

**Low-rank tensor reconstruction of concentrated densities with  
application to Bayesian inversion**

Martin Eigel, Robert Gruhlke, Manuel Marschall

submitted: December 20, 2019 (revision: August 12, 2020)

Weierstrass Institute  
Mohrenstr. 39  
10117 Berlin  
Germany  
E-Mail: martin.eigel@wias-berlin.de  
robert.gruhlke@wias-berlin.de  
manuel.marschall@wias-berlin.de

No. 2672  
Berlin 2019



---

2010 *Mathematics Subject Classification.* 62F15, 62G07, 35R60, 60H35, 65C20, 65N12, 65N22, 65J10.

*Key words and phrases.* Tensor train, uncertainty quantification, VMC, low-rank, reduced order model, Bayesian inversion, partial differential equations with random coefficients.

M. Eigel and R. Gruhlke have been supported by the DFG SPP1886 “Polymorphic uncertainty modelling for the numerical design of structures“.

Edited by  
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)  
Leibniz-Institut im Forschungsverbund Berlin e. V.  
Mohrenstraße 39  
10117 Berlin  
Germany

Fax: +49 30 20372-303  
E-Mail: [preprint@wias-berlin.de](mailto:preprint@wias-berlin.de)  
World Wide Web: <http://www.wias-berlin.de/>

# Low-rank tensor reconstruction of concentrated densities with application to Bayesian inversion

Martin Eigel, Robert Gruhlke, Manuel Marschall

**ABSTRACT.** Transport maps have become a popular mechanic to express complicated probability densities using sample propagation through an optimized push-forward. Beside their broad applicability and well-known success, transport maps suffer from several drawbacks such as numerical inaccuracies induced by the optimization process and the fact that sampling schemes have to be employed when quantities of interest, e.g. moments are to compute. This paper presents a novel method for the accurate functional approximation of probability density functions (PDF) that copes with those issues. By interpreting the pull-back result of a target PDF through an inexact transport map as a perturbed reference density, a subsequent functional representation in a more accessible format allows for efficient and more accurate computation of the desired quantities. We introduce a layer-based approximation of the perturbed reference density in an appropriate coordinate system to split the high-dimensional representation problem into a set of independent approximations for which separately chosen orthonormal basis functions are available. This effectively motivates the notion of  $h$ - and  $p$ -refinement (i.e. “mesh size” and polynomial degree) for the approximation of high-dimensional PDFs. To circumvent the curse of dimensionality and enable sampling-free access to certain quantities of interest, a low-rank reconstruction in the tensor train format is employed via the Variational Monte Carlo method. An a priori convergence analysis of the developed approach is derived in terms of Hellinger distance and the Kullback-Leibler divergence. Applications comprising Bayesian inverse problems and several degrees of concentrated densities illuminate the (superior) convergence in comparison to Monte Carlo and Markov-Chain Monte Carlo methods.

## 1. OVERVIEW

We derive a novel numerical method for the functional representation of complicated (in particular highly concentrated) probability densities. This difficult task usually is attacked with Markov Chain Monte Carlo (MCMC) methods which yield samples of the posterior. Despite their popularity, the convergence rate of these methods is ultimately limited by the employed Monte Carlo sampling technique, see e.g. [13] for recent multilevel techniques in this context. Moreover, practical issues e.g. regarding the initial number of samples (burn-in) or a specific convergence assessment arise.

In this work, we propose a new approach based on *function space representations with efficient surrogate models* in several instances. This is motivated by our previous work on adaptive low-rank approximations of solutions of parametric random PDEs with Adaptive Stochastic Galerkin FEM (ASGFEM, see e.g. [20, 17]) and in particular the sampling-free Bayesian inversion presented in [18] where the setting of uniform random variables was examined. A generalization to the important case of Gaussian random variables turns out to be non-trivial from a computational point of view due to the difficulties of representing highly concentrated densities in a compressing tensor format which is required in order to cope with the high dimensionality of the problem. As a consequence, we develop a discretization approach which takes into account the potentially problematic structure of the probability density at hand by a combination of several transformations and approximations that can be chosen adaptively to counteract the interplay of the employed numerical approximations. With the computed functional representation of the density, the evaluation of moments or other statistical quantities of interest can be carried out efficiently and with high accuracy.

The central idea of the method is to obtain a map which transports the target density to some convenient reference density and employ low-rank regression techniques to obtain a functional representation, for which accurate numerical methods are available. Transport maps for probability densities are a classical topic in mathematics, cf.[56, 49]. They are under active research in particular in the

area of optimal transport [56, 49] and also have become popular in current machine learning research [55, 46, 12]. A main application we have in mind is Bayesian inversion where, given a prior density and some observations of the forward model, a posterior density should be determined. In this context, the rescaling approaches in [50, 51] based on the Laplace approximation can be considered as transport maps of a certain (affine) form. More general transport maps have been examined extensively in [22, 44] and other works of the research group. Obtaining a transport map is in general realized by minimizing a certain loss functional, e.g. the Kullback-Leibler distance, between the target and a pushed-forward reference density. This process has been analyzed and improved using iterative maps [6] or multi-scale approaches [43]. However, the optimization, the loss functional and the chosen model class for the transport map yield only an approximation to an *exact* transport. We hence suppose that, in general, an inexact transport is available. By a pull-back argument, this can be interpreted as starting from a different or slightly perturbed reference density. The degree of perturbation has then to be coped with in subsequent approximation steps to enable an explicit representation of this new reference and make quantities of interest (QoI) directly accessible. Finding a suitable approximation relies on concepts from adaptive finite element methods (FEM). In addition to the selection of (local) approximation spaces of a certain degree (“p-refinement”), we introduce a spatial decomposition of the density representation into layers (“h-refinement”) around some center of mass of the considered density. This enables to exploit the decay behavior of the approximated density. Overall, this “hp-refinement” allows to balance inaccuracies and hence perturbations of the reference density by putting more effort into the discretization part. One hence has the freedom to decide whether more effort should be invested into computing an exact transport map or into a more elaborate discretization (with more layers and larger basis) of the perturbed reference density.

For eventual computations with the devised (possibly high-dimensional) functional density representation, an efficient representation format is required. In our context, hierarchical tensors and in particular tensor trains (TT) prove to be advantageous, cf. [3, 41]. These low-rank formats enable to alleviate the curse of dimensionality under suitable conditions and allow for efficient evaluations of high-dimensional objects. For each layer of the discretization we aim to obtain a low-rank tensor representation of the respective perturbed reference density. In certain ideal cases, such as transporting to the standard Gaussian density, a rank-one representation is sufficient. Having a perturbed reference density that is Gaussian but not standard normal, the theory in [47] applies. In more general cases, a low-rank representability may be observed numerically. To allow for tensor methods to be applicable, the desired discretization layers have to be tensor domains. Therefore, the underlying perturbed reference density is transformed to an alternative coordinate system which benefits the representation and allows to exploit the regularity and decay behavior of the density. To generate a tensor train representation (coupled with a function basis which is then also called extended or functional TT format [28]), the Variational Monte Carlo (VMC) method [21] is employed. It basically is a tensor regression approach based on function samples for which a convergence analysis is available. Notably, depending on the chosen loss functional, it leads to the best approximation in the respective model space. It has previously been examined in the context of random PDEs in [21] as an alternative nonintrusive numerical approach to Stochastic Galerkin FEM in the TT format [20, 17]. The approximation of [17] is used in one of the presented examples for Bayesian inversion with the random Darcy equation with lognormal coefficient. We note that surrogate models of the forward model have been used in the context of MCMC e.g. in [34] and tensor representations (obtained by cross approximation) were used in [14] to improve the efficiency of MCMC sampling.

The derivation of our method is supported by an a priori convergence analysis with respect to the Hellinger distance and the Kullback-Leibler divergence. In the analysis, different error sources have to be considered, in particular a layer truncation error depending on decay properties of the density, a low-rank truncation error and model space approximations are introduced. Moreover, the VMC error analysis [21] comprising statistical estimation and numerical approximation errors is adjusted to be

applicable to the devised approach. While not usable for an a posteriori error control in its current initial form, the derived analysis leads the way to more elaborate results for this promising method in future research.

With the constructed functional density surrogate, sampling-free computations of statistical quantities of interest such as moments or marginals become feasible by fast tensor contractions, even for highly concentrated or (depending on the available transport map) nonlinearly transformed high-dimensional densities.

