

# A transport-based multifidelity preconditioner for Markov chain Monte Carlo\*

Benjamin Peherstorfer<sup>†</sup>      Youssef Marzouk<sup>‡</sup>

May 2019

## Abstract

Markov chain Monte Carlo (MCMC) sampling of posterior distributions arising in Bayesian inverse problems is challenging when evaluations of the forward model are computationally expensive. Replacing the forward model with a low-cost, low-fidelity model often significantly reduces computational cost; however, employing a low-fidelity model alone means that the stationary distribution of the MCMC chain is the posterior distribution corresponding to the low-fidelity model, rather than the original posterior distribution corresponding to the high-fidelity model. We propose a multifidelity approach that combines, rather than replaces, the high-fidelity model with a low-fidelity model. First, the low-fidelity model is used to construct a transport map that deterministically couples a reference Gaussian distribution with an approximation of the low-fidelity posterior. Then, the high-fidelity posterior distribution is explored using a non-Gaussian proposal distribution derived from the transport map. This multifidelity “preconditioned” MCMC approach seeks efficient sampling via a proposal that is explicitly tailored to the posterior at hand and that is constructed efficiently with the low-fidelity model. By relying on the low-fidelity model only to construct the proposal distribution, our approach guarantees that the stationary distribution of the MCMC chain is the high-fidelity posterior. In our numerical examples, our multifidelity approach achieves significant speedups compared to single-fidelity MCMC sampling methods.

**Keywords:** Bayesian inverse problems; transport maps; multifidelity; model reduction; Markov chain Monte Carlo

---

\*The authors acknowledge support of the AFOSR MURI on multi-information sources of multi-physics systems under Award Number FA9550-15-1-0038.

<sup>†</sup>Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (pehersto@cims.nyu.edu)

<sup>‡</sup>Massachusetts Institute of Technology, Cambridge, MA 02139 (ymarz@mit.edu)

# 1 Introduction

Bayesian inference provides a framework to quantify uncertainties in the solutions of inverse problems [31, 62, 63]. The Bayesian approach to inverse problems combines observed data, a forward model that maps parameters to observations, a prior distribution on the parameters, and a statistical model for the mismatch between model predictions and observations to define the *posterior distribution* via Bayes’ theorem. The posterior distribution characterizes the parameter values and their uncertainties, given these ingredients. Practically “solving” a Bayesian inverse problem, however, entails exploring the posterior distribution, e.g., computing posterior expectations. A flexible and widely used approach for exploring posterior distributions is to draw samples with Markov chain Monte Carlo (MCMC) methods [64, 27]. Using these methods, the forward model typically must be evaluated multiple times at different parameter values for each sample that is drawn, such that MCMC sampling quickly becomes computationally infeasible if each forward model solve is expensive.

In this paper, we propose a multifidelity preconditioner to increase the efficiency of MCMC sampling. Our multifidelity approach exploits low-cost, low-fidelity models to construct a proposal distribution that approximates the posterior distribution at hand, then uses this proposal distribution to perform MCMC sampling of the original (high-fidelity) posterior distribution. The proposal distribution is derived from a *transport map* that transforms the potentially complex posterior distribution into another distribution from which samples can be drawn more easily; in particular, we seek a map that transforms the posterior into a more Gaussian and more isotropic distribution. See [41, 37, 45] for an introduction to transport maps in the context of Bayesian inverse problems. The key idea in the present work is to construct an invertible map using low-fidelity models, but to apply it to the high-fidelity posterior. The map then acts as a *preconditioner* for the high-fidelity posterior, preserving information (due to its invertibility) while enabling MCMC sampling to proceed more efficiently. An alternative but equivalent perspective is that pushing a simple proposal through the inverse of this transport map—for instance, using a Metropolis independence sampler with a standard Gaussian proposal—yields potentially non-Gaussian and tailored proposals that can be efficient for the high-fidelity posterior. The end result is the same: more efficient MCMC sampling, allowing the number of high-fidelity forward model evaluations to be reduced. Since the low-fidelity model is used only for preconditioning (or equivalently, constructing the proposal), the stationary distribution of the chain obtained with our multifidelity approach is guaranteed to be the posterior distribution corresponding to the high-fidelity model.

There is a long tradition of exploiting low-fidelity models to speed up MCMC sampling for Bayesian inverse problems. There is work [69, 2, 39, 38, 34, 14] that *replaces* the high-fidelity model with a low-fidelity model (see also the survey [26]); because of this replacement, however, samples are drawn from a different distribution—the posterior induced by the low-fidelity model. Under certain assumptions, it has been shown that the posterior distribution induced by the

low-fidelity model converges to the original/high-fidelity posterior distribution as the low-fidelity approximation is refined [40, 19]. Other approaches *adapt* low-fidelity models [20, 33] over a finite interval of posterior exploration, or quantify the error introduced by sampling the low-fidelity posterior distribution [22, 36]. Yet another family of approaches incrementally and infinitely refines approximations of the forward model on-the-fly during MCMC sampling [17, 16]; under appropriate conditions, these schemes guarantee that the MCMC chain asymptotically samples the high-fidelity posterior distribution.

Instead of replacing the high-fidelity model with low-fidelity models, multifidelity methods combine high- and low-fidelity models. The aim is to leverage low-fidelity models for speeding up computations while allowing occasional recourse to the high-fidelity model to establish accuracy guarantees [50]. A variety of multifidelity methods have been developed for uncertainty propagation [8, 7, 42, 24, 12, 47, 49, 52, 48]; see the survey [50]. For the solution of Bayesian inverse problems, there are multi-stage MCMC methods that aim to reduce the number of high-fidelity model evaluations by first screening proposed moves with low-fidelity models [15, 23]. Another line of work builds on hierarchies of low-fidelity models, typically derived from different discretizations of partial differential equations (PDEs) underlying the high-fidelity model, to reduce sampling costs [4, 21, 32].

An alternative to multi-stage and hierarchical methods for increasing the efficiency of MCMC sampling is to use notions of transport to construct more effective proposal distributions. Effective proposals in MCMC should reflect the local or global geometry of the target distribution [27, 9]. In keeping with this idea, Parno et al. [45] use transport maps to precondition MCMC sampling. As described earlier, preconditioning involves constructing transport maps that “Gaussianize” the target distribution, such that it can be sampled more effectively by standard MCMC algorithms; these maps thus encode the geometry of the target. Parno et al. [45] build and refine such transport maps in an online fashion during MCMC sampling, as more and more samples are obtained. This approach can be seen as a form of *adaptive MCMC* [30, 54], wherein a (non-Gaussian) proposal distribution is adapted as the MCMC sampling proceeds. The transport maps in [45] are constructed via the solution of a convex and separable optimization problem, which is simple and fast to obtain numerically. However, this sampling approach faces some of the usual pitfalls of adaptive MCMC. One of these issues is that a certain amount of initial mixing is necessary for the adaptation to be effective, because the online-adapted transport map depends on past samples.

The transport approach of Moselhy et al. [41] instead follows an offline/online decomposition of the computation. In an offline phase, a reference distribution is selected from which independent samples can be drawn cheaply and then a transport map is constructed that pushes forward the reference distribution to the posterior distribution. Then, in the online phase, the transport map is used to transform samples from the reference distribution into samples from the posterior distribution. If a large number of samples are transformed in the online phase, then the one-time high cost of constructing the transport map

offline is compensated. The optimization problem for constructing the transport map can employ derivative information from the posterior density, and thus yields accurate maps even if the posterior is concentrated. Furthermore, no samples of the posterior distribution are necessary to solve the optimization problem, in contrast with adaptive MCMC techniques such as [45]. However, evaluations of the objective of the optimization problem entail evaluations of the unnormalized posterior density and thus of the forward model; moreover, this problem is in general not separable across dimensions and not convex (unless the posterior density is log-concave). As a result, map construction in [41] is typically far more computationally demanding than map construction in [45]. Furthermore, since the construction of the transport map involves numerical approximations, the pushforward of the reference distribution by the map in general only *approximates* the posterior distribution.

We propose a multifidelity approach that combines several of the advantages of the approaches introduced by Parno [45] and Moselhy [41]. We construct transport maps building on the offline/online approach of [41] and mitigate the high computational costs of solving the optimization problem by relying on low-cost, low-fidelity models. Our approach can exploit a wide range of low-fidelity models, including projection-based reduced models [58, 56, 28, 3, 13], data-fit interpolation and regression models [25], machine-learning-based models [18, 70, 57, 65], and simplified-physics models [1, 43]. Then, in the online phase, the transport map is used to precondition MCMC sampling of the high-fidelity posterior, as in [45]. The corresponding MCMC scheme is ergodic for the high-fidelity posterior. In other words, the Metropolis step corrects errors that otherwise would be introduced by relying on the transport map alone to push forward the reference to the posterior. Thus, we obtain a multifidelity approach that uses low-fidelity models to speed up computations while making occasional recourse to the high-fidelity model to establish convergence, in the sense that the stationary distribution of our MCMC chain is the high-fidelity posterior distribution. In contrast to [45], our MCMC algorithm is not adaptive, because the map is built once offline and then stays fixed during MCMC sampling. In particular, no samples of the high-fidelity posterior distribution are needed to construct the transport maps.

Section 2 describes the problem setup and briefly reviews transport maps in the context of Bayesian inverse problems. In Section 3, we introduce our multifidelity preconditioner and discuss it in the context of sampling with the Metropolis-Hastings algorithm. Section 4 demonstrates our multifidelity approach on two examples where we achieve significant speedups compared to using the high-fidelity model alone. Conclusions are given in Section 5.

## 2 Preliminaries

Section 2.1 and Section 2.2 define our Bayesian inverse problem setting and describe the Metropolis-Hastings algorithm for MCMC sampling. We refer to, e.g., [31, 62, 63], for details on Bayesian approaches to inverse problems. Section 2.3

discusses transport maps for coupling probability distributions in the context of Bayesian inverse problems. The problem description is given in Section 2.4.

## 2.1 Bayesian inverse problems

Consider the high-fidelity forward model (parameter-to-observable map)  $G : \mathcal{P} \rightarrow \mathcal{Y}$ , with parameter  $\boldsymbol{\theta} \in \mathcal{P}$ , where  $\mathcal{P} \subseteq \mathbb{R}^d$ , and observable  $\mathbf{y} \in \mathcal{Y}$ , where  $\mathcal{Y} \subseteq \mathbb{R}^{d'}$ . Thus, the parameter  $\boldsymbol{\theta}$  and the observable  $\mathbf{y}$  are  $d$ -dimensional and  $d'$ -dimensional vectors, respectively. Consider now observed data

$$\mathbf{y} = G(\boldsymbol{\theta}^*) + \boldsymbol{\epsilon},$$

corresponding to some parameter value  $\boldsymbol{\theta}^* \in \mathcal{P}$ . The noise  $\boldsymbol{\epsilon}$  is assumed to be a realization of a zero-mean Gaussian random variable with covariance matrix  $\boldsymbol{\Sigma}_\epsilon \in \mathbb{R}^{d' \times d'}$ . Define the data-misfit function

$$\Phi_{\mathbf{y}}(\boldsymbol{\theta}) = \frac{1}{2} \left\| \boldsymbol{\Sigma}_\epsilon^{-\frac{1}{2}} (G(\boldsymbol{\theta}) - \mathbf{y}) \right\|_2^2,$$

with the Euclidean norm  $\|\cdot\|_2$ . The likelihood function  $L_{\mathbf{y}} : \mathcal{P} \rightarrow \mathbb{R}$  is

$$L_{\mathbf{y}}(\boldsymbol{\theta}) = \exp(-\Phi_{\mathbf{y}}(\boldsymbol{\theta})).$$

Combining the prior distribution with density  $\pi_0$  and the likelihood  $L$  via Bayes' theorem gives the posterior density up to a normalizing constant

$$\pi(\boldsymbol{\theta}) \propto L_{\mathbf{y}}(\boldsymbol{\theta})\pi_0(\boldsymbol{\theta}).$$

Above and for the remainder of this paper, we assume that the prior and posterior measures are absolutely continuous with respect to the Lebesgue measure.

