),$$

where $\mathbf{Q}_1$ and $\mathbf{Q}_2$ are defined as in (B.6). Thus we have

$$\begin{aligned}\log\left(\frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(w, w')\right) &= -\frac{1}{2}\langle w', \mathbf{D}_B^{-2}(\mathbf{D}_A^2 + \mathbf{D}_B^2 - \mathbf{I}) w' \rangle \\ &\quad + \frac{1}{2}\langle w, \mathbf{D}_B^{-2}(\mathbf{D}_A^2 + \mathbf{D}_B^2 - \mathbf{I}) w \rangle.\end{aligned}$$

Both terms in the above equation are μ_0 -a.s. finite by condition (4) of Theorem 3.1.

Following Theorem 2.21 of [50], we have

$$\begin{aligned}\log\left(\frac{dq(w, \cdot)}{d\tilde{q}(w, \cdot)}(w')\right) &= -\langle \mathbf{D}_B^{-1}\mathbf{D}_G\nabla_w\eta(w; y), \mathbf{D}_B^{-1}(w' - \mathbf{D}_A w) \rangle \\ &\quad - \frac{1}{2}\|\mathbf{D}_B^{-1}\mathbf{D}_G\nabla_w\eta(w; y)\|^2.\end{aligned}$$

Following the same argument as in [Appendix B](#), we can show that the Gaussian laws of $q(w, dw')$ and $\tilde{q}(w, dw')$ are equivalent, and hence $\log(\frac{dq(w, \cdot)}{d\tilde{q}(w, \cdot)}(w'))$ is also guaranteed to be μ_0 -a.s. finite. This can be seen from the equation above, by the conditions of Theorem 3.1 and Assumption 2.1(3), since the inner product of an element of $\mathcal{H}$ with white noise is finite almost surely.

Transforming $\log\left(\frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(w, w')\right)$ back to (v, v') , we have

$$\begin{aligned}\log\left(\frac{d\tilde{\nu}^\perp}{d\tilde{\nu}}(v, v')\right) &= -\frac{1}{2}\langle v', \mathbf{B}^{-2}(\mathbf{A}^2 + \mathbf{B}^2 - \mathbf{I}) v' \rangle \\ &\quad + \frac{1}{2}\langle v, \mathbf{B}^{-2}(\mathbf{A}^2 + \mathbf{B}^2 - \mathbf{I}) v \rangle.\end{aligned}$$

Transforming $\log\left(\frac{dq(w, \cdot)}{d\tilde{q}(w, \cdot)}(w')\right)$ back to (v, v') and substituting into (C.2), we have

$$\beta(v, v') = -\eta(v; y) - \langle v_{\text{ref}}, v \rangle - \langle \mathbf{B}^{-1}\mathbf{G}\nabla_v\eta(v; y), \mathbf{B}^{-1}(v' - \mathbf{A}v) \rangle - \frac{1}{2}\|\mathbf{B}^{-1}\mathbf{G}\nabla_v\eta(v; y)\|^2.$$

By collecting terms we can define

$$\begin{aligned}\rho(v, v') &= -\eta(v; y) - \langle v_{\text{ref}}, v \rangle - \frac{1}{2}\langle v, \mathbf{B}^{-2}(\mathbf{A}^2 + \mathbf{B}^2 - \mathbf{I}) v \rangle \\ &\quad - \langle \mathbf{B}^{-1}\mathbf{G}\nabla_v\eta(v; y), \mathbf{B}^{-1}(v' - \mathbf{A}v) \rangle - \frac{1}{2}\|\mathbf{B}^{-1}\mathbf{G}\nabla_v\eta(v; y)\|^2,\end{aligned}$$

and the acceptance probability has the form

$$\alpha(v, v') = \min \{1, \exp (\rho(v', v) - \rho(v, v'))\} .$$