While several assumptions have to be satisfied for this method to work most efficiently, the approach is rather general and can be further adapted to the considered problem. Moreover, it should be emphasized that by constructing a functional representation, structural properties of the density at hand (in particular smoothness, sparsity, low-rank approximability and decay behavior in different parameters) can be exploited in a much more extensive way than what is possible with sampling based methods such as MCMC, leading to more accurate statistical computations and better convergence rates. We note that the perturbed posterior surrogate can be used to efficiently generate samples by rejection sampling or within a MCMC scheme. Since the perturbed transport can be seen as a preconditioner, the sample generation can be based on the perturbed prior. These samples can then be pushed forward to the posterior. As a prospective extension, the constructed posterior density could directly be used in a Stochastic Galerkin FEM based on the integral structure, closing the loop of forward and inverse problem, resulting in the inferred forward problem with model data determined by Bayesian inversion from the observed data.

The structure of the paper is as follows. Section 2 is concerned with the representation of probability densities and introduces a relation between a target and a reference density. Such a transport map can be determined numerically by approximation in a chosen class of functions and with an assumed structure, leading to the concept of perturbed reference densities. To counteract the perturbation, a layered truncated discretization is introduced. An efficient low-rank representation of the mappings is described in Section 3 where the tensor train format is discussed. In order to obtain this nonintrusively, the Variational Monte Carlo (VMC) tensor reconstruction is reviewed. A priori convergence results with respect to the Hellinger distance and Kullback-Leibler divergence are derived in Section 4. For practical purposes, the proposed method is described in terms of an algorithm in Section 5. Possible applications we have in mind are examined in Section 6. In particular, the setting of Bayesian inverse problems is recalled. Moreover, the computation of moments and marginals is scrutinized. Section 7 illustrates the performance of the proposed method. In addition to an examination of the numerical sensitivity of the accuracy with respect to the perturbation of the transport maps, a typical model problem from Uncertainty Quantification (UQ) is depicted, namely the identification of a parametrization for the random Darcy equation with lognormal coefficient given as solution of a stochastic Galerkin FEM.

## 2. DENSITY REPRESENTATION

The aim of this section is to introduce the central ideas of the proposed approximation of densities. For this task, two established concepts are reviewed, namely *transport maps* [22, 6], which are closely related to the notion of optimal transport [56, 49], and *hierarchical low-rank tensor representations* [41, 30, 3]. By the combination of these techniques, assuming the access to a suitable transformation, the developed approach yields a functional representation of the density in a format which is suited to computations with high-dimensional functions. In particular, we are able to handle highly concentrated posterior densities, e.g. appearing in the context of Bayesian inverse problems. While transport maps on their own in principle enable the generation of samples of some target distribution, the combination with a functional low-rank representation allows for integral quantities such as (centered) moments to become computable. Given an approximate transport map, the low-rank representation can be seen

as a further approximation step (improving the inaccuracy of the used transport) to gain direct access to the target density.

Consider a target measure  $\pi$  with Radon-Nikodym derivative with respect to the Lebesgue measure  $\lambda$  denoted as  $f$  with support in  $\mathbb{R}^d$ ,  $d < \infty$ , i.e.

$$(1) \quad f(y) := \frac{d\pi}{d\lambda}(y), \quad y \in Y := \mathbb{R}^d.$$

In the following we assume that point evaluations of  $f$  are available up to a multiplicative constant, motivated by the framework of Bayesian posterior density representation with unknown normalization constant. Furthermore, let  $\pi_0$  be some reference measure exhibiting a Radon-Nikodym derivative with respect to the Lebesgue measure denoted as  $f_0$ . This is motivated by the prior measure and density in the context of Bayesian inference.

**2.1. Transport Maps.** The notion of density transport is classical and with optimal transport has become a popular field recently, see e.g. [56, 49]. It has been employed to improve numerical approaches for Bayesian inverse problems for instance in [22, 6, 14]. Similar approaches are discussed in terms of sample transport e.g. for Stein's method [35, 11] or multi-layer maps [6]. We review the properties required for our approach in what follows. Note that since our target application is Bayesian inversion, we usually use the terms ‘‘prior’’ and ‘‘posterior’’ instead of the more general ‘‘reference’’ and ‘‘target’’ densities.

Let  $X := \mathbb{R}^d$  and assume that there exists an exact transport map

$$(2) \quad T: X \rightarrow Y,$$

which is a diffeomorphism<sup>1</sup> that relates  $\pi$  and  $\pi_0$  via pullback, i.e.

$$(3) \quad f_0(x) = f(T(x)) |\det \mathcal{J}_T(x)|, \quad x \in X.$$

Then, computations might be carried out in terms of the measure  $\pi_0$ , which is commonly assumed to be of a simpler structure. For instance the moment computation with respect to some multiindex  $\alpha$  reads as follows,

$$(4) \quad \int_Y y^\alpha d\pi(y) = \int_X T(x)^\alpha d\pi_0(x) \int_X T(x)^\alpha f_0(x) d\lambda(x).$$

Note that the computation of the right-hand side in (4) may still be a challenging task depending on the actual structure of  $T$ . In [22]  $T$  is expanded in chaos polynomials with respect to  $\pi_0$ . From a practical point of view, this provides access to lower-order moments using orthogonality of the underlying polynomial system.

Here we follow an alternative strategy with the aim to efficiently compute moments of some target density based on a functional representation. Notably we assume a convenient (simple) structures of  $T$  with the potential drawback of reduced accuracy, i.e. an inexact (pull-back) transport from the target to an auxiliary density (instead of the exact reference). Motivated by the Bayesian context, we call such a pull-back of some posterior density the *perturbed prior density*, see Section 2.2. Given a simple transport structure, the possibly demanding computational task is shifted to the accurate approximation of the perturbed prior. For this, there is justified hope of feasibility in some appropriate (alternative) coordinate system. In order to tackle moment computations, other posterior statistics or to generate posterior samples, we hence devise a numerical approach that enables a *workload balancing between the reconstruction of some problem-dependent transport structure and the accurate evaluation of the perturbed prior*. In the following we list some examples of transport maps.

<sup>1</sup>note that the requirements on  $T$  can be weakened, e.g. to local Lipschitz

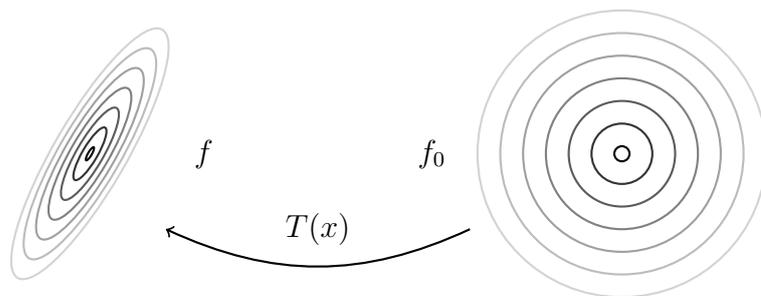


FIGURE 1. Illustration of affine transport: translation, rotation and rescaling.

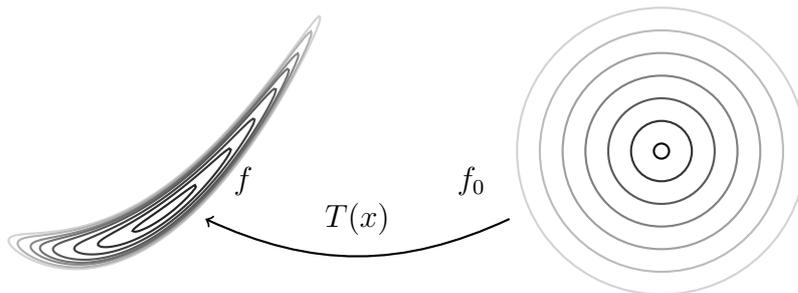


FIGURE 2. Illustration of quadratic transport: affine properties and bending.

2.1.1. *Affine transport.* In [50, 51] the authors employ an affine linear preconditioning for acceleration of MCMC or sparse-grid integration in the context of highly informative and concentrated Bayesian posterior densities, using a s.p.d. matrix  $H \in \mathbb{R}^{d,d}$  and  $M \in \mathbb{R}^d$ . In the mentioned articles, up to a multiplicative constant,  $H$  corresponds to the inverse square root of the Hessian at the MAP (maximum a posteriori probability)  $M$ , i.e. the location of the local optimum of a Laplace approximation of the posterior density. This rather simple construction, under the assumption of an unimodal density, leads to stable numerical algorithms for the computation of quantities of interest as the posterior mass concentrates. When considering the push-forward of a reference density  $f_0$  to a target density  $f$  this concept coincides with an affine transport

$$(5) \quad y = T(x) = Hx + M, \quad x \in X.$$

In the transport setting  $H$  and  $M$  may be computed for instance via some minimization of the Kullback-Leibler divergence as in [22]. Note that  $H$  and  $M$  do not necessarily have to be the inverse square root of the Hessian or the MAP. Figure 1 illustrates the concept of an affine transport.