## 2.2 The Metropolis-Hastings algorithm

MCMC methods are widely used to sample posterior distributions that arise in Bayesian inverse problems. The Metropolis-Hastings algorithm defines a wide class of MCMC methods, on which we will build in the following. Algorithm 1 describes the Metropolis-Hastings approach. In our Bayesian inverse problem setting, inputs are the likelihood  $L$ , the prior density  $\pi_0$ , a proposal density  $q$ , and the number of iterations  $M \in \mathbb{N}$ . In each iteration, a proposal sample  $\boldsymbol{\theta}'$  is drawn from the proposal distribution, which may depend on the previous sample  $\boldsymbol{\theta}_{i-1}$ . Then, the acceptance probability  $\alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}')$  is computed, which requires evaluating the likelihood  $L$  and prior  $\pi_0$  at the candidate sample  $\boldsymbol{\theta}'$ . The proposal sample  $\boldsymbol{\theta}'$  is accepted with probability  $\alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}')$  and rejected with probability  $1 - \alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}')$ . This process is repeated for  $M$  iterations and the samples  $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$  are returned. The corresponding Markov chain is, by construction, reversible for  $\pi$  and thus has  $\pi$  as a stationary distribution. With some relatively simple additional conditions on  $\pi$  and  $q$ , one can show that the chain converges to  $\pi$ , from any starting point; see [53] for a full discussion.

---

**Algorithm 1** Metropolis-Hastings

---

- 1: **procedure** METROPOLISHASTINGS( $L, \pi_0, q, M$ )
- 2:   Choose a starting point  $\boldsymbol{\theta}_0$
- 3:   **for**  $i = 1, \dots, M$  **do**
- 4:     Draw candidate  $\boldsymbol{\theta}'$  from proposal  $q(\cdot|\boldsymbol{\theta}_{i-1})$
- 5:     Compute acceptance probability

$$\alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}') = \min \left\{ 1, \frac{q(\boldsymbol{\theta}_{i-1}|\boldsymbol{\theta}')L_{\mathbf{y}}(\boldsymbol{\theta}')\pi_0(\boldsymbol{\theta}')}{q(\boldsymbol{\theta}'|\boldsymbol{\theta}_{i-1})L_{\mathbf{y}}(\boldsymbol{\theta}_{i-1})\pi_0(\boldsymbol{\theta}_{i-1})} \right\}$$

- 6:     Set the sample  $\boldsymbol{\theta}_i$  to

$$\boldsymbol{\theta}_i = \begin{cases} \boldsymbol{\theta}', & \text{with probability } \alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}'), \\ \boldsymbol{\theta}_{i-1}, & \text{with probability } 1 - \alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}') \end{cases}$$

- 7:   **end for**
  - 8:   **return**  $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$
  - 9: **end procedure**
- 

In general, the samples produced by MCMC are correlated; this correlation inflates the variance of any expectations estimated with MCMC samples, relative to an expectation estimated with uncorrelated Monte Carlo sample sets of the same size. The efficiency of an MCMC sampler can thus be measured with the *effective sample size* (ESS) of any sample set it produces, which is inversely proportional to the *integrated autocorrelation time* of the chain [35]. To define the ESS, consider a function  $f : \mathcal{P} \rightarrow \mathbb{R}$  that is measurable with respect to the Lebesgue measure, and let us assume we are interested in estimating the expected value

$$\mathbb{E}[f] = \int_{\mathcal{P}} f(\boldsymbol{\theta})\pi(\boldsymbol{\theta}) \, d\boldsymbol{\theta},$$

with respect to the posterior distribution  $\pi$ . Consider now the Monte Carlo estimator of  $\mathbb{E}[f]$  that uses  $n \in \mathbb{N}$  samples  $\{\boldsymbol{\theta}_i\}_{i=1}^n$ ,

$$E_n f = \frac{1}{n} \sum_{i=1}^n f(\boldsymbol{\theta}_i).$$

The ESS of  $\{f(\boldsymbol{\theta}_i)\}_{i=1}^n$  is  $n^* \in \mathbb{R}$  such that

$$\text{Var}[E_n f] = \frac{\text{Var}[f]}{n^*},$$

where  $\text{Var}[E_n f]$  is the variance of the estimator  $E_n f$  and  $\text{Var}[f]$  is the variance of  $f(\boldsymbol{\theta})$  for  $\boldsymbol{\theta} \sim \pi$ . In other words,  $n^* \leq n$  is the number of independent Monte Carlo samples from  $\pi$  that would be required to obtain an estimator with the same variance as  $E_n f$ . It can be shown that  $n^* = n/\tau$ , where  $\tau$  is the integrated autocorrelation time associated with the chain  $\{f(\boldsymbol{\theta}_t)\}_t$ . Better MCMC mixing corresponds to smaller  $\tau$  and larger ESS.

## 2.3 Transport maps

The use of transport maps in the context of Bayesian inference was introduced in [41]. In particular, [41] proposed a variational Bayesian approach involving transport. Rather than using importance sampling or MCMC to characterize the posterior distribution, this approach seeks a transport map that pushes forward a tractable “reference” distribution to the posterior, such that samples drawn from the reference and acted on by the map are distributed according to the posterior. Below we follow [37] to introduce the notion of transport maps.

### 2.3.1 Definition of transport maps

Let  $\mu_\pi$  and  $\mu_\eta$  be two probability measures on  $\mathbb{R}^d$  that are absolutely continuous with respect to the Lebesgue measure. In the following,  $\mu_\pi$  is the target measure, which corresponds to the posterior distribution in our case, and  $\mu_\eta$  is the reference measure, which typically is a Gaussian or another distribution from which we can draw independent samples efficiently. The probability density function corresponding to the target measure  $\mu_\pi$  is the posterior  $\pi$ , and the probability density function corresponding to  $\mu_\eta$  is denoted by  $\eta$ . A transport map is a function  $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$  that pushes forward the reference  $\mu_\eta$  to the target  $\mu_\pi$ , which we write as

$$\mu_\pi = T_\# \mu_\eta, \quad (1)$$

and which means that for any Borel set  $A \subseteq \mathbb{R}^d$ , it holds  $\mu_\pi(A) = \mu_\eta(T^{-1}(A))$ . Existence of such maps is ensured by the absolute continuity of the reference and the target measure. There may be infinitely many transport maps that push forward a given reference to the target of interest. Uniqueness can be enforced by introducing a cost function that is minimized while imposing the constraint (1). This construction leads to the notion of *optimal transport*; see, e.g., [66, 67, 68].

Instead of introducing a cost function to regularize the problem of finding a transport map, we directly impose structure on the map  $T$  as in [41, 45, 37]. In particular, we will seek lower triangular maps that are monotone increasing. The lower triangular structure of  $T$  is as follows

$$T(\vartheta_1, \vartheta_2, \dots, \vartheta_d) = \begin{bmatrix} T_1(\vartheta_1) \\ T_2(\vartheta_1, \vartheta_2) \\ T_3(\vartheta_1, \vartheta_2, \vartheta_3) \\ \vdots \\ T_d(\vartheta_1, \vartheta_2, \dots, \vartheta_d) \end{bmatrix}, \quad (2)$$

where  $\vartheta_i$  denotes the  $i$ th component of  $\boldsymbol{\vartheta} = [\vartheta_1, \dots, \vartheta_d]^T \in \mathbb{R}^d$  and where  $T_i : \mathbb{R}^i \rightarrow \mathbb{R}$  is the  $i$ th component function of the map  $T$ . Monotonicity in this context corresponds to the condition that  $T_i$  is a monotone increasing function of  $\vartheta_i$ , for all  $i = 1 \dots d$ . This condition ensures that  $\nabla T \succeq 0$  and  $\det \nabla T \geq 0$ ; see [37, 61] for more detail. Since we assume that the reference and the target measures are absolutely continuous, existence and uniqueness of such a

lower-triangular transport map is guaranteed; this map is in fact the Knothe–Rosenblatt rearrangement [6, 11, 55].

### 2.3.2 Numerical approximations of transport maps

Following [41], we will obtain numerical approximations of the Knothe–Rosenblatt rearrangement by solving an optimization problem. Let  $\mathcal{T}$  be a finite-dimensional subspace of the space of all smooth lower triangular functions (2) from  $\mathbb{R}^d$  into  $\mathbb{R}^d$ . Then, an approximation  $\tilde{T} \in \mathcal{T}$  of a transport map  $T$  can be obtained via numerical optimization over the coefficients of the representation of  $\tilde{T}$  in a basis of  $\mathcal{T}$ . To set up the optimization problem, consider the *pullback*  $(\tilde{T}^{-1})_{\#}\mu_{\pi} \equiv \tilde{T}^{\#}\mu_{\pi}$  of  $\mu_{\pi}$  through a map  $\tilde{T}$ ; the density of this pullback measure can be written as

$$\tilde{\eta}(\boldsymbol{\vartheta}) = \pi(\tilde{T}(\boldsymbol{\vartheta}))|\det \nabla \tilde{T}(\boldsymbol{\vartheta})|, \quad (3)$$

where  $|\det \nabla \tilde{T}(\boldsymbol{\vartheta})|$  is the absolute value of the determinant of the Jacobian  $\nabla \tilde{T}(\boldsymbol{\vartheta})$  of  $\tilde{T}$  at  $\boldsymbol{\vartheta}$ . Note that the functions in  $\mathcal{T}$  are smooth in the sense that  $\nabla \tilde{T}$  exists and is sufficiently regular; see, e.g., [37, 61]. Let

$$D_{\text{KL}}(\pi_1||\pi_2) = \mathbb{E}_{\pi_1} \left( \log \frac{\pi_1}{\pi_2} \right)$$

denote the Kullback–Leibler (KL) divergence of  $\pi_1$  from  $\pi_2$  (where  $\pi_1$  and  $\pi_2$  in the log term are densities). Then, a solution  $\tilde{T}^* \in \mathcal{T}$  of the optimization problem

$$\begin{aligned} \min_{\tilde{T} \in \mathcal{T}} \quad & D_{\text{KL}}(\eta||\tilde{\eta}), \\ \text{s.t.} \quad & \nabla \tilde{T} \succ 0, \end{aligned} \quad (4)$$

is an approximation of a transport map that pushes forward the reference measure  $\mu_{\eta}$  to the target  $\mu_{\pi}$ . The constraint  $\nabla \tilde{T} \succ 0$  means that the Jacobian of  $\tilde{T}$  is positive definite. If the approximation space  $\mathcal{T}$  is sufficiently rich such that  $D_{\text{KL}}(\eta||\tilde{\eta}) = 0$ , then we have  $\tilde{T}^{*\#}\mu_{\pi} = \mu_{\eta}$  and  $\tilde{T}^*_{\#}\mu_{\eta} = \mu_{\pi}$  [37].

The KL divergence is not symmetric. The direction of the KL divergence here is chosen such that the expected value is taken with respect to the reference  $\mu_{\eta}$ , which is selected so that it can easily be sampled. Furthermore, the objective in (4) can be minimized without knowledge of the normalizing constant of  $\pi$ . To see this, transform the objective  $D_{\text{KL}}(\eta||\tilde{\eta})$  into

$$\begin{aligned} D_{\text{KL}}(\eta||\tilde{\eta}) &= \mathbb{E}_{\eta} \left[ \log \left( \frac{\eta}{\tilde{\eta}} \right) \right] \\ &= \mathbb{E}_{\eta} \left[ \log \eta - \log \pi \circ \tilde{T} - \log |\det \nabla \tilde{T}| \right], \end{aligned} \quad (5)$$

where we used the definition of  $\tilde{\eta}$  in (3). The expectation  $\mathbb{E}_{\eta}[\log \eta]$  is independent of the map  $\tilde{T}$  and therefore can be ignored when minimizing the objective of (4). Similarly, the normalizing constant of  $\pi$  leads to a constant term in (5) that is independent of  $\tilde{T}$  and therefore unnormalized evaluations of  $\pi$  are sufficient to optimize (4).

## 2.4 Problem description

We identify two challenges of directly relying on a solution  $\tilde{T}^*$  of (4) to solve a Bayesian inverse problem, as proposed in [41]. First, the KL divergence objective (5), which contains an expectation with respect to the reference measure  $\mu_\eta$ , is typically estimated with a Monte Carlo method because no closed form expression is available for general  $\pi$ . Thus, the (unnormalized version of the) pullback density (3) must be evaluated at a potentially large number of samples for each optimization iteration. Since each evaluation of the pullback entails an evaluation of the (unnormalized) posterior distribution  $\pi$ , and thus of the forward model  $G$ , the optimization can become computationally intractable if  $G$  is expensive to evaluate. Second, the push forward  $\tilde{T}_\#^* \mu_\eta$  is only an approximation of the posterior  $\mu_\pi$ . In particular, approximation errors may follow from the choice of the finite-dimensional approximation space  $\mathcal{T}$ , the finite number of Monte Carlo samples used to discretize the expectation in the objective, and any other errors in the numerical optimization. While this error can be estimated (see [41, 37]), reducing this error to an arbitrarily small threshold—e.g., by enriching  $\mathcal{T}$ —can be computationally expensive.

## 3 Multifidelity preconditioned Metropolis-Hastings

We propose a multifidelity preconditioned Metropolis-Hastings (MFMH). Similar to the approach of Moselhy [41], our multifidelity approach consists of an offline and an online phase. In the offline phase, a transport map is constructed rapidly by using a low-cost, low-fidelity approximation of the high-fidelity model. Then, in the online phase, a proposal distribution is derived from the transport map and used to sample the high-fidelity posterior via a modified Metropolis-Hastings algorithm, following the work of Parno [45]. If the low-fidelity model is an accurate approximation of the high-fidelity model, and if the transport map captures the essential structure of  $\pi$ , then the ESS of this MFMH will be higher and fewer online evaluations of the high-fidelity model will be required to achieve a given accuracy. In any case, any error associated with the transport map is corrected by Metropolization in the online phase; samples are thus drawn (asymptotically) from the posterior distribution corresponding to the high-fidelity model  $G$ . In other words, the MFMH approach offers the same convergence guarantees as standard MCMC.

### 3.1 Approximation spaces and numerical optimization

This section provides details on the numerical construction of transport maps from the reference measure to the target measure.

#### 3.1.1 Integrated squared parametrization

Each component function  $\tilde{T}_i, i = 1, \dots, d$  of the approximation  $\tilde{T}$  of the map  $T$  defined in (2) is parameterized with the integrated-squared ansatz [5], which

is similar to the integrated-exponential ansatz introduced in [61]. Component function  $\tilde{T}_i$  is parameterized as

$$\begin{aligned} \tilde{T}_i(\vartheta_1, \dots, \vartheta_i; \boldsymbol{\beta}_i) \\ = \tilde{T}_i^{(L)}(\vartheta_1, \dots, \vartheta_{i-1}; \boldsymbol{\beta}_i^{(L)}) + \int_0^{\vartheta_i} \left( \tilde{T}_i^{(R)}(\vartheta_1, \dots, \vartheta_{i-1}, t; \boldsymbol{\beta}_i^{(R)}) \right)^2 dt, \end{aligned} \quad (6)$$

where

$$\boldsymbol{\beta}_i = \begin{bmatrix} \boldsymbol{\beta}_i^{(L)} & \boldsymbol{\beta}_i^{(R)} \end{bmatrix} \quad (7)$$

is a parameter vector and  $\tilde{T}_i^{(L)} : \mathbb{R}^{i-1} \rightarrow \mathbb{R}$  and  $\tilde{T}_i^{(R)} : \mathbb{R}^i \rightarrow \mathbb{R}$  are functions that are parameterized by the parameters  $\boldsymbol{\beta}_i^{(L)}$  and  $\boldsymbol{\beta}_i^{(R)}$ , respectively. The integrated-squared parameterization (6) guarantees that the map  $\tilde{T}$  is monotone and therefore automatically satisfies the constraint  $\nabla \tilde{T} \succ 0$  in (4) [61].

### 3.1.2 Approximation space

Following [45, 37], we represent the functions  $\tilde{T}_i^{(L)}$  and  $\tilde{T}_i^{(R)}$  in each component  $\tilde{T}_i, i = 1, \dots, d$  of the map  $\tilde{T}$  as multivariate polynomials. Let  $\mathbf{j} = [j_1, \dots, j_d]^T \in \mathbb{N}^d$  be a multi-index and let  $\phi_{j_i}$  be a univariate polynomial with degree  $j_i$  for  $i = 1, \dots, d$ . Define the multivariate polynomial function as

$$\phi_{\mathbf{j}}(\boldsymbol{\vartheta}) = \prod_{i=1}^d \phi_{j_i}(\vartheta_i). \quad (8)$$

In the following, the multivariate polynomial functions (8) are simply the monomials. Note that other polynomial families can be used, e.g., Hermite polynomials. (One might also use Hermite functions, as in [61], for better control of tail behavior.) Consider now the sets  $\mathcal{J}_i \subset \mathbb{N}^d$  for  $i = 1, \dots, d$

$$\mathcal{J}_i = \{\mathbf{j} \mid \|\mathbf{j}\|_1 \leq \ell, j_k = 0, \forall k > i\}, \quad (9)$$

which correspond to the total-degree polynomials of maximal degree  $\ell \in \mathbb{N}$ . The constraint  $j_k = 0, \forall k > i$  in the definition of  $\mathcal{J}_i$  in (9) imposes the lower-triangular structure of the map  $\tilde{T}$  as defined in (2). The set  $\mathcal{J}_i$  leads to the definition of the approximation space  $\mathcal{T}_i$

$$\mathcal{T}_i = \text{span}\{\phi_{\mathbf{j}} \mid \mathbf{j} \in \mathcal{J}_i\},$$

of the functions  $\tilde{T}_i^{(R)}$  and  $\tilde{T}_{i+1}^{(L)}$ , respectively, of the  $i$ th and  $i + 1$ st component of  $\tilde{T}$ . Thus,  $\tilde{T}_i^{(L)} \in \mathcal{T}_{i-1}$  and  $\tilde{T}_i^{(R)} \in \mathcal{T}_i$  can be represented as

$$\tilde{T}_i^{(L)}(\vartheta_1, \dots, \vartheta_{i-1}; \boldsymbol{\beta}_i^{(L)}) = \sum_{\mathbf{j} \in \mathcal{J}_{i-1}} \beta_{i,\mathbf{j}}^{(L)} \phi_{\mathbf{j}}(\vartheta_1, \dots, \vartheta_{i-1}), \quad (10)$$

and

$$\tilde{T}_i^{(R)}(\vartheta_1, \dots, \vartheta_{i-1}, t; \beta_i^{(R)}) = \sum_{\mathbf{j} \in \mathcal{J}_i} \beta_{i,\mathbf{j}}^{(R)} \phi_{\mathbf{j}}(\vartheta_1, \dots, \vartheta_{i-1}, t), \quad (11)$$

where  $\beta_i^{(L)}$  and  $\beta_i^{(R)}$  are the vectors of the coefficients  $\beta_{i,\mathbf{j}}^{(L)}$  with  $\mathbf{j} \in \mathcal{J}_{i-1}$  and  $\beta_{i,\mathbf{j}}^{(R)}$  with  $\mathbf{j} \in \mathcal{J}_i$ , respectively. Note that we combine  $\beta_i^{(L)}$  and  $\beta_i^{(R)}$  into the vector  $\beta_i$  as defined in (7). The approximation space  $\mathcal{T}$  is the product space

$$\mathcal{T} = (\mathcal{T}_0 \oplus \mathcal{T}_1) \otimes \dots \otimes (\mathcal{T}_{d-1} \oplus \mathcal{T}_d),$$

where  $\oplus$  and  $\otimes$  denote the sum and the product of two spaces.

### 3.1.3 Numerical solution of optimization problem

Consider the optimization problem (4) and the transformation (5) that shows that it is sufficient to minimize

$$\mathbb{E}_\eta \left[ -\log \pi \circ \tilde{T} - \log \det \nabla \tilde{T} \right] \quad (12)$$

with respect to  $\tilde{T}$ . Note that evaluations of the unnormalized version of  $\pi$  are sufficient to minimize (12); see Section 2.3.2. We have dropped the absolute value above since the maps  $\tilde{T}$  are guaranteed to be monotone, via our parameterization. We now replace the expected value with its sample-average approximation [59], i.e., a Monte Carlo estimator employing independent draws  $\{\vartheta_i\}_{i=1}^n$  from the reference distribution  $\eta$ . Making the dependence on the coefficients  $\beta_1, \dots, \beta_d$  explicit, we obtain the optimization problem

$$\min_{\beta_1, \dots, \beta_d} \frac{1}{n} \sum_{i=1}^n \left[ -\log \pi(\tilde{T}(\vartheta_i; \beta_1, \dots, \beta_d)) - \log \det \nabla \tilde{T}(\vartheta_i; \beta_1, \dots, \beta_d) \right]. \quad (13)$$

Our optimization problem (13) is unconstrained because the constraint  $\nabla \tilde{T} \succ 0$  of (4) is automatically satisfied via the squared-integrated parameterization (6).

## 3.2 Constructing transport maps with low-fidelity models

We propose to approximate the high-fidelity forward model  $G$  with a low-cost, low-fidelity model to reduce the computational costs of constructing a transport map via the optimization problem (13). Note that replacing evaluations of  $G$  with evaluations of low-fidelity models in (13) will introduce an error that we must correct later; see Section 3.4.

Let  $\hat{G} : \mathcal{P} \rightarrow \mathcal{Y}$  be a low-fidelity approximation of  $G$  that maps the parameter onto an observable. The model  $\hat{G}$  gives rise to a low-fidelity potential function

$$\hat{\Phi}_{\mathbf{y}}(\boldsymbol{\theta}) = \frac{1}{2} \left\| \Sigma_{\epsilon}^{-\frac{1}{2}} \left( \hat{G}(\boldsymbol{\theta}) - \mathbf{y} \right) \right\|_2^2,$$

and to the low-fidelity likelihood function  $\hat{L}_{\mathbf{y}} : \mathcal{P} \rightarrow \mathbb{R}$

$$\hat{L}_{\mathbf{y}}(\boldsymbol{\theta}) = \exp \left( -\hat{\Phi}_{\mathbf{y}}(\boldsymbol{\theta}) \right).$$

The probability density of the corresponding low-fidelity posterior distribution is, up to a normalizing constant,

$$\hat{\pi}(\boldsymbol{\theta}) \propto \hat{L}_{\mathbf{y}}(\boldsymbol{\theta})\pi_0(\boldsymbol{\theta}).$$

We now use the low-fidelity posterior density  $\hat{\pi}$  in our variational construction of the map  $\hat{T} \in \mathcal{T}$ . Consider the density  $\hat{\eta}$  that is the pullback of the low-fidelity posterior density  $\hat{\pi}$  through a map  $\hat{T}$ ,

$$\hat{\eta}(\boldsymbol{\vartheta}) = \hat{\pi}(\hat{T}(\boldsymbol{\vartheta}))|\det \nabla \hat{T}(\boldsymbol{\vartheta})|.$$

We can find  $\hat{T}$  by minimizing  $D_{\text{KL}}(\eta||\hat{\eta})$ ; this yields the following optimization problem, analogous to (13) but with  $\hat{\pi}$  replacing  $\pi$ ,

$$\min_{\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_d} \frac{1}{n} \sum_{i=1}^n \left[ -\log \hat{\pi}(\hat{T}(\boldsymbol{\vartheta}_i; \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_d)) - \log \det \nabla \hat{T}(\boldsymbol{\vartheta}_i; \boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_d) \right], \quad (14)$$

with the coefficients  $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_d$  defining  $\hat{T}^* \in \mathcal{T}$ .