2.1.2. *Quadratic transport.* A more general class of polynomial transport exhibits the form

$$(6) \quad T(x) = \frac{1}{2}x : \mathbf{A} : x + Hx + M, \quad x \in X,$$

with  $\mathbf{A} \in \mathbb{R}^{d,d,d}$ ,  $H \in \mathbb{R}^{d,d}$ ,  $M \in \mathbb{R}^d$ . Such a quadratic transport may be used for simple nonlinear transformations as depicted in Figure 2.

2.1.3. *More general transport maps.* The parametrization of transport maps can be chosen quite liberally as long as certain criteria are satisfied, which are either directly imposed in the ansatz space  $\mathcal{T}$  of the maps or added as constraints during optimization. In particular, the approximate transport map has to be invertible, which can be ensured by requiring a positive Jacobian. A commonly used measure

for transport optimization is the Kullback-Leibler divergence<sup>2</sup> leading to the optimization problem

$$(7) \quad \min_{T \in \mathcal{T}} d_{\text{KL}}(Y; T\pi_0, \pi) \quad \text{such that} \quad \det \nabla T > 0 \quad \pi\text{-a.e.}$$

Several suggestions regarding simplifications and special choices of function spaces  $\mathcal{T}$  such as smooth triangular maps based on higher-order polynomials or radial basis functions can for instance be found in the review article [22]. An interesting idea is to subdivide the task into the iterative computation of simple correction maps which are then composed as proposed in [6]. We again emphasize that while an accurate transport map is desirable, any approximation of such a map can in principle be used with the proposed method. In fact one can decide whether it is beneficial to spend more effort on the approximation of the perturbed density or on a better representation of the transport.

**2.2. Inexact transport and the perturbed prior.** In general, the transport map  $T$  is unknown or difficult to determine and hence has to be approximated by some  $\tilde{T}: X \rightarrow Y$ , e.g. using a polynomial chaos representation with respect to  $\pi_0$  [22] or with a more advanced composition of simple maps in a reduced space such as in [6]. As a consequence, it holds

$$(8) \quad \int_Y y^\alpha d\pi(y) \approx \int_X \tilde{T}(x)^\alpha d\pi_0(x)$$

subject to the accuracy of the involved approximation of  $T$ . One can also view  $\tilde{T}$  as the push-forward of some measure  $\tilde{\pi}_0$  with density  $\tilde{f}_0$  to  $\pi$  given by

$$(9) \quad \tilde{f}_0(x) = f(\tilde{T}(x)) |\det \mathcal{J}_{\tilde{T}}(x)|.$$

We henceforth refer to (9) as the auxiliary reference or *perturbed prior density*. Using this construction, the moment computation reads

$$(10) \quad \int_Y y^\alpha d\pi(y) = \int_X \tilde{T}(x)^\alpha d\tilde{\pi}_0 = \int_X \tilde{T}(x)^\alpha \tilde{f}_0(x) d\lambda(x).$$

If one would know  $\tilde{f}_0$ , by (9) and (10) one would also have access to the exact posterior.

Equation (10) is the starting point of the proposed method by approximating  $\tilde{f}_0$  in another coordinate system which is better adapted to the structure of the approximate (perturbed) prior. Consider a (fixed) diffeomorphism

$$(11) \quad \Phi: \hat{X} \subset \mathbb{R}^d \rightarrow X, \quad \hat{x} \mapsto x = \Phi(\hat{x})$$

with Jacobian  $\hat{x} \mapsto |\det \mathcal{J}_\Phi(\hat{x})|$  and define the *perturbed transformed prior*

$$(12) \quad \hat{f}_0: \hat{X} \mapsto \mathbb{R}_+, \quad \hat{x} \mapsto \hat{f}_0(\hat{x}) := \tilde{f}_0(\Phi(\hat{x})).$$

In case (12) can be approximated accurately by some function  $\hat{f}_0^h$  then

$$(13) \quad \int_Y y^\alpha d\pi(y) \approx \int_{\hat{X}} \tilde{T}(\Phi(\hat{x}))^\alpha \hat{f}_0^h(\hat{x}) |\det \mathcal{J}_\Phi(\hat{x})| d\lambda(\hat{x})$$

with accuracy determined only by the approximation quality of  $\hat{f}_0^h$ . Thus, (12) and (13) enable a balancing between the construction of the transport map approximation  $\tilde{T}$  of  $T$  to shift its complexity given the underlying diffeomorphism  $\Phi$  to the approximation of (12) in a new coordinate system intrinsic to  $\hat{X}$ .

The construction of  $\tilde{T}$  and a suitable map in (11) may be used to obtain a convenient transformed auxiliary reference density given in (12). An approximation thereof can be significantly simpler compared

<sup>2</sup>although in machine learning Wasserstein or Sinkhorn distances have become very popular when so-called normalizing flows are computed

to a possibly complicated and concentrated target density  $f$  or the computation of the exact transport  $T$ . This e.g. is satisfied if

- $f_0$  is a Gaussian density and  $\tilde{T}$  maps  $f$  to  $\tilde{f}_0$  which is in some sense near to a Gaussian density. In this case,  $\Phi$  from (11) may be chosen as the  $d$ -dimensional spherical transformation and extended low-rank tensor formats are employed for the approximation, see Section 3. In this setting, the introduction of an adapted coordinate system allows to shift the exponential decay to the one dimensional radial parameter. The accuracy of an approximation can then be improved easily by additional  $h$ -refinements as described in Section 2.3.
- The reference density  $f_0$  has a complicated form and might be replaced by  $\tilde{f}_0$  to become computationally accessible.

In the following we state an important property that needs to be fulfilled by the perturbed prior  $\tilde{f}_0$  in order to lead to a convergent method with the employed approximations.

**Definition 2.1. (outer polynomial exponential decay)** A function  $\tilde{f}_0: X \rightarrow \mathbb{R}^+$  has outer polynomial exponential decay if there exists a simply connected compact  $K \subset X$  with a polynomial  $\pi^+$  being positive on  $X \setminus K$  and some  $C > 0$  such that

$$(14) \quad \tilde{f}_0(x) \leq C \exp(-\pi^+(x)), \quad x \in X \setminus K.$$

**2.3. Layer based representation.** To further refine and motivate the notion of an adapted coordinate system, let  $L \in \mathbb{N}$  and  $(X^\ell)_{\ell=1}^L$  be pairwise disjoint domains in  $X$  s.t.

$$(15) \quad K := \bigcup_{\ell=1}^L \overline{X^\ell}$$

is simply connected and compact and define  $X^{L+1} := X \setminus K$ . Then, for given  $L \in \mathbb{N}$  we may decompose the perturbed prior  $\tilde{f}_0$  as

$$(16) \quad \tilde{f}_0(x) = \sum_{\ell=1}^{L+1} \tilde{f}_0^\ell(x) \quad \text{with} \quad \tilde{f}_0^\ell := \chi_\ell \tilde{f}_0,$$

where  $\chi_\ell$  denotes the indicator function on  $X^\ell$ . Moreover, for any tensor set  $\hat{X}^\ell := \times_{i=1}^d \hat{X}_i^\ell$  and diffeomorphism  $\Phi^\ell: \hat{X}^\ell \mapsto X^\ell$ ,  $1 \leq \ell \leq L+1$ , we may represent the *localized perturbed prior*  $\tilde{f}_0^\ell$  as a pull-back function

$$(17) \quad \tilde{f}_0^\ell = \hat{f}_0^\ell \circ \Phi^{\ell-1},$$

where  $\hat{f}_0^\ell$  is a map defined on  $\hat{X}^\ell$  as in (12). We consider the following example.

**Example 2.2. (multivariate polar transformation)**

The  $d$ -dimensional spherical coordinate system allows for simple layer layouts in terms of hyperspherical shells. In particular, for  $\ell = 1, \dots, L+1 < \infty$ , with  $0 = \rho_1 < \rho_2 < \dots < \rho_{L+1} < \rho_{L+2} = \infty$ , let

$$\begin{aligned} \hat{X}^\ell &:= [\rho_\ell, \rho_{\ell+1}] \times [0, 2\pi] \times \times_{i=2}^{d-2} [0, \pi], \\ X^\ell &:= B_{\rho_{\ell+1}}(0) \setminus B_{\rho_\ell}(0) \subset X, \end{aligned}$$

i.e.  $\hat{X}^\ell$  and  $X^\ell$  denote the corresponding adopted (transformed) and the original parameter space, respectively. Then, for  $\hat{x} = (\rho, \theta_0, \boldsymbol{\theta}) \in \hat{X}$ ,  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{d-2})$ , the polar transformation  $\Phi^\ell: \hat{X}^\ell \rightarrow$

$X^\ell$  reads

$$(18) \quad \Phi^\ell(\hat{x}) = \Phi^\ell(\rho, \theta_0, \boldsymbol{\theta}) = \rho \begin{bmatrix} \cos \theta_0 \sin \theta_1 \sin \theta_2 \cdots \sin \theta_{d-3} \sin \theta_{d-2} \\ \sin \theta_0 \sin \theta_1 \sin \theta_2 \cdots \sin \theta_{d-3} \sin \theta_{d-2} \\ \cos \theta_1 \sin \theta_2 \cdots \sin \theta_{d-3} \sin \theta_{d-2} \\ \cos \theta_2 \cdots \sin \theta_{d-3} \sin \theta_{d-2} \\ \vdots \\ \cos \theta_{d-3} \sin \theta_{d-2} \\ \cos \theta_{d-2} \end{bmatrix}.$$

Moreover, the Jacobian is given by

$$(19) \quad \det \mathcal{J}_{\Phi^\ell}(\rho, \theta_0, \boldsymbol{\theta}) = \rho^{d-1} \prod_{i=1}^{d-2} \sin^i \theta_i.$$

This layer based coordinate change enables a representation of the density on bounded domains. Even though the remainder layer is unbounded, we assume that  $K$  is sufficiently large to cover all probability mass of  $\tilde{f}_0$  except for a negligible higher-order error.