### 3.3 Constructing deep transport maps

The optimization problem (14) finds a transport map in the approximation space  $\mathcal{T}$ . Thus, if the approximation space  $\mathcal{T}$  is chosen too coarse, then the pullback of the target distribution through the map is only a poor approximation of the reference distribution. Instead of choosing richer approximation spaces to find more accurate transport maps, which would lead to a large number of coefficients to be optimized for in (14), the work [41, 44] proposes to take compositions of transport maps in coarse approximation spaces.

Let  $T^{(1)}$  be a transport map that pushes forward the reference  $\mu_\eta$  onto the posterior  $\mu_{\hat{\pi}}$  and let  $\hat{T}^{(1)} \in \mathcal{T}$  be a numerical approximation of  $T^{(1)}$  derived with optimization problem (14). Since  $\hat{T}^{(1)}$  is an approximation of the map  $T^{(1)}$ , we obtain

$$\mu_{\hat{\pi}} \approx \mu_{\hat{\pi}}^{(1)} = \hat{T}_{\#}^{(1)} \mu_\eta.$$

Following [45], to account for the discrepancy between  $\mu_{\hat{\pi}}$  and  $\mu_{\hat{\pi}}^{(1)}$ , a second map  $\hat{T}^{(2)} \in \mathcal{T}$  is constructed to push forward  $\mu_{\hat{\pi}}^{(1)}$  to  $\mu_{\hat{\pi}}$ , i.e.,

$$\mu_{\hat{\pi}} \approx \mu_{\hat{\pi}}^{(2)} = \hat{T}_{\#}^{(2)} \mu_{\hat{\pi}}^{(1)}.$$

The aim is that the map  $\hat{T}^{(2)}$  should capture only a small correction from  $\mu_{\hat{\pi}}^{(1)}$  to  $\mu_{\hat{\pi}}^{(2)}$ , so that the composition  $\hat{T}^{(1,2)} = \hat{T}^{(2)} \circ \hat{T}^{(1)}$  more accurately pushes forward  $\mu_\eta$  to  $\mu_{\hat{\pi}}$  (in the sense of KL divergence). This process is repeated  $k$  times to obtain the “deep” map

$$\hat{T}^* = \hat{T}^{(k)} \circ \dots \circ \hat{T}^{(1)}. \quad (15)$$

Note that (15) is only one possible way of constructing deep transport maps. Another possibility is to reverse the order of (15) and compose maps from the right, rather than from the left, by keeping the  $\mu_\eta$  as the reference and pulling back the target  $\mu_\pi$  by the current map approximation. This construction is discussed in [61].

### 3.4 MFMH algorithm

We now use the transport map  $\widehat{T}^*$  to derive an MCMC proposal with the aim of improving sampling efficiency (e.g., reducing integrated autocorrelation time) and at the same time guaranteeing that the stationary distribution of the Markov chain is the posterior distribution corresponding to the high-fidelity model  $G$ .

Since the transport map  $\widehat{T}^*$  is monotone and lower triangular, as enforced by the approximation space  $\mathcal{T}$ , the inverse map  $\widehat{S}^* = \widehat{T}^{*-1}$  is cheap to evaluate by solving  $d$  one-dimensional root-finding problems. We refer to [45, 37] for details on how to invert lower-triangular, monotone maps. The inverse map  $\widehat{S}^*$  pushes forward the true posterior  $\mu_\pi$  onto a distribution that *approximates* the reference distribution. In other words,  $\widehat{S}_\#^* \mu_\pi = \mu_\eta$  only if the map were exact, which in general it is not; otherwise,  $\widehat{S}_\#^* \mu_\pi$  is simply closer to  $\mu_\eta$  (in the sense of KL divergence) than  $\mu_\pi$  was. If  $\mu_\eta$  is chosen to be a Gaussian, then the map  $\widehat{S}^*$  approximately ‘‘Gaussianizes’’  $\mu_\pi$ . As shown in [45], and as will be demonstrated with our numerical results in Section 4, MCMC sampling from the approximate reference distribution  $\widehat{S}_\#^* \mu_\pi$  generally yields higher ESSs for a given computational effort than sampling directly from the posterior distribution. Crucially, because the map  $\widehat{S}^* = \widehat{T}^{*-1}$  is invertible, samples from the approximate reference distribution can be *exactly* (up to machine precision) pushed forward to samples from the posterior distribution via the map  $\widehat{T}^*$ .

Our MFMH algorithm is summarized in Algorithm 2 and follows the preconditioned Metropolis-Hastings algorithm introduced in [45]. The MFMH algorithm has the same inputs as the single-fidelity Metropolis-Hastings in Algorithm 1, except that additionally the maps  $\widehat{T}^*$  and  $\widehat{S}^*$  are required. The current state  $\boldsymbol{\theta}_{i-1}$  is mapped with  $\widehat{S}^*$  onto the approximate reference to obtain  $\boldsymbol{\vartheta}_{i-1}$ . Then, a candidate sample is drawn from the proposal distribution based on  $\boldsymbol{\vartheta}_{i-1}$ , and that candidate sample is mapped back with  $\widehat{T}^*$  to a sample of the posterior  $\pi$ .

Whether the candidate sample is accepted or rejected is based on the probability  $\alpha$  that is determined using the high-fidelity posterior density  $\pi$ ; the maps, derived from the low-fidelity posterior, only affect the proposal distribution. Thus it is guaranteed that the stationary distribution of the chain is the posterior distribution  $\pi$ . Note that there is considerable flexibility in the choice of the reference-space proposal distribution  $q$ . Below we will mostly choose  $q(\cdot|\boldsymbol{\vartheta}_{i-1})$  to be a  $d$ -dimensional Gaussian, *independent* of  $\boldsymbol{\vartheta}_{i-1}$ . Thus our MCMC algorithm reduces to a Metropolis independence sampler [53], whether viewed on the reference space or on the target space. Given the monotonicity of the transport

---

**Algorithm 2** Multifidelity preconditioned Metropolis-Hastings (MFMH)
 

---

- 1: **procedure** PRECONDITIONEDMETROPOLISHASTINGS( $L_{\mathbf{y}}, \pi_0, q, M, \widehat{T}^*, \widehat{S}^*$ )
- 2:   Choose a starting point  $\boldsymbol{\theta}_0$
- 3:   **for**  $i = 1, \dots, M$  **do**
- 4:     Map state  $\boldsymbol{\theta}_{i-1}$  onto reference with  $\widehat{S}^*(\boldsymbol{\theta}_{i-1}) = \boldsymbol{\vartheta}_{i-1}$
- 5:     Draw candidate  $\boldsymbol{\vartheta}'$  from proposal  $q(\cdot | \boldsymbol{\vartheta}_{i-1})$
- 6:     Map candidate  $\boldsymbol{\vartheta}'$  onto target with  $\widehat{T}^*(\boldsymbol{\vartheta}') = \boldsymbol{\theta}'$
- 7:     Compute acceptance probability

$$\alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}') = \min \left\{ 1, \frac{q(\boldsymbol{\vartheta}_{i-1} | \boldsymbol{\vartheta}') L_{\mathbf{y}}(\boldsymbol{\theta}') \pi_0(\boldsymbol{\theta}') |\det \nabla \widehat{T}^*(\boldsymbol{\theta}')|}{q(\boldsymbol{\vartheta}' | \boldsymbol{\vartheta}_{i-1}) L_{\mathbf{y}}(\boldsymbol{\theta}_{i-1}) \pi_0(\boldsymbol{\theta}_{i-1}) |\det \nabla \widehat{T}^*(\boldsymbol{\theta}_{i-1})|} \right\}$$

- 8:     Set the sample  $\boldsymbol{\theta}_i$  to

$$\boldsymbol{\theta}_i = \begin{cases} \boldsymbol{\theta}', & \text{with probability } \alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}'), \\ \boldsymbol{\theta}_{i-1}, & \text{with probability } 1 - \alpha(\boldsymbol{\theta}_{i-1}, \boldsymbol{\theta}') \end{cases}$$

- 9:   **end for**
  - 10:   **return**  $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M$
  - 11: **end procedure**
- 

maps and the full support of the Gaussian reference on  $\mathbb{R}^d$ , the proposal distribution on the target space is guaranteed to dominate the posterior distribution, as required for Metropolis independence sampling to converge.

## 4 Numerical results

This section demonstrates our multifidelity approach on two examples. All runtime measurements were performed on compute nodes with Intel Xeon E5-1620 and 64GB RAM on a single core using a MATLAB implementation.

### 4.1 Diffusion equation with reaction term

We first consider a model with diffusion and a nonlinear reaction term, where our goal is to infer the parameters of the reaction term.

#### 4.1.1 Problem setup

Let  $\Omega = (0, 1)^2 \subseteq \mathbb{R}^2$  and  $\mathcal{P} = \mathbb{R}^2$  and consider the PDE

$$-\nabla^2 u(x_1, x_2; \boldsymbol{\theta}) + g(u(x_1, x_2; \boldsymbol{\theta}), \boldsymbol{\theta}) = 100 \sin(2\pi x_1) \sin(2\pi x_2), \quad \mathbf{x} \in \Omega, \quad (16)$$

with homogeneous Dirichlet boundary conditions, where  $\mathbf{x} = [x_1, x_2]^T$ ,  $\boldsymbol{\theta} = [\theta_1, \theta_2]^T \in \mathcal{P} = \mathbb{R}^2$ , and  $u : \Omega \times \mathcal{P} \rightarrow \mathbb{R}$  is the solution function. The nonlinear

function  $g$  is

$$g(u(\mathbf{x}; \boldsymbol{\theta}), \boldsymbol{\theta}) = (0.1 \sin(\theta_1) + 2) \exp(-2.7\theta_1^2) (\exp(1.8\theta_2 u(\mathbf{x}; \boldsymbol{\theta})) - 1) .$$

We discretize (16) with finite differences on a grid with equidistant grid points and mesh width  $h > 0$ . The corresponding system of nonlinear equations is solved with Newton’s method and inexact line search based on the Armijo condition. The model  $G_h : \mathcal{P} \rightarrow \mathcal{Y}$  derived with mesh width  $h$  maps from  $\mathcal{P}$  into  $\mathcal{Y} = \mathbb{R}^{12}$ . The components of the observable  $\mathbf{y}$  correspond to the values of the approximated solution function at the spatial coordinates  $[0.25i, 0.2j]^T \in \Omega$  with  $i = 1, 2, 3$  and  $j = 1, 2, 3, 4$ .

A low-fidelity model  $\widehat{G}_h$  of  $G_h$  is derived via projection-based model reduction [3]. Solutions of (16) for parameters on an  $100 \times 100$  equidistant grid in  $[-\pi/2, -\pi/2] \times [1, 5] \subset \mathcal{P}$  and mesh width  $h$  are computed and a 20-dimensional reduced space with proper orthogonal decomposition is constructed. The operators corresponding to the high-fidelity  $G_h$  are projected via Galerkin projection onto the reduced space and the low-fidelity model  $\widehat{G}_h$  is obtained. We use a reduced model based on proper orthogonal decomposition in this example; however, we note that our MFMH approach is applicable with other low-fidelity models, e.g., coarse-grid approximations, data-fit surrogate models, and simplified-physics models. Even if the low-fidelity model turns out to be a poor approximation of the high-fidelity model, our MFMH approach guarantees that samples are asymptotically drawn from the high-fidelity posterior distribution, cf. Section 3.4. In case of a poor low-fidelity model, our MFMH approach deteriorates to a Metropolis-Hastings sampler with a poor proposal distribution.

#### 4.1.2 Setup of inverse problems

We set  $\boldsymbol{\theta}^* = [0.5, 2]^T$  and consider the data  $\mathbf{y} = G_H(\boldsymbol{\theta}^*) + \boldsymbol{\epsilon}$ , where  $H = 1/64$  and  $\boldsymbol{\epsilon}$  adds Gaussian noise with zero mean and variance 0.0026, which corresponds to 0.1% noise with respect to the Euclidean norm of the output  $G_H(\boldsymbol{\theta}^*)$ . We then use the model  $G_h$  with  $h = 1/32$  and the corresponding low-fidelity model  $\widehat{G}_h$  for inference. Note that the data  $\mathbf{y}$  is computed with a discretization of the PDE (16) with mesh width  $H = 1/64$ , while we use a mesh width  $h = 1/32$  for inference. The prior distribution is a Gaussian distribution with mean  $[\pi/4, 1.2]^T$  and covariance matrix

$$\boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} = \begin{bmatrix} 64 & 0 \\ 0 & 16 \end{bmatrix} .$$

The low-fidelity  $\widehat{G}_h$  is about 80 times faster to evaluate than the high-fidelity model  $G_h$  for  $h = 1/32$ .

We construct a transport map from the reference Gaussian distribution to the posterior corresponding to the low-fidelity model  $\widehat{G}_h$ . The reference Gaussian distribution has zero mean and standard deviation 0.1. Note that instead we could have scaled and centered the posterior distribution to have zero mean

and marginal variances of one and then used a standard Gaussian as reference distribution.

The optimization is performed with MATLAB’s `fmincon` optimizer, where the tolerance (`TolX`) is set to  $10^{-3}$  and where we use  $n = 250$  samples of the reference distribution to approximate the expected value in the objective function, see Section 3.1.3. We compose two maps following Section 3.3. The approximation space for the first map  $\hat{T}^{(1)}$  corresponds to first-order polynomials (linear) and the approximation space of the second map  $\hat{T}^{(2)}$  corresponds to second-order polynomials (quadratic). The starting point for the optimization is the identity map. The transport map is then used to precondition Metropolis-Hastings as shown in Algorithm 2. The proposal is the reference distribution, i.e., we obtain an independence sampler with a proposal that is independent of the current state of the chain. Because we expect that the transport map approximately pushes the reference onto the posterior, it is reasonable to consider an independence sampler with the reference distribution as proposal. We discard every other sample, which means that we perform  $M = 2m$  iterations if we want  $m$  samples.

We compare sampling with our MFMH approach to delayed-rejection adaptive-metropolis (DRAM) sampling [29]. The DRAM implementation<sup>1</sup> that we use is based on two stages, with the second proposal equivalent to the first proposal scaled down by a factor of 3. The length of the non-adaptive period is 100. All other parameters of DRAM are set to their default values in the implementation mentioned above. We initialize DRAM with a Gaussian proposal that has a diagonal covariance with all elements on the diagonal being equal. We start DRAM with diagonal elements in  $\{10^{-4}, 5 \times 10^{-4}, 10^{-3}, 5 \times 10^{-3}, 10^{-2}, 2 \times 10^{-2}, 3 \times 10^{-2}, 4 \times 10^{-2}, 5 \times 10^{-2}, 10^{-1}, 5 \times 10^{-1}\}$  and then select the run with the highest ESS; cf. Section 2.2. The first  $10^4$  samples are discarded as burn-in and then every other sample is used. This means that DRAM performs  $M = 2m + 10^4$  iterations if we want  $m$  samples. Note that the same thinning of discarding every other sample is applied to the samples obtained with MFMH and DRAM.

### 4.1.3 Results

Figure 1 shows samples drawn with DRAM from the high-fidelity posterior and compares them to samples drawn with MFMH. Both reflect the “banana”-like shape of the posterior. Mapping the samples from Figure 1c to the reference distribution results in the samples shown in Figure 1d. Figure 2 reports the ESS of samples drawn with DRAM and our MFMH approach. For MFMH, we report results for  $m \in \{5 \times 10^2, 10^3, 5 \times 10^3, 10^4\}$  samples and for DRAM for  $m \in \{5 \times 10^3, 10^4, 5 \times 10^4\}$  samples. In this example, MFMH achieves a higher ESS than DRAM with respect to runtime. The reported runtime is the time needed for model evaluations and, in case of MFMH, for constructing and evaluating the transport map. Note that we use a burn-in of  $10^4$  samples for DRAM, whereas such a burn-in is unnecessary in case of our MFMH sampler because it is an

---

<sup>1</sup><http://helios.fmi.fi/~lainema/dram/>

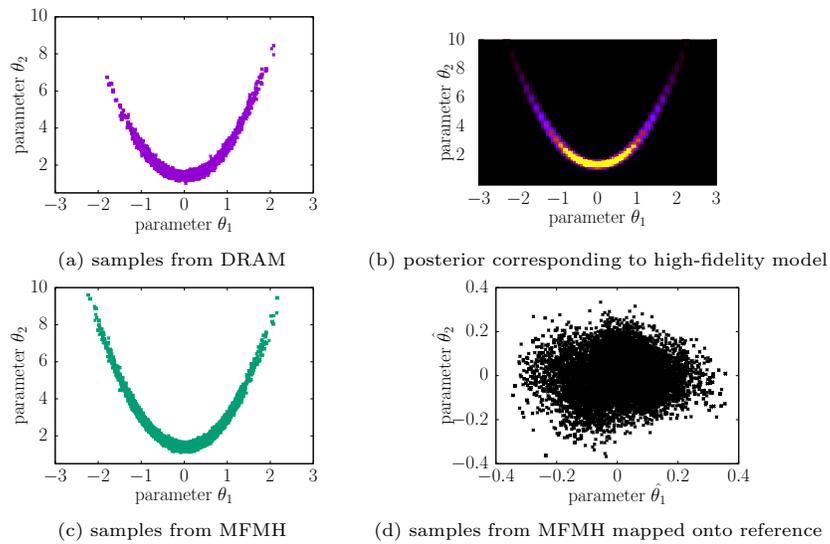


Figure 1: Diffusion-reaction problem: Plots in (a) and (c) visualize samples drawn with DRAM and MFMH, respectively. The “banana”-like shape of the posterior is reflected in both sets of samples, see plot (b). Plot (d) shows the samples from (c) mapped onto the reference distribution, which demonstrates that our transport map captures well the high-fidelity posterior distribution even though it was constructed with a low-fidelity model.

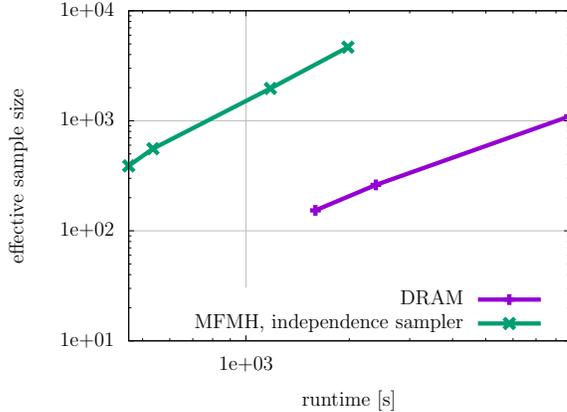


Figure 2: Diffusion-reaction problem: The ESS of samples drawn with our MFMH approach is higher than with DRAM, in this example. The reported runtime includes the runtime of constructing the transport map in case of our MFMH approach. Note that there is no burn-in time for MFMH because it is based on an independence sampler in this example.

independence sampler. Still, this is a fair comparison because MFMH incurs offline costs to find a proposal distribution for independence sampling, instead of a burn-in time as DRAM. The offline costs of our MFMH approach, just as the costs of the burn-in time of DRAM, are included in the runtime plotted in Figure 2. Figure 3a and Figure 3b are trace plots corresponding to DRAM and our MFMH approach, respectively. Both samplers lead to fast mixing. The autocorrelation functions plotted in Figure 3c and Figure 3d for DRAM and MFMH, respectively, indicate that MFMH achieves a lower autocorrelation in this example, which agrees with the finding that our MFMH approach achieves a higher ESS than DRAM in this example, cf. Figure 2. The acceptance rate of our MFMH independence sampler is about 80%, which confirms that the proposal constructed with the transport map is so effective that it can be used for independence sampling.

#### 4.1.4 Results for inference from observations with higher noise

We now add 1% noise with respect to the Euclidean norm of the output  $G_H(\theta^*)$ , which corresponds to a 10 times higher noise level than considered in Section 4.1.2. We additionally modify the nonlinear term in (16) to

$$g(u(\mathbf{x}; \boldsymbol{\theta}), \boldsymbol{\theta}) = 5(0.1 \sin(\theta_1) + 2) \exp(-2.7\theta_1^2) (\exp(2.3\theta_2 u(\mathbf{x}; \boldsymbol{\theta})) - 1) .$$

to prevent the signal from being lost in the noise. Note that this is necessary because otherwise the prior dominates the posterior so that samples are essentially drawn from the prior distribution; independent of the sampling scheme

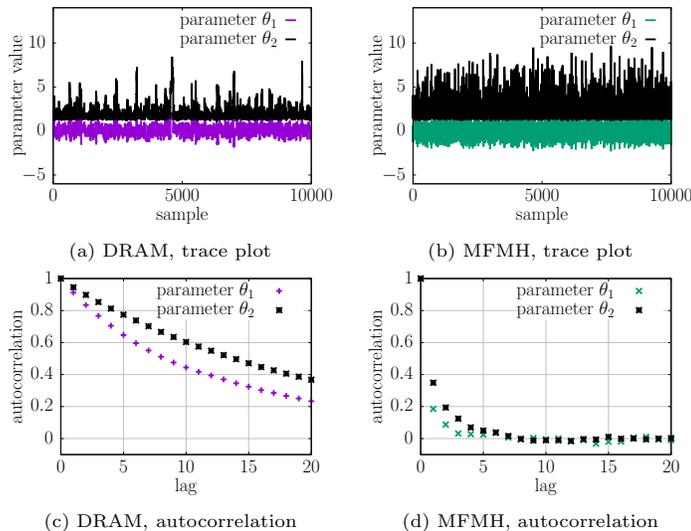


Figure 3: Diffusion-reaction problem: The trace plots in (a) and (b) show fast mixing for DRAM and our MFMH approach. The autocorrelation functions plotted in (c) and (d) indicate a lower autocorrelation of the samples drawn with our MFMH approach than with DRAM in this example.

that is used. The rest of the setup is the same as described in Section 4.1.1 and Section 4.1.2.