Up to this point, the choice of transformation  $\Phi^\ell$ ,  $\ell = 1, \dots, L + 1$ , is fairly general. However, for the further development of the method we assume the following property.

**Definition 2.3. (rank 1 stability)**

Let  $\mathcal{X}, \hat{\mathcal{X}} = \times_{i=1}^d \hat{\mathcal{X}}_i \subset \mathbb{R}^d$  be open and bounded sets. A diffeomorphism  $\Phi: \hat{\mathcal{X}} \mapsto \mathcal{X}$  is called rank 1 stable if  $\Phi$  and the absolute value of its Jacobian  $\det \mathcal{J}_\Phi$  have rank 1, i.e. there exists univariate functions  $\Phi_i: \hat{\mathcal{X}}_i \rightarrow \mathcal{X}$ ,  $h_i: \hat{\mathcal{X}} \rightarrow \mathbb{R}$ ,  $i = 1, \dots, d$ , such that for  $\hat{x} \in \hat{\mathcal{X}}$

$$(20) \quad \Phi(\hat{x}) = \prod_{i=1}^d \Phi_i(\hat{x}_i), \quad |\det \mathcal{J}_\Phi(\hat{x})| = \prod_{i=1}^d h_i(\hat{x}_i).$$

**Proposition 2.4.** *The multivariate polar coordinate transformation from Example 2.2 is rank 1 stable.*

Due to the notion of rank 1 stable transformations, the map  $\hat{x} \mapsto T(\Phi(\hat{x}))$  in (13) inherits the rank structure of  $T$ , see Section 3. Furthermore, since the Jacobian  $\hat{x} \mapsto |\det \mathcal{J}_\Phi(\hat{x})|$  is rank 1, we can construct tensorized orthonormal basis functions which may be used to approximate the perturbed transformed prior in (12).

**Remark 2.5.** *The described concept can be extended to any rank  $r \in \mathbb{N}$  Jacobians of  $\Phi$ , i.e.*

$$(21) \quad |\det \mathcal{J}_\Phi(\hat{x})| = \sum_{k=1}^r \prod_{i=1}^d h_{i,k}(\hat{x}_i).$$

Motivated by the right-hand side in (13), one may use different approximations of the perturbed transformed prior  $\tilde{f}_0 \circ \Phi$  in  $r$  distinct tensorized spaces, each associated to the rank 1 weight  $\prod_{i=1}^d h_{i,k}$ .

**2.4. Layer truncation.** This paragraph is devoted to the treatment of the last (remainder or truncation) layer introduced in (16) with the aim to suggest some approximation choices.

If  $\tilde{f}_0$  is represented in the layer format (16), it is convenient to simply extend the function to zero after layer  $L \in \mathbb{N}$ . By this, the remaining (possibly small) probability mass is neglected. Such a procedure is typically employed in numerical applications and does not impose any computational issues since events on the outer truncated domain are usually exponentially unlikely for truncation value chosen sufficiently large. Nevertheless, in order to present a rigorous treatment, we require properties like absolute continuity, which would be lost by using a cut-off function. Inspired by [51] regarding the

information limit of unimodal posterior densities<sup>3</sup>, we suggest a Gaussian approximation for the last layer  $L + 1$  on the unbounded domain  $X^{L+1}$ , i.e. for some s.p.d.  $\Sigma \in \mathbb{R}^{d,d}$  and  $\mu \in \mathbb{R}^d$  we define the *hybrid representation of the perturbed prior* by

$$(22) \quad \tilde{f}_0^{\text{Trun}}(x) := C_L \begin{cases} \tilde{f}_0^\ell(x), & x \in X^\ell, \ell = 1, \dots, L, \\ f_{\Sigma, \mu}(x), & x \in X^{L+1}, \end{cases}$$

with  $C_L = (C_L^< + C_L^>)^{-1}$ , where

$$(23) \quad C_L^< := \int_{X \setminus K} f_{\Sigma, \mu}(x) \, d\lambda(x),$$

$$(24) \quad C_L^> := \sum_{\ell=1}^L \int_{X^\ell} \tilde{f}_0^\ell(x) \, d\lambda(x),$$

and  $f_{\Sigma, \mu}$  denotes the Gaussian probability density function with mean  $\mu$  and covariance matrix  $\Sigma$ .

**Remark 2.6.** A good choice for  $\mu$  and  $\Sigma$  would be the mean and covariance of the exact perturbed prior  $\tilde{f}_0$ , which however is not accessible a priori. Thus, in numerical simulations one may choose  $\mu$  and  $\Sigma$  as (centralized) moments of the normalized truncated perturbed prior density  $\tilde{f}_0^{\text{Trun}}|_K$  or as the MAP point and the corresponding square root of the numerically computed Hessian as a result of an optimization algorithm on  $\tilde{f}_0$ .

**Lemma 2.7. (truncation error)** For  $\mu \in \mathbb{R}^d$  and  $\Sigma \in \mathbb{R}^{d,d}$  let  $\tilde{f}_0$  have outer polynomial exponential decay with positive polynomial  $\tilde{\pi}^+$  and  $\tilde{C} > 0$  with  $K = \overline{B_R(\mu)}$  for some  $R > 0$ . Then, for  $C_\Sigma = 1/2\lambda_{\min}(\Sigma^{-1})$  there exists  $C = C(\tilde{C}, \Sigma, d, C_\Sigma) > 0$  such that

$$\|\tilde{f}_0 - \tilde{f}_0^{\text{Trun}}\|_{L^1(X \setminus K)} \lesssim \|\exp(-\tilde{\pi}^+)\|_{L^1(X \setminus K)} + \Gamma(d/2, C_\Sigma R^2)$$

and

$$\left| \int_{X \setminus K} \log \left( \frac{\tilde{f}_0}{f_{\Sigma, \mu}} \right) \tilde{f}_0 \, dx \right| \leq \int_{X \setminus K} \left( \frac{1}{2} \|x\|_{\Sigma^{-1}}^2 + \tilde{\pi}^+(x) \right) e^{-\tilde{\pi}^+(x)} \, d\lambda(x)$$

with the incomplete Gamma function  $\Gamma$ .

*Proof.* The proof follows immediately from the definition of  $\tilde{f}_0^{\text{Trun}}$ .  $\square$

In the case that the perturbed prior is close to a Gaussian standard normal distribution, it holds  $c \approx 1$ .

Note that the constant  $C_L^<$  in (23) may exhibit an analytic form whereas computing  $C_L^>$  suffers from the curse of dimensionality and is in general not available. To circumvent this issue and render further use of the representation (22) feasible, we introduce a suitable low-rank approximation in the next section.

### 3. LOW-RANK TENSOR TRAIN FORMAT

The computation of high-dimensional integrals and the efficient construction of surrogates is a challenging task with a multitude of approaches. Some of these techniques are sparse grid methods [7, 26], collocation [23, 39, 25] or modern sampling techniques [27, 48, 38]. As motivated by  $C_L^>$  in (24), we aim for a model to adequately approximate the localized perturbed prior maps  $\tilde{f}_0^\ell$ . The introduction of an adapted coordinate system enables the use of low-rank representations such as the tensor train

<sup>3</sup>A result of [51] is that under suitable conditions the posterior distribution converges to a Gaussian in the limit of zero noise and infinite measurements.

(TT) format [41, 31, 30] described in this section. We highlight a “non-intrusive” sample-based technique to obtain such a representation of arbitrary maps, namely the Variational Monte Carlo (VMC) method [21].