Figure 4a shows samples drawn with DRAM. The increased noise leads to a wider banana-like shape of the posterior than in Section 4.1.2 and Figure 1. The samples drawn with MFMH are shown in Figure 4b. Figure 4c compares the ESS of DRAM with  $m \in \{10^3, 5 \times 10^3, 10^4\}$  samples and MFMH with  $m \in \{5 \times 10^2, 10^3, 5 \times 10^3, 10^4\}$  samples, where our MFMH approach achieves similar speedups as in Figure 2 in Section 4.1.3. Thus, the higher noise level seems to have little impact on the performance of our MFMH approach in this example.

## 4.2 Euler Bernoulli beam problem

We now infer the effective stiffness of an Euler Bernoulli beam, for which a model is available on GitHub<sup>2</sup>. The model was developed by Matthew Parno for the 2018 Gene Golub SIAM Summer School on “Inverse Problems: Systematic Integration of Data with Models under Uncertainty.”

<sup>2</sup><https://github.com/g2s3-2018/labs>

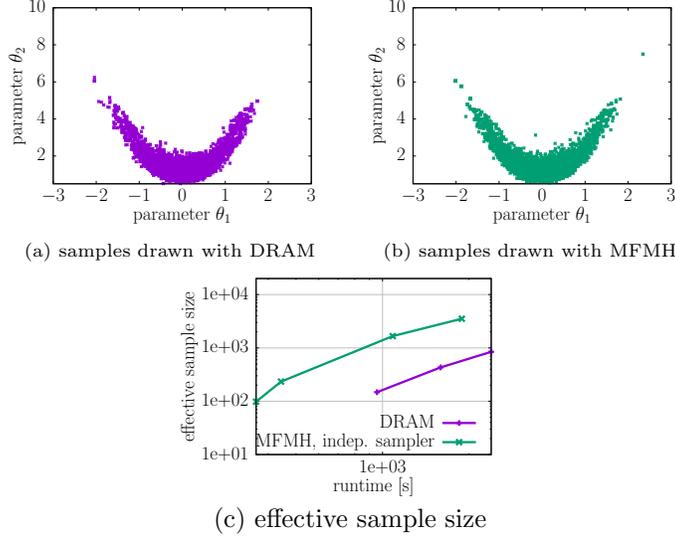


Figure 4: Diffusion-reaction problem: Increasing the noise level to 1% in the observations leads to a wider banana-like shape of the posterior. The speedup achieved with our MFMH approach is similar to the speedup obtained in Figure 2, where the noise level is 0.1%.

#### 4.2.1 Problem setup

Let  $L > 0$  be the length of the beam and define  $\Omega = [0, L] \subset \mathbb{R}$ . Consider the PDE

$$\frac{\partial^2}{\partial x^2} \left( E(x) \frac{\partial^2}{\partial x^2} u(x) \right) = f(x), \quad x \in \bar{\Omega}, \quad (17)$$

where  $u : \Omega \rightarrow \mathbb{R}$  is the vertical deflection of the beam,  $f : \Omega \rightarrow \mathbb{R}$  is the load, and  $\bar{\Omega} = (0, L)$ . The effective stiffness of the beam is given by  $E : \Omega \rightarrow \mathbb{R}$  and describes beam geometry and material properties. The beam is in cantilever configuration, where the left boundary is fixed and the right boundary is free, i.e., the boundary conditions are

$$u(0) = 0, \quad \frac{\partial}{\partial x} u \Big|_{x=0} = 0, \quad \frac{\partial^2}{\partial x^2} u \Big|_{x=L} = 0, \quad \frac{\partial^3}{\partial x^3} u \Big|_{x=L} = 0.$$

The length of the beam is  $L = 1$  in the following. The PDE (17) is discretized with finite differences on a mesh of  $N = 601$  equidistant grid points in  $\Omega$ . The same effective stiffness  $E$  as available in the GitHub repository<sup>1</sup> is used by interpolating on the grid points. Let  $x_{\text{obs}}^{(1)}, \dots, x_{\text{obs}}^{(41)}$  be equidistant points of the  $N = 601$  grid points. The observation  $\mathbf{y} \in \mathbb{R}^{41}$  is the displacement  $u$  at the 41 points  $x_{\text{obs}}^{(1)}, \dots, x_{\text{obs}}^{(41)}$  polluted with zero-mean Gaussian noise with variance  $10^{-4}$ .

### 4.2.2 Models for the Euler Bernoulli problem

We now derive the high-fidelity forward model. Consider the function  $I : \mathbb{R} \times \Omega \rightarrow \mathbb{R}$  defined as

$$I(x, \alpha) = \left( 1 + \exp \left( -\frac{x - \alpha}{0.005} \right) \right)^{-1},$$

with

$$\lim_{x \rightarrow -\infty} I(x, \alpha) = 0, \quad \lim_{x \rightarrow \infty} I(x, \alpha) = 1,$$

such that there is a smooth transition from 0 to 1 at  $\alpha$ . Define further  $k = 3$  and let  $\alpha_1, \dots, \alpha_{k+1}$  be the  $k + 1$  equidistant points in  $\Omega$ . Let  $\mathbb{R}_+ = \{z \in \mathbb{R} : z > 0\}$  and consider the parameter  $\boldsymbol{\theta} = [\theta_1, \dots, \theta_k]^T \in \mathbb{R}_+^k$ . Define the function  $\widehat{E}_i : \Omega \times \mathbb{R} \rightarrow \mathbb{R}$  as

$$\widehat{E}_i(x, \theta_i) = (1 - I(x, \alpha_i))\widehat{E}_{i-1}(x, \theta_{i-1}) + I(x, \alpha_i)\theta_i,$$

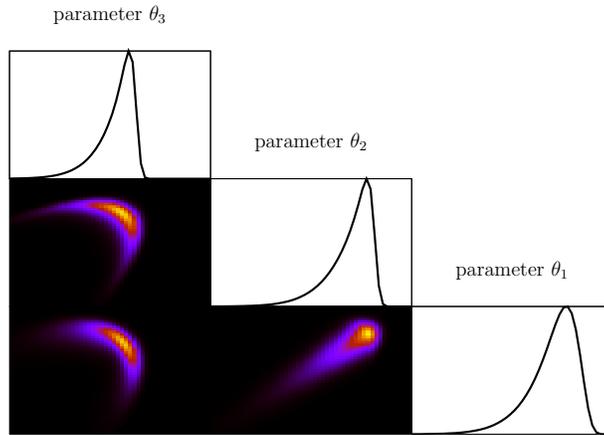
for  $i = 2, \dots, k$  and  $\widehat{E}_1(x, \theta_1) = \theta_1$ . Given a parameter  $\boldsymbol{\theta}$ , the function  $\widehat{E}_k$  is a smooth approximation of the piecewise constant function  $\sum_{i=1}^k \theta_i \mathbb{I}_{(\alpha_i, \alpha_{i+1}]}$ , where  $\mathbb{I}_{(\alpha_i, \alpha_{i+1}]}$  is the indicator function of the interval  $(\alpha_i, \alpha_{i+1}] \subset \mathbb{R}$ . The high-fidelity forward model  $G$  maps a parameter  $\boldsymbol{\theta} \in \mathbb{R}_+^k$  onto the displacement  $u$  with effective stiffness  $\widehat{E}_k$  at the observation points  $x_{\text{obs}}^{(1)}, \dots, x_{\text{obs}}^{(41)}$ . The same discretization as described in Section 4.2.1 is used. The low-fidelity model  $\widehat{G}$  is a spline interpolant of  $G$  on a logarithmically spaced grid in the domain  $[0.5, 4]^3$ . The interpolant is obtained with the `griddedInterpolant` procedure available in MATLAB. Extrapolation is turned off. The low-fidelity model  $\widehat{G}$  is about 1,100 times faster to evaluate than the high-fidelity model  $G$ .

### 4.2.3 Setup of inverse problem

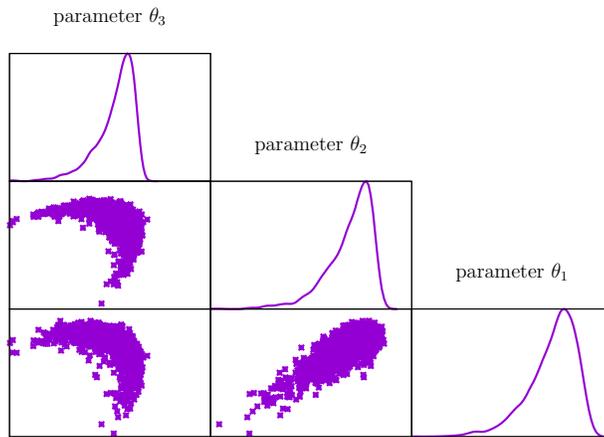
The observation  $\mathbf{y}$  is obtained as described in Section 4.2.1. The prior is a log-normal distribution with mean  $[1, 1, 1]^T$  and covariance matrix

$$\boldsymbol{\Sigma}_\epsilon = \begin{bmatrix} 0.05 & 0 & 0 \\ 0 & 0.05 & 0 \\ 0 & 0 & 0.05 \end{bmatrix}.$$

We construct a transport map  $\widehat{T}$  using the low-fidelity model  $\widehat{G}$ . The transport map  $\widehat{T}$  is based on quadratic polynomials and approximately maps the reference Gaussian distribution with mean  $[1, 1, 1]^T$  and covariance with diagonal entries 0.1 onto the low-fidelity posterior distribution. Note that a standard Gaussian can be used as reference distribution if the posterior distribution is centered to zero mean and scaled to have marginal variances of one. This effectively would perform a linear transformation. We optimize for the coefficients of  $\widehat{T}$  with `fmincon` available in MATLAB with tolerance set to  $10^{-8}$  (`TolX` option) and with 500 samples from the reference distribution. The initial point of the



(a) dimension-wise heat map of posterior density function



(b) samples drawn from posterior distribution with DRAM

Figure 5: Euler Bernoulli beam: Plot (a) visualizes the posterior density function. Plot (b) shows samples drawn with DRAM from the posterior distribution. The plots on the diagonal of (b) show marginal densities estimated from the samples.

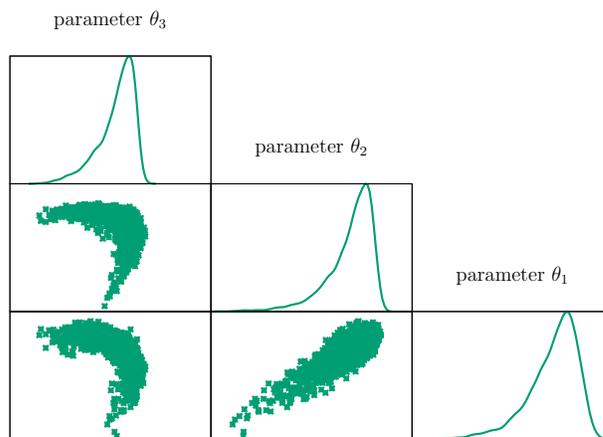
optimization corresponds to the identity map. We then use the transport map in Algorithm 2 with a random-walk Metropolis algorithm on the reference space, i.e., with a local Gaussian proposal  $\boldsymbol{\vartheta}' \sim \mathcal{N}(\boldsymbol{\vartheta}_{i-1}, \boldsymbol{\Sigma}')$ . The covariance matrix  $\boldsymbol{\Sigma}'$  is diagonal. To select the diagonal of  $\boldsymbol{\Sigma}'$ , we ran MFMH for proposals with variance in  $\{10^{-4}, 5 \times 10^{-4}, 10^{-3}, 5 \times 10^{-3}, 10^{-2}, 2 \times 10^{-2}, 3 \times 10^{-2}, 4 \times 10^{-2}, 5 \times 10^{-2}, 10^{-1}, 5 \times 10^{-1}\}$  and then selected the variance that leads to the highest ESS, cf. Section 4.1.2. The corresponding results are reported in the following. Additionally, we consider an independence MFMH sampler where the reference distribution together with the transport map serves as a proposal distribution that is independent of the previous sample. The rest of the setup is the same as in Section 4.1.2.

#### 4.2.4 Results

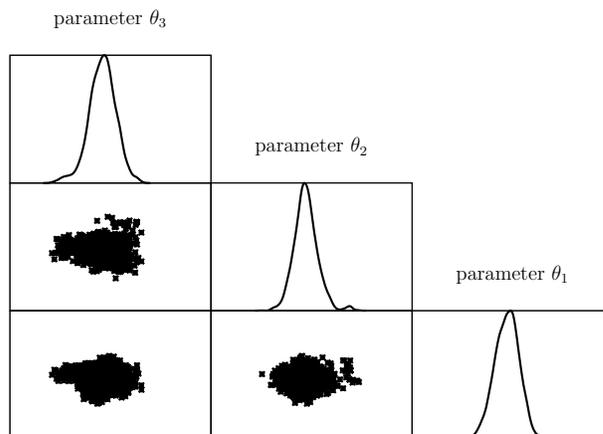
Figure 5a visualizes the posterior density function corresponding to the high-fidelity model. The plot indicates that the posterior distribution is a non-Gaussian distribution. We use the same procedure as in Section 4.1.2 to draw samples from the posterior with DRAM. The samples with the highest ESS are shown in Figure 5b.

Let us now consider our MFMH approach. First, we consider the independence MFMH sampler that uses a proposal that is independent of the previous sample, see Figure 6a. The samples are in agreement with the samples drawn with DRAM from the high-fidelity posterior distribution, cf. Figure 5b. Figure 6b plots the samples mapped with the transport map onto an approximation of the reference Gaussian distribution, which demonstrates that the transport map  $\widehat{T}^*$  captures well the posterior distribution. The ESSs of  $m \in \{10^5, 5 \times 10^5, 10^6, 5 \times 10^6\}$  samples drawn with DRAM and the independence MFMH sampler are compared in Figure 7. Our independence MFMH sampler achieves a higher ESS than DRAM in this example. Figure 7 also shows that a higher ESS compared to DRAM is achieved only if sufficiently many samples are drawn because otherwise the offline costs of constructing the transport maps are not compensated. The MFMH sampler with a local random walk proposal leads to a higher ESS than DRAM too, but the improvement is smaller compared to the independence MFMH sampler. Additionally, an even larger number of samples is necessary to compensate the offline costs of constructing the transport maps. Note that the MFMH sampler with a local random walk proposal and the DRAM sampler both use a burn-in of  $10^4$  samples as in Section 4.1.2.

Trace plots for our MFMH samplers are provided in Figure 8. Both samplers achieve fast mixing in all three parameters  $\theta_1, \theta_2$ , and  $\theta_3$ . Figure 9 shows the autocorrelation functions for our MFMH samplers, which seem to decay faster than the autocorrelation functions corresponding to DRAM in this example. The faster decay of the autocorrelation functions agrees with the higher ESS of our MFMH samplers reported in Figure 7. The acceptance rate of the MFMH sampler with a local random walk proposal is about 41% and that of the MFMH independence sampler is about 75%.



(a) samples drawn with MFMH



(b) samples mapped onto an approximation of the reference Gaussian distribution

Figure 6: Euler Bernoulli beam: Plot (a) shows samples drawn with our MFMH method using a transport map constructed with a low-fidelity model. Mapping the samples shown in (a) with the transport map gives samples of an approximate reference Gaussian distribution, see plot (b). The plots on the diagonal of (a) and (b) show marginal densities estimated from the samples.

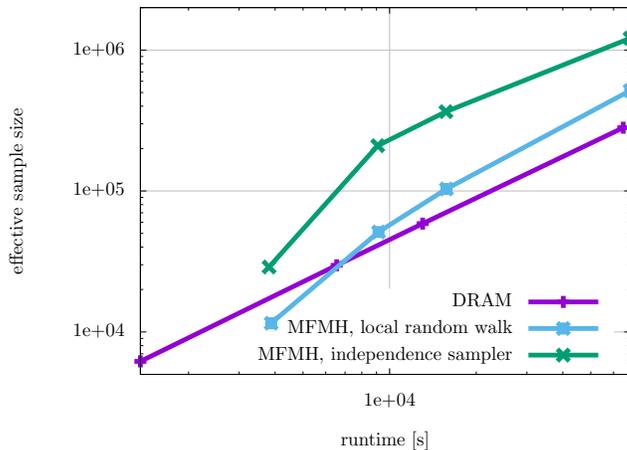


Figure 7: Euler Bernoulli beam: Our MFMH approach leads to a higher ESS than DRAM after the offline costs of constructing the transport map are compensated. The runtime is the total time of drawing samples, including the runtime of constructing the transport map with the low-fidelity model in case of sampling with MFMH. The points on the curves correspond to  $m \in \{10^5, 5 \times 10^5, 10^6, 5 \times 10^6\}$  samples, respectively.

## 5 Conclusions

Our MFMH approach leverages low-fidelity models to precondition MCMC sampling, with the aim of improving MCMC mixing while guaranteeing that the stationary distribution of the chain is the posterior distribution corresponding to the high-fidelity model. In a one-time expensive step, a low-fidelity model is used to construct a transport map that approximately maps an easy-to-sample reference distribution to the posterior distribution corresponding to the low-fidelity model. In the second step, the transport map is used to precondition the posterior distribution corresponding to the high-fidelity model during sampling with Metropolis-Hastings. Since the low-fidelity model is used for preconditioning only, the stationary distribution of the chain obtained in the second step is the posterior distribution corresponding to the high-fidelity model. Our MFMH approach achieves significant speedups compared to single-fidelity sampling with DRAM in our numerical examples.

This work focused on problems where evaluating the likelihoods are computationally expensive because solving the forward models is demanding. A direction of future research is to investigate our MFMH approach for problems with higher-dimensional parameter spaces. In [45] and [41], transport maps are used in up to 10 and 139 parameter dimensions, respectively, which shows that computations with transport maps similar to our MFMH approach are tractable in higher dimensions. A new building block for our MFMH approach in higher dimensions will be theoretical results in [61, 60] which show that conditional

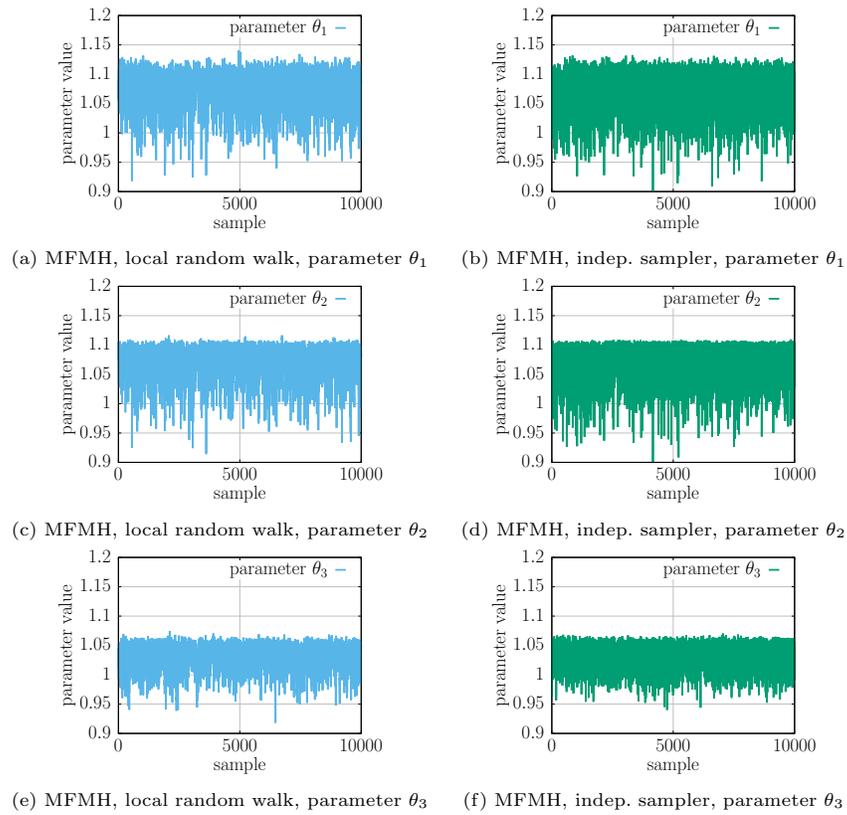


Figure 8: Euler Bernoulli beam: Trace plots of our MFMH samplers for all three parameters  $\theta_1, \theta_2$ , and  $\theta_3$ .

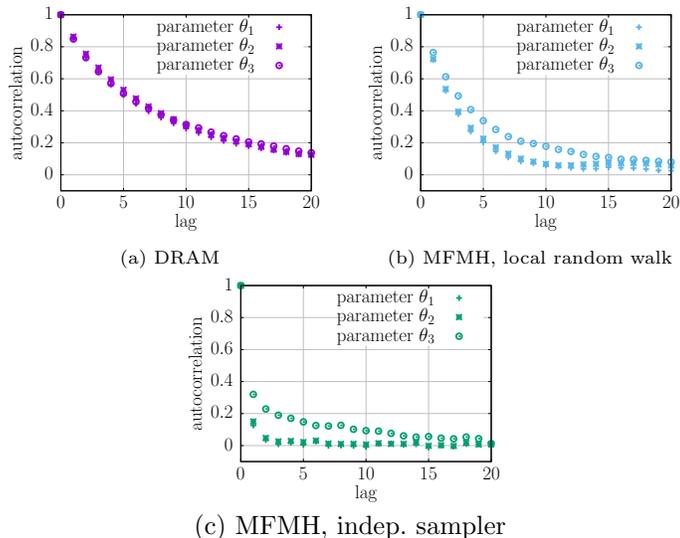


Figure 9: Euler Bernoulli beam: The autocorrelation functions corresponding to our MFMH samplers seem to decrease faster than the autocorrelation function corresponding to DRAM in this example, which confirms the higher ESS achieved by our MFMH approach; see Figure 7.

independence structure in the posterior leads to sparse or low-rank transports, making representing and computing transport maps in even higher dimensions tractable. Another option that we will explore is to discretize transport maps on sparse grids [10, 51]. In a different line of future work, we will investigate the optimal trade-off between low-fidelity model accuracy and low-fidelity model costs so that the low-fidelity model is sufficiently accurate for improving, e.g., the ESS with MFMH while achieving speedups by keeping the offline costs low. A similar question of finding optimal models for multifidelity Monte Carlo estimation has been investigated in [46]. We plan to build on [46] to derive similar complexity trade-offs for MFMH.