Let  $\hat{X} = \bigotimes_{i=1}^d \hat{X}_i$  be a tensor space of separable Banach spaces  $\hat{X}_i, i \in [d] := \{1, \dots, d\}$ , and consider a map  $g: \hat{X} \rightarrow \mathbb{R}$ . The function  $g$  can be represented in the TT format if there exists a *rank vector*  $\mathbf{r} = (r_1, \dots, r_{d-1}) \in \mathbb{N}^{d-1}$  and univariate functions  $g^i[k_{i-1}, k_i]: \hat{X}_i \rightarrow \mathbb{R}$  for  $k_i \in [r_i], i \in [d]$ , such that for all  $\hat{x} \in \hat{X}$

$$(25) \quad g(\hat{x}) = \sum_{\mathbf{k}=1}^{\mathbf{r}} \prod_{i=1}^d g^i[k_{i-1}, k_i](\hat{x}_i), \quad \mathbf{k} := (k_1, \dots, k_{d-1}).$$

For ease of notation it is convenient to set  $k_0 = k_d = 1$ . In the forthcoming sections we consider weighted tensorized Lebesgue spaces. In particular, for a non-negative weight function  $w: \hat{X} \rightarrow \mathbb{R}$  with  $w = \bigotimes_{i=1}^d w_i, w \in L^1(\hat{X})$ , define the tensorization of  $L^2(\hat{X}, w) = \bigotimes_{i=1}^d L^2(\hat{X}_i, w_i)$  by

$$(26) \quad \mathcal{V}(\hat{X}) := L^2(\hat{X}, w) = \left\{ v: \hat{X} \rightarrow \mathbb{R} \mid \|v\|_{\mathcal{V}}^2 := \int_{\hat{X}} v(\hat{x})^2 w(\hat{x}) d\lambda(\hat{x}) < \infty \right\}.$$

We assume that there exists an complete orthonormal basis  $\{P_k^i : k \in \mathbb{N}\}$  in  $L^2(\hat{X}_i, w_i)$  for every  $i \in [d]$  which is known a priori. For discretization purposes, we introduce the finite dimensional subspaces

$$(27) \quad \mathcal{V}_{i, n_i} := \overline{\text{span} \{P_1^i, \dots, P_{n_i}^i\}} \subseteq L^2(\hat{X}_i, w_i)$$

for  $i = 1, \dots, d$ , and  $n_i \in \mathbb{N}$ . On these we formulate the *extended tensor train format* in terms of the coefficient tensors

$$(28) \quad G^i: [r_{i-1}] \times [n_i] \times [r_i] \rightarrow \mathbb{R}, \quad (k_{i-1}, j, k_i) \quad \mapsto G^i[k_{i-1}, j, k_i], \quad i \in [d],$$

such that every univariate function  $g^i \in \mathcal{V}_{i, n_i}$  can be written as

$$(29) \quad g^i[k_{i-1}, k_i](\hat{x}_i) = \sum_{j=1}^{n_i} G^i[k_{i-1}, j, k_i] P_j^i(\hat{x}_i) \quad \text{for } \hat{x} \in \hat{X}_i.$$

For the full tensor format the function

$$(30) \quad g \in \mathcal{V}_\Lambda := \bigotimes_{i=1}^d \mathcal{V}_{i, n_i} \subseteq \mathcal{V}(\hat{X})$$

can be expressed by a high dimensional algebraic tensor  $G: \Lambda := \times_{i=1}^d [n_i] \rightarrow \mathbb{R}$  and tensorized functions  $P_\alpha := \bigotimes_{i=1}^d P_{\alpha_i}$  for  $\alpha = (\alpha_1, \dots, \alpha_d) \in \Lambda$  such that

$$(31) \quad g(\hat{x}) = \sum_{\alpha \in \Lambda} G[\alpha_1, \dots, \alpha_d] \prod_{i=1}^d P_{\alpha_i}(\hat{x}_i).$$

In contrast to this, the format given by (25) and (29) admits a linear structure in the dimension. More precisely, the memory complexity of  $\mathcal{O}(\max\{n_1, \dots, n_d\}^d)$  in (31) reduces to

$$(32) \quad \mathcal{O}(\max\{r_1, \dots, r_{d-1}\}^2 \cdot d \cdot \max\{n_1, \dots, n_d\}).$$

This observation raises the question of expressibility for certain classes of functions and the existence of a low-rank vector  $\mathbf{r}$  where  $\max\{r_1, \dots, r_{d-1}\}$  is sufficiently small for practical computations. This issue is e.g. addressed in [52, 2, 29] under certain assumptions on the regularity and in [24, 41, 4, 21] explicit (algorithmic) constructions of the format are discussed even in case that  $g$  has no analytic representation.

For later reference we define the finite dimensional low-rank manifold of rank  $r$  tensor trains by

$$(33) \quad \mathcal{M}_r := \{g \in \mathcal{V}(\hat{X}) \mid g \text{ as in (25) with } g^i \text{ as in (29)}\}.$$

This is an embedded manifold in the finite full tensor space  $\mathcal{V}_\Lambda$  from (30) admitting the cone property. We also require the concept of the algebraic (full) tensor space

$$(34) \quad \mathbb{T} := \{G: \mathbb{N}^d \rightarrow \mathbb{R}\}$$

and the corresponding low-rank form for given  $r \in \mathbb{N}^{d-1}$  defined by

$$(35) \quad \mathbb{TT}_r := \left\{ G: \Lambda \rightarrow \mathbb{R} \mid G[\alpha] = \sum_{k=1}^r \prod_{i=1}^d G[k_{i-1}, \alpha_i, k_i] \right\}.$$

Without going into detail, we mention the higher order singular value decomposition (HOSVD), which is used to decompose a full algebraic tensor into a low-rank tensor train. The algorithm is based on successive unfoldings of the full tensor into matrices, which are orthogonalized and possibly truncated by a singular value decomposition, see [40] for details. This algorithm enables us to state the following Lemma.

**Lemma 3.1** ([40, Theorem 2.2]). *For any  $g \in \mathcal{V}_\Lambda$  and  $r \in \mathbb{R}^{d-1}$  there exists an extended low-rank tensor train  $g_r \in \mathcal{M}_r$  with*

$$(36) \quad \|g - g_r\|_{\mathcal{V}(\hat{X})}^2 \leq \sum_{i=1}^{d-1} \sigma_i^2,$$

where  $\sigma_i$  is the distance of the  $i$ -th unfolding matrix of the coefficient tensor of  $g$  in the HOSVD to its best rank  $r_i$  approximation in the Frobenius norm.

*Proof.* The proof follows from the best approximation result of the usual matrix SVD with respect to the Frobenius norm and the orthonormality of the chosen basis.  $\square$

**Remark 3.2.** *Estimate (36) is rather unspecific as the  $\sigma_i$  cannot be quantified a priori. In the special case of Gaussian densities we refer to [47] for an examination of the low-rank representation depending on the covariance structure. By considering a transport  $\tilde{T}$  that maps the considered density only “close” to a standard Gaussian, the results can be applied immediately to our setting and more precise estimates are possible.*

**3.1. Tensor train regression by Variational Monte Carlo.** We review the sampling-based VMC method presented in [21] which is employed to construct TT representations of the local maps  $\Phi^\ell$  as in (17). The approach generalizes the concept of randomized tensor completion [19] and its analysis relies on the theory of statistical learning, leading to a priori convergence results. It can also be seen as a generalized tensor least squares technique. An alternative cross-interpolation method for probability densities is presented in [14].

For the VMC framework, consider the *model class*  $\mathcal{M}_r(\underline{c}, \bar{c}) \subset \mathcal{M}_r$  of truncated rank  $r \in \mathbb{R}^{d-1}$  tensor trains which is given for  $0 \leq \underline{c} < \bar{c} \leq \infty$  by

$$(37) \quad \mathcal{M}_r(\underline{c}, \bar{c}) := \left\{ g \in \mathcal{M}_r \mid \underline{c} \leq g(\hat{x}) \leq \bar{c} \text{ a.e. in } \hat{X} \right\}.$$

The model class  $\mathcal{M}_r(\underline{c}, \bar{c})$  is a finite subset of the truncated nonlinear space  $\mathcal{V}(\hat{X}, \underline{c}, \bar{c}) \subseteq \mathcal{V}(\hat{X})$  defined as

$$(38) \quad \mathcal{V}(\hat{X}, \underline{c}, \bar{c}) := \{v \in L^2(\hat{X}, w) \mid \underline{c} \leq v(\hat{x}) \leq \bar{c} \text{ a.e. in } \hat{X}\},$$

which we equip with the metric  $d_{\mathcal{V}(\hat{X}, \underline{c}, \bar{c})}(v, w) := \|v - w\|_{\mathcal{V}}$ .