## References

- [1] Allaire, D., Willcox, K.: A mathematical and computational framework for multifidelity design and analysis with computer models. *International Journal for Uncertainty Quantification* **4**(1), 1–20 (2014)
- [2] Balakrishnan, S., Roy, A., Ierapetritou, M.G., Flach, G.P., Georgopoulos, P.G.: Uncertainty reduction and characterization for complex environmental fate and transport models: An empirical Bayesian framework incorporating the stochastic response surface method. *Water Resources Research* **39**(12), 1–13 (2003)

- [3] Benner, P., Gugercin, S., Willcox, K.: A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review* **57**(4), 483–531 (2015)
- [4] Beskos, A., Jasra, A., Law, K., Tempone, R., Zhou, Y.: Multilevel sequential Monte Carlo samplers. *Stochastic Processes and their Applications* **127**(5), 1417 – 1440 (2017)
- [5] Bigoni, D., Spantini, A., Morrison, R., Baptista, R.M.: Documentation of TransportMaps software package. Tech. rep., Uncertainty Quantification Group, Massachusetts Institute of Technology (2018). URL <http://transportmaps.mit.edu>
- [6] Bonnotte, N.: From Knothe’s rearrangement to Brenier’s optimal transport map. *SIAM Journal on Mathematical Analysis* **45**(1), 64–87 (2013)
- [7] Boyaval, S.: A fast Monte–Carlo method with a reduced basis of control variates applied to uncertainty propagation and Bayesian estimation. *Computer Methods in Applied Mechanics and Engineering* **241–244**, 190–205 (2012)
- [8] Boyaval, S., Lelièvre, T.: A variance reduction method for parametrized stochastic differential equations using the reduced basis paradigm. *Communications in Mathematical Sciences* **8**(3), 735–762 (2010)
- [9] Brooks, S., Gelman, A., Jones, G.L., Meng, X.L.: *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC (2011)
- [10] Bungartz, H.J., Griebel, M.: Sparse grids. *Acta Numerica* **13**, 147–269 (2004)
- [11] Carlier, G., Galichon, A., Santambrogio, F.: From Knothe’s transport to Brenier’s map and a continuation method for optimal transport. *SIAM Journal on Mathematical Analysis* **41**(6), 2554–2576 (2010)
- [12] Chen, P., Quarteroni, A.: Accurate and efficient evaluation of failure probability for partial differential equations with random input data. *Computer Methods in Applied Mechanics and Engineering* **267**, 233 – 260 (2013)
- [13] Chen, P., Quarteroni, A., Rozza, G.: Reduced basis methods for uncertainty quantification. *SIAM/ASA Journal on Uncertainty Quantification* **5**(1), 813–869 (2017)
- [14] Chen, P., Schwab, C.: Sparse-grid, reduced-basis Bayesian inversion. *Computer Methods in Applied Mechanics and Engineering* **297**, 84 – 115 (2015)
- [15] Christen, J.A., Fox, C.: Markov chain Monte Carlo using an approximation. *Journal of Computational and Graphical Statistics* **14**(4), 795–810 (2005)

- [16] Conrad, P., Davis, A., Marzouk, Y., Pillai, N., Smith, A.: Parallel local approximation MCMC for expensive models. *SIAM/ASA Journal on Uncertainty Quantification* **6**(1), 339–373 (2018)
- [17] Conrad, P.R., Marzouk, Y., Pillai, N.S., Smith, A.: Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. *Journal of the American Statistical Association* **111**(516), 1591–1607 (2016)
- [18] Cortes, C., Vapnik, V.: Support-vector networks. *Machine Learning* **20**(3), 273–297 (1995)
- [19] Cotter, S., Dashti, M., Stuart, A.: Approximation of Bayesian inverse problems for PDEs. *SIAM Journal on Numerical Analysis* **48**(1), 322–345 (2010)
- [20] Cui, T., Marzouk, Y., Willcox, K.: Data-driven model reduction for the Bayesian solution of inverse problems. *International Journal for Numerical Methods in Engineering* **102**(5), 966–990 (2015)
- [21] Dodwell, T.J., Ketelsen, C., Scheichl, R., Teckentrup, A.L.: A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow. *SIAM/ASA Journal on Uncertainty Quantification* **3**(1), 1075–1108 (2015)
- [22] Drohmann, M., Carlberg, K.: The ROMES method for statistical modeling of reduced-order-model error. *SIAM/ASA Journal on Uncertainty Quantification* **3**(1), 116–145 (2015)
- [23] Efendiev, Y., Hou, T., Luo, W.: Preconditioning Markov chain Monte Carlo simulations using coarse-scale models. *SIAM Journal on Scientific Computing* **28**(2), 776–803 (2006)
- [24] Eldred, M.S., Ng, L.W.T., Barone, M.F., Domino, S.P.: Multifidelity uncertainty quantification using spectral stochastic discrepancy models. In: R. Ghanem, D. Higdon, H. Owhadi (eds.) *Handbook of Uncertainty Quantification*, pp. 1–45. Springer International Publishing, Cham (2016)
- [25] Forrester, A.I.J., Keane, A.J.: Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences* **45**(1–3), 50–79 (2009)
- [26] Frangos, M., Marzouk, Y., Willcox, K., van Bloemen Waanders, B.: Surrogate and Reduced-Order Modeling: A Comparison of Approaches for Large-Scale Statistical Inverse Problems, pp. 123–149. John Wiley & Sons, Ltd (2010)
- [27] Gilks, W.R., Richardson, S., Spiegelhalter, D.J.: *Markov Chain Monte Carlo in Practice*. Chapman & Hall (1996)
- [28] Gugercin, S., Antoulas, A.: A survey of model reduction by balanced truncation and some new results. *International Journal of Control* **77**(8), 748–766 (2004)

- [29] Haario, H., Laine, M., Mira, A., Saksman, E.: DRAM: Efficient adaptive MCMC. *Statistics and Computing* **16**(4), 339–354 (2006)
- [30] Haario, H., Saksman, E., Tamminen, J.: An adaptive Metropolis algorithm. *Bernoulli* **7**(2), 223–242 (2001)
- [31] Kaipio, J., Somersalo, E.: *Statistical and computational inverse problems*. Springer (2005)
- [32] Latz, J., Papaioannou, I., Ullmann, E.: Multilevel sequential<sup>2</sup> Monte Carlo for Bayesian inverse problems. *Journal of Computational Physics* **368**, 154 – 178 (2018)
- [33] Li, J., Marzouk, Y.: Adaptive construction of surrogates for the Bayesian solution of inverse problems. *SIAM Journal on Scientific Computing* **36**(3), A1163–A1186 (2014)
- [34] Lieberman, C., Willcox, K., Ghattas, O.: Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing* **32**(5), 2523–2542 (2010)
- [35] Liu, J.S.: *Monte Carlo Strategies in Scientific Computing*. Springer (2008)
- [36] Manzoni, A., Pagani, S., Lassila, T.: Accurate solution of Bayesian inverse uncertainty quantification problems combining reduced basis methods and reduction error models. *SIAM/ASA Journal on Uncertainty Quantification* **4**(1), 380–412 (2016)
- [37] Marzouk, Y., Moselhy, T., Parno, M., Spantini, A.: Sampling via measure transport: An introduction. In: R. Ghanem, D. Higdon, H. Owhadi (eds.) *Handbook of Uncertainty Quantification*, pp. 1–41. Springer International Publishing, Cham (2016)
- [38] Marzouk, Y., Najm, H.N.: Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *Journal of Computational Physics* **228**(6), 1862 – 1902 (2009)
- [39] Marzouk, Y., Najm, H.N., Rahn, L.A.: Stochastic spectral methods for efficient Bayesian solution of inverse problems. *Journal of Computational Physics* **224**(2), 560 – 586 (2007)
- [40] Marzouk, Y., Xiu, D.: A stochastic collocation approach to Bayesian inference in inverse problems. *Commun. Comput. Phys.* **6**, 826–847 (2009)
- [41] Moselhy, T.A.E., Marzouk, Y.: Bayesian inference with optimal maps. *Journal of Computational Physics* **231**(23), 7815 – 7850 (2012)
- [42] Ng, L., Willcox, K.: Multifidelity approaches for optimization under uncertainty. *International Journal for Numerical Methods in Engineering* **100**(10), 746–772 (2014)

- [43] Ng, L., Willcox, K.: Monte-Carlo information-reuse approach to aircraft conceptual design optimization under uncertainty. *Journal of Aircraft* pp. 1–12 (2015)
- [44] Parno, M.: Transport maps for accelerated Bayesian computation. Ph.D. thesis, Massachusetts Institute of Technology (2015)
- [45] Parno, M., Marzouk, Y.: Transport map accelerated Markov chain Monte Carlo. *SIAM/ASA Journal on Uncertainty Quantification* **6**(2), 645–682 (2018)
- [46] Peherstorfer, B.: Multifidelity Monte Carlo estimation with adaptive low-fidelity models. *SIAM/ASA Journal on Uncertainty Quantification* (2019). (accepted)
- [47] Peherstorfer, B., Cui, T., Marzouk, Y., Willcox, K.: Multifidelity importance sampling. *Computer Methods in Applied Mechanics and Engineering* **300**, 490–509 (2016)
- [48] Peherstorfer, B., Kramer, B., Willcox, K.: Multifidelity preconditioning of the cross-entropy method for rare event simulation and failure probability estimation. *SIAM/ASA Journal on Uncertainty Quantification* **6**(2), 737–761 (2018)
- [49] Peherstorfer, B., Willcox, K., Gunzburger, M.: Optimal model management for multifidelity Monte Carlo estimation. *SIAM Journal on Scientific Computing* **38**(5), A3163–A3194 (2016)
- [50] Peherstorfer, B., Willcox, K., Gunzburger, M.: Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review* (2018)
- [51] Pflüger, D., Peherstorfer, B., Bungartz, H.: Spatially adaptive sparse grids for high-dimensional data-driven problems. *Journal of Complexity* **26**(5), 508–522 (2010)
- [52] Qian, E., Peherstorfer, B., O’Malley, D., Vesselinov, V.V., Willcox, K.: Multifidelity Monte Carlo estimation of variance and sensitivity indices. *SIAM/ASA Journal on Uncertainty Quantification* **6**(2), 683–706 (2018)
- [53] Robert, C.P., Casella, G.: *Monte Carlo Statistical Methods*. Springer (2004)
- [54] Roberts, G.O., Rosenthal, J.S.: Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics* **18**(2), 349–367 (2009)
- [55] Rosenblatt, M.: Remarks on a multivariate transformation. *Ann. Math. Statist.* **23**(3), 470–472 (1952)

- [56] Rozza, G., Huynh, D., Patera, A.: Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Archives of Computational Methods in Engineering* **15**(3), 1–47 (2007)
- [57] Santin, G., Wittwar, D., Haasdonk, B.: Greedy regularized kernel interpolation. *ArXiv e-prints* 1807.09575 (2018)
- [58] Sirovich, L.: Turbulence and the dynamics of coherent structures. *Quarterly of Applied Mathematics* **45**, 561–571 (1987)
- [59] Spall, J.: *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. John Wiley & Sons (2003)
- [60] Spantini, A.: *On the low-dimensional structure of Bayesian inference*. Ph.D. thesis, Massachusetts Institute of Technology (2017)
- [61] Spantini, A., Bigoni, D., Marzouk, Y.: Inference via low-dimensional couplings. *Journal of Machine Learning Research* **19**(66), 1–71 (2018)
- [62] Stuart, A.M.: *Inverse problems: A Bayesian perspective*. *Acta Numerica* **19**, 451–559 (2010)
- [63] Tarantola, A.: *Inverse Problem Theory*. Elsevier (1987)
- [64] Tierney, L.: Markov chains for exploring posterior distributions. *Ann. Statist.* **22**(4), 1701–1728 (1994)
- [65] Vapnik, V.: *Statistical Learning Theory*. Wiley (1998)
- [66] Vershik, A.M.: Long history of the Monge-Kantorovich transportation problem. *The Mathematical Intelligencer* **35**(4), 1–9 (2013)
- [67] Villani, C.: *Topics in Optimal Transportation*. American Mathematical Society (2003)
- [68] Villani, C.: *Optimal Transport: Old and New*. Springer (2009)
- [69] Wang, J., Zabarvas, N.: Using Bayesian statistics in the estimation of heat source in radiation. *International Journal of Heat and Mass Transfer* **48**(1), 15 – 29 (2005)
- [70] Wirtz, D., Haasdonk, B.: A vectorial kernel orthogonal greedy algorithm. *Dolomites Research Notes on Approximation* **6**, 83–100 (2013)