Alternatively, for numerical purposes we may characterize  $\mathcal{M}_r(\underline{c}, \bar{c})$  and  $\mathcal{V}(\hat{X}, \underline{c}, \bar{c})$  in terms of constraints on the coefficients of the underlying representation with respect to  $\{P_\alpha\}$ . For  $\ell^2(\mathbb{T}) := \{G \in \mathbb{T} \mid \sum_{\alpha \in \mathbb{N}^d} G[\alpha]^2 < \infty\}$  we have

$$(39) \quad \mathcal{V}(\hat{X}, \underline{c}, \bar{c}) = \{v(\hat{x}) = \sum_{\alpha \in \mathbb{N}^d} G[\alpha] \cdot P_\alpha(\hat{x}) \mid G \in \ell^2(\mathbb{T}), \underline{F}(G) \geq 0, \bar{F}(G) \leq 0\},$$

$$(40) \quad \mathcal{M}(\underline{c}, \bar{c}) = \{v(\hat{x}) = \sum_{\alpha \in \Lambda} G[\alpha] \cdot P_\alpha(\hat{x}) \mid G \in \mathbb{T}\mathbb{T}_r, \underline{F}^r(G) \geq 0, \bar{F}^r(G) \leq 0\},$$

for constraint functions  $\underline{F}, \bar{F}: \ell^2(\mathbb{T}) \rightarrow \mathbb{R}$  and

$\underline{F}^r, \bar{F}^r: \ell^2(\mathbb{T}\mathbb{T}_r) \rightarrow \mathbb{R}$  implicitly bounding the coefficient tensors. Note that due to the orthonormality of  $\{P_\alpha\}_{\alpha \in \mathbb{N}^d}$  in  $\mathcal{V}(\hat{X})$  for every  $v \in \mathcal{V}(\hat{X})$  it holds

$$(41) \quad \|v\|_{\mathcal{V}} = \|G\|_{\ell^2(\mathbb{T})} \quad \text{with} \quad v = \sum_{\alpha} G[\alpha] P_\alpha \in \mathcal{V}.$$

Additionally, we define a *loss function*  $\iota: \mathcal{V}(\hat{X}, \underline{c}, \bar{c}) \times \hat{X} \rightarrow \mathbb{R}$  such that  $\iota(\cdot, \hat{x})$  is continuous for almost all  $\hat{x} \in \hat{X}$  and  $\iota(v, \cdot)$  is integrable with respect to the weight function  $w$  of  $\mathcal{V}(\hat{X})$  for every  $v \in \mathcal{V}(\hat{X}, \underline{c}, \bar{c})$ . Then, we consider the *cost functional*  $\mathcal{J}: \mathcal{V}(\hat{X}, \underline{c}, \bar{c}) \rightarrow \mathbb{R}$  given by

$$(42) \quad \mathcal{J}(v) := \int_{\hat{X}} \iota(v, \hat{x}) w(\hat{x}) d\lambda(\hat{x}).$$

To further analyze the approximability in the given TT format using sampling techniques, we define two common discrepancy measures for probability density functions.

**Lemma 3.3.** (*KL loss compatibility*) Let  $h^* \in \mathcal{V}(\hat{X}, 0, c^*)$  for  $c^* < \infty$  and  $0 < \underline{c} < \bar{c} < \infty$ . Then

$$(43) \quad \mathcal{V}(\hat{X}, \underline{c}, \bar{c}) \ni g \mapsto \iota(g, \hat{x}) = \iota(g, \hat{x}, h^*) := -\log(g(\hat{x})) h^*(\hat{x})$$

is uniformly bounded and Lipschitz continuous on

$\mathcal{M}_r(\underline{c}, \bar{c})$  if  $P_\alpha \in L^\infty(\hat{X})$  for every  $\alpha \in \Lambda$ . Furthermore,  $\mathcal{J}$  is globally Lipschitz continuous on the metric space  $(\mathcal{V}(\hat{X}, \underline{c}, \bar{c}), d_{\mathcal{V}(\hat{X}, \underline{c}, \bar{c})})$ .

*Proof.* The loss  $\iota$  is bounded on  $\mathcal{M}_r(\underline{c}, \bar{c})$  since  $0 < \underline{c} < \bar{c} < \infty$ . Let  $g_1, g_2 \in \mathcal{V}_r(\hat{X}, \underline{c}, \bar{c})$  with coefficient tensors  $G_1$  and  $G_2 \in \mathbb{T}\mathbb{T}_r$ , then

$$(44) \quad |\iota(g_1, \hat{x}) - \iota(g_2, \hat{x})| \leq \underbrace{\frac{1}{\underline{c}} \sup_{\hat{x} \in \hat{X}} \{h^*(\hat{x})\}}_{:= C^* < \infty} |g_1(\hat{x}) - g_2(\hat{x})|.$$

The global Lipschitz continuity of  $\mathcal{J}$  follows by using (44) and

$$(45) \quad |\mathcal{J}(g_1) - \mathcal{J}(g_2)| \leq C^* \|g_1 - g_2\|_{L^1(\hat{X}, w)} \leq CC^* d_{\mathcal{V}(\hat{X}, \underline{c}, \bar{c})}(g_1, g_2),$$

with a constant  $C$  related to the embedding of  $L^2(\hat{X}, w)$  into  $L^1(\hat{X}, w)$ . If  $g_1, g_2$  are in  $\mathcal{M}_r(\underline{c}, \bar{c})$  then by Parseval's identity and the finite dimensionality of  $\mathcal{M}_r(\underline{c}, \bar{c})$  there exists  $c = c\left(\sup_{\alpha \in \Lambda} \|P_\alpha\|_{L^\infty(\hat{X})}\right) > 0$  such that

$$(46) \quad |g_1(x) - g_2(x)| \leq c \|G_1 - G_2\|_{\ell^2(\mathbb{T})} = c \|g_1 - g_2\|_{\mathcal{V}} = c d_{\mathcal{V}(\hat{X}, \underline{c}, \bar{c})}(g_1, g_2),$$

which yields the Lipschitz continuity on  $\mathcal{M}_r(\underline{c}, \bar{c})$ . Now let  $g_1, g_2 \in \mathcal{V}(\hat{X}, \underline{c}, \bar{c})$ . The global Lipschitz continuity of  $\mathcal{J}$  follows by using (44) and

$$(47) \quad |\mathcal{J}(g_1) - \mathcal{J}(g_2)| \leq C^* \|g_1 - g_2\|_{L^1(\hat{X}, w)} \leq CC^* d_{\mathcal{V}(\hat{X}, \underline{c}, \bar{c})}(g_1, g_2),$$

with a constant  $C$  related to the embedding of  $L^2(\hat{X}, w)$  into  $L^1(\hat{X}, w)$ .  $\square$

**Lemma 3.4.** ( *$L^2$ -loss compatibility*) Let  $h^* \in \mathcal{V}(\hat{X}, 0, \bar{c})$  for  $\bar{c} < \infty$ . Then

$$(48) \quad \mathcal{V}(\hat{X}, 0, \bar{c}) \ni g \mapsto \iota(g, \hat{x}) = \iota(g, \hat{x}, h^*) := |g(\hat{x}) - h^*(\hat{x})|^2$$

is uniformly bounded and Lipschitz continuous on  $\mathcal{M}_r(0, \bar{c})$  provided  $P_\alpha \in L^\infty(\hat{X})$  for every  $\alpha \in \Lambda$ .

*Proof.* Let  $g_1, g_2 \in \mathcal{V}(\hat{X}, 0, \bar{c})$ . Then

$$(49) \quad |\iota(g_1, \hat{x}) - \iota(g_2, \hat{x})| \leq |g_1(\hat{x}) - g_2(\hat{x})| \cdot |g_2(\hat{x}) + g_2(\hat{x})| + 2|g_1(\hat{x}) - g_2(\hat{x})| h^*(\hat{x}).$$

Due to  $\bar{c} < \infty$  the Lipschitz property follows as in the proof of Lemma 3.3 if  $g_1, g_2$  in  $\mathcal{M}_r(\underline{c}, \bar{c})$ .  $\square$

To examine the VMC convergence in our setting, we recall the analysis of [21] in a slightly more general manner. The target objective of the method is to find a minimizer

$$(50) \quad v^* \in \operatorname{argmin}_{v \in \mathcal{V}(\hat{X}, \underline{c}, \bar{c})} \mathcal{J}(v).$$

Due to the infinite dimensional setting we confine the minimization problem in (50) to our model class  $\mathcal{M} = \mathcal{M}_r(\underline{c}, \bar{c})$ . This yields the minimization problem

$$(51) \quad \text{find } v_{\mathcal{M}}^* \in \operatorname{argmin}_{v \in \mathcal{M}} \mathcal{J}(v).$$

A crucial step is then to consider the empirical functional instead of the integral in  $\mathcal{J}$ , namely

$$(52) \quad \mathcal{J}_N(v) := \frac{1}{N} \sum_{k=1}^N \iota(v; \hat{x}^k),$$

with independent samples  $\{\hat{x}^k\}_{k \leq N}$  distributed according to the measure  $w\lambda$  with a (possibly rescaled) weight function  $w$  with respect to the Lebesgue measure  $\lambda$ . The corresponding empirical optimization problem then takes the form

$$(53) \quad \text{find } v_{\mathcal{M}, N}^* \in \operatorname{argmin}_{v \in \mathcal{M}} \mathcal{J}_N(v).$$

The analysis examines different errors with respect to  $h^* \in \mathcal{V}(\hat{X}, 0, \bar{c})$  defined by

$$(54) \quad \mathcal{E} := |\mathcal{J}(h^*) - \mathcal{J}(v_{\mathcal{M}, N}^*)|,$$

$$(55) \quad \mathcal{E}_{\text{app}} := |\mathcal{J}(h^*) - \mathcal{J}(v_{\mathcal{M}}^*)|,$$

$$(56) \quad \mathcal{E}_{\text{gen}} := |\mathcal{J}(v_{\mathcal{M}}^*) - \mathcal{J}(v_{\mathcal{M}, N}^*)|,$$

denoting the VMC-, approximation- and generalization error respectively. By a simple splitting, the VMC error can be bounded by the approximation and the generalization error, namely

$$(57) \quad \mathcal{E} \leq \mathcal{E}_{\text{app}} + \mathcal{E}_{\text{gen}}.$$

Due to the global Lipschitz property on  $\mathcal{V}(\hat{X}, \underline{c}, \bar{c})$  with  $\underline{c} > 0$  in the setting of (43) or  $\underline{c} \geq 0$  as in (48), the approximation error can be bounded by the best approximation in  $\mathcal{M}$ . In particular there exists  $C > 0$  such that

$$(58) \quad \mathcal{E}_{\text{app}} \leq C \inf_{v \in \mathcal{M}} \|h^* - v\|_{\mathcal{V}(\hat{X})}^2.$$

We note that such an estimation by the best approximation in  $\mathcal{M}$  with respect to the  $\mathcal{V}(\hat{X})$ -norm may not be required when using the Kullback-Leibler divergence if one is interested directly in the best approximation in this divergence. Then the assumption  $\underline{c} > 0$  can be relaxed in the construction of  $\mathcal{V}(\hat{X}, \underline{c}, \bar{c})$  since no global Lipschitz continuity of  $\mathcal{J}$  in Lemma 3.3 is required. Thus the more natural subspace of  $\mathcal{V}(\hat{X}, 0, \bar{c})$  of absolutely continuous functions with respect to  $h^*$  may be considered instead.

It remains to bound the statistical generalization error  $\mathcal{E}_{\text{gen}}$ . For this the notion of covering numbers is required. Let  $(\Omega, \mathcal{F}, \mathbb{P})$  be an abstract probability space.

**Definition 3.5. (covering number)** Let  $\epsilon > 0$ . The covering number  $\nu(\mathcal{M}, \epsilon)$  denotes the minimal number of open balls of radius  $\epsilon$  with respect to the metric  $d_{\nu(\hat{X}, \underline{c}, \bar{c})}$  needed to cover  $\mathcal{M}$ .

**Lemma 3.6.** Let  $\iota$  be defined as in (43) or (48). Then there exist  $C_1, C_2 > 0$  only depending on the uniform bound and the Lipschitz constant of  $\mathcal{M}$  given in Lemma 3.3 and 3.4, respectively, such that for  $\epsilon > 0$  and  $N \in \mathbb{N}$  denoting the number of samples in the empirical cost functional in (52) it holds

$$(59) \quad \mathbb{P}[\mathcal{E}_{\text{gen}} > \epsilon] \leq 2\nu(\mathcal{M}, C_2^{-1}\epsilon)\delta(1/4\epsilon, N),$$

with  $\delta(\epsilon, N) \leq 2 \exp(-2\epsilon^2 N/C_1^2)$ .

*Proof.* The claim follows immediately from Lemmas 3.3 and 3.4, respectively, and [21, Thm. 4.12, Cor. 4.19].  $\square$

**Remark 3.7** (choice of  $\underline{c}, \bar{c}$  and  $\hat{X}$ ). Due to the layer based representation in (16) and (22) on each layer  $\hat{X}^\ell = \Phi^{-1}(X^\ell)$  we have the freedom to choose  $\underline{c}$  separately. In particular, assuming that the perturbed prior  $\tilde{f}_0$  decays per layer, we can choose  $\underline{c}$  according to the decay and with this control the constant in (44).

#### 4. ERROR ESTIMATES

This section is devoted to the derivation of a priori error estimates for the previously introduced construction in terms of the Hellinger distance and Kullback-Leibler divergence. We employ the VMC approach from Section 3.1 to the density layer approximation which leads to a convergence result.

Recall that our goal is to approximate the perturbed prior  $\tilde{f}_0$  given some transport  $\tilde{T}$  represented by a function  $\tilde{f}_0^{\text{Trun,TT}}$  defined by

$$(60) \quad \tilde{f}_0^{\text{Trun,TT}}(x) := C_L^{\text{TT}} \begin{cases} \tilde{f}_0^{\ell, \text{TT}}(x), & x \in X^\ell, \ell = 1, \dots, L, \\ f_{\Sigma, \mu}(x), & x \in X^{L+1}. \end{cases}$$

Here,  $C_L^{\text{TT}} := (C_L^< + C_L^>, \text{TT})^{-1}$  with  $C_L^<$  from (23) and

$$(61) \quad C_L^>, \text{TT} := \sum_{\ell=1}^L \int_{X^\ell} \tilde{f}_0^{\ell, \text{TT}}(x) \, d\lambda(x).$$

Furthermore,  $\tilde{f}_0^{\ell, \text{TT}} = \hat{f}_0^{\ell, \text{TT}, N_\ell} \circ (\Phi^\ell)^{-1}$  is the pullback of a function  $\hat{f}_0^{\ell, \text{TT}, N_\ell}$  in  $\mathcal{M}^\ell = \mathcal{M}(\underline{c}_\ell, \bar{c}_\ell)$  over  $\hat{X}^\ell$ . Analog to the empirical minimization problem (52) with  $w_\ell = |\det \mathcal{J}_{\Phi^\ell}|$ , we choose  $\hat{f}_0^{\ell, \text{TT}, N_\ell}$  as

$$(62) \quad \hat{f}_0^{\ell, \text{TT}, N_\ell} \in \operatorname{argmin}_{v \in \mathcal{M}^\ell} \frac{1}{N_\ell} \sum_{k=1}^{N_\ell} \iota(v, \hat{x}^k, \hat{f}_0),$$

with samples  $\{\hat{x}^k\}_{k=1}^{N_\ell}$  drawn from the (possibly rescaled) finite measure  $w_\ell \lambda$ . The connection to the actual approximation of the target density  $f$  given by

$$(63) \quad \tilde{f}^{\text{TT}} := \tilde{f}_0^{\text{Trun,TT}} \circ \tilde{T}^{-1} \otimes |\mathcal{J}_{\tilde{T}^{-1}}|$$

is reviewed in the following. We refer to Figure 3 for a visual presentation of the involved objects, approximations and transformations.

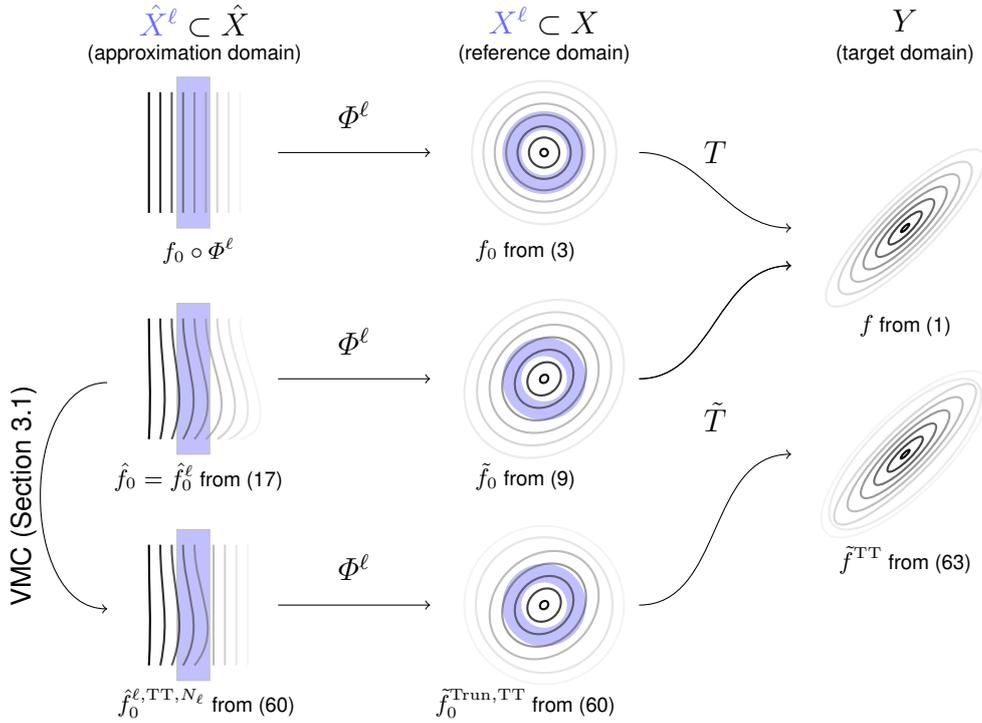


FIGURE 3. Overview of the presented method sketching the different involved transformations and approximations with references to the respective equations.

We first consider the relation of a target density  $f$  and its perturbed prior  $\tilde{f}_0$ . Since the transport  $\tilde{T}$  maps  $X$  to  $Y$ , an error functional  $d(Y; \cdot, \cdot)$  has to satisfy

$$(64) \quad d(Y; f, \tilde{f}^{TT}) = d(X; \tilde{f}_0, \tilde{f}_0^{\text{Trun, TT}}).$$

This property ensures that control of the error of the approximation in terms of the perturbed prior with respect to  $d(X; \cdot, \cdot)$  transfers directly to  $f$ . Note that this criterion is canonical as passing to the image space of some measurable function is fundamental in probability theory.

Prominent measures of discrepancy for two absolutely continuous Lebesgue probability density functions  $h_1$  and  $h_2$  on some measurable space  $Z$  are the Hellinger distance

$$(65) \quad d_{\text{Hell}}(Z, h_1, h_2) = \int_Z \left( \sqrt{h_1(z)} - \sqrt{h_2(z)} \right)^2 d\lambda(z),$$

and the Kullback-Leibler divergence

$$(66) \quad d_{\text{KL}}(Z, h_1, h_2) = \int_Z \log \left( \frac{h_1(z)}{h_2(z)} \right) h_1(z) d\lambda(z).$$

For the Hellinger distance, the absolute continuity assumption can be dropped from an analytical point of view. Observe that both  $d_{\text{Hell}}$  and  $d_{\text{KL}}$  both satisfy (64).

**Lemma 4.1.** *Let  $\sharp \in \{\text{Hell}, \text{KL}\}$ , then it holds*

$$(67) \quad d_{\sharp}(Y; f, \tilde{f}^{TT}) = d_{\sharp}(X; \tilde{f}_0, \tilde{f}_0^{\text{Trun, TT}}).$$

*Proof.* We only show (67) for  $\sharp = \text{KL}$  since  $\sharp = \text{Hell}$  follows by similar arguments. By definition

$$(68) \quad d_{\text{KL}}(Y; f, \tilde{f}^{TT}) = \int_Y \log \left( \frac{f(y)}{\tilde{f}^{TT}(y)} \right) f(y) d\lambda(y),$$

and the introduction of the transport map  $\tilde{T}$  yields the claim

$$(69) \quad \int_X \log \left( \frac{f \circ \tilde{T}(x)}{\tilde{f}^{\text{TT}} \circ \tilde{T}(x)} \cdot \frac{|\det \mathcal{J}_{\tilde{T}}(x)|}{|\det \mathcal{J}_{\tilde{T}}(x)|} \right) \tilde{f}_0(x) d\lambda(x) = d_{\text{KL}}(X; \tilde{f}_0, \tilde{f}_0^{\text{Trun,TT}}).$$

□

With the previous results and notations, the following assumption turns out to be required for the convergence result.

**Assumption 4.2.** For a target density  $f: Y \rightarrow \mathbb{R}_+$  and a transport map  $\tilde{T}: X \rightarrow Y$ , there exists a simply connected compact domain  $K$  such that  $\tilde{f}_0 = (f \circ \tilde{T}) \otimes |\det \mathcal{J}_{\tilde{T}}| \in L^2(K)$  has outer polynomial exponential decay with polynomial  $\pi^+$  on  $X \setminus K$ . Consider the symmetric positive definite matrix  $\Sigma \in \mathbb{R}^{d,d}$  and  $\mu \in \mathbb{R}^d$  as the covariance and mean for the outer approximation  $f_{\Sigma, \mu}$ . Furthermore, let  $K = \bigcup_{\ell=1}^L \overline{X^\ell}$  with  $X^\ell$  being the image of a rank-1 stable diffeomorphism  $\Phi^\ell: \hat{X}^\ell \rightarrow X^\ell$  for every  $\ell = 1, \dots, L$ .

We can now formulate the main theorem of this section regarding the convergence of the developed approximation with respect to the Hellinger distance and the KL divergence.

**Theorem 4.3. (A priori convergence)** Let Assumption 4.2 hold and let a sequence of sample sizes  $(N^\ell)_{\ell=1}^L \subset \mathbb{N}$  be given. For every  $\ell = 1, \dots, L$ , consider bounds  $0 < \underline{c}^\ell < \bar{c}^\ell < \infty$  and let  $\tilde{f}^{\text{TT}}$  be defined as in (63). Then there exist constants  $C, C_\Sigma, C^\ell, C_t^\ell > 0$ ,  $\ell = 1, \dots, L$ , such that for  $\sharp \in \{\text{KL}, \text{Hell}\}$

$$(70) \quad d_\sharp(Y, f, \tilde{f}^{\text{TT}}) \leq C \left( \sum_{\ell=1}^L (\mathcal{E}_{\text{best}}^\ell + \mathcal{E}_{\text{sing}}^\ell + \mathcal{E}_{\text{gen}}^\ell) + \mathcal{E}_{\text{trun}}^\sharp \right).$$

Here,  $\mathcal{E}_{\text{best}}^\ell$  denotes the error of the best approximation  $v_A^\ell$  to  $\hat{f}_0^\ell$  in the full truncated polynomial space  $\mathcal{V}_A^\ell(\underline{c}^\ell, \bar{c}^\ell) = \mathcal{V}_A \cap \mathcal{V}(\hat{X}^\ell, \underline{c}^\ell, \bar{c}^\ell)$  given by

$$\mathcal{E}_{\text{best}}^\ell := \|\hat{f}_0^\ell - v_A^\ell\|_{\mathcal{V}(\hat{X}^\ell)} = \inf_{v \in \mathcal{V}_A^\ell(\underline{c}^\ell, \bar{c}^\ell)} \|\hat{f}_0^\ell - v\|_{\mathcal{V}(\hat{X}^\ell)},$$

$\mathcal{E}_{\text{sing}}^\ell$  is the low-rank approximation error of the algebraic tensor associated to  $v_A^\ell$  and the truncation error  $\mathcal{E}_{\text{trun}}$  is given by

$$\begin{aligned} (\mathcal{E}_{\text{trun}}^{\text{Hell}})^2 &:= \|\exp(-\pi^+)\|_{L^1(X \setminus K)} + \Gamma(d/2, C_\Sigma R^2), \\ \mathcal{E}_{\text{trun}}^{\text{KL}} &:= \int_{X \setminus K} \left( \frac{1}{2} \|x\|_{\Sigma^{-1}}^2 + \tilde{\pi}^+(x) \right) e^{-\tilde{\pi}^+(x)} d\lambda(x). \end{aligned}$$

Furthermore, for any  $(\epsilon^\ell)_{\ell=1}^L \subset \mathbb{R}_+$  the generalization errors  $\mathcal{E}_{\text{gen}}^\ell$  can be bounded in probability

$$\mathbb{P}(\mathcal{E}_{\text{gen}}^\ell > \epsilon^\ell) \leq 2\nu(\mathcal{M}^\ell, C^\ell \epsilon^\ell) \delta^\ell(1/4\epsilon^\ell, N^\ell)$$

with  $\nu$  denoting the covering number from Definition 3.5 and  $\delta^\ell(\epsilon, N) \leq 2 \exp(-2\epsilon^2 N / C^\ell)$ .

*Proof.* We first prove (70) for  $\sharp = \text{Hell}$  and point out that the Hellinger distance can be bounded by the  $L^2$  norm. Note that  $|\sqrt{a} - \sqrt{b}| \leq \sqrt{|a - b|}$  for  $a, b \geq 0$  and with Lemma 4.1 it holds

$$d_{\text{Hell}}(Y; f, \tilde{f}^{\text{TT}}) = d_{\text{Hell}}(X; \tilde{f}_0, \tilde{f}_0^{\text{Trun,TT}}) \leq \|\tilde{f}_0 - \tilde{f}_0^{\text{Trun,TT}}\|_{L^1(K)} + \|\tilde{f}_0 - \tilde{f}_0^{\text{Trun,TT}}\|_{L^1(X \setminus K)}.$$

Since  $K = \cup_{\ell=1}^L X^\ell$  and  $X^\ell$  are bounded, there exist constants  $C(X^\ell) > 0$ ,  $\ell = 1, \dots, L$ , such that

$$\|\tilde{f}_0 - \tilde{f}_0^{\text{Trun,TT}}\|_{L^1(K)} = \sum_{\ell=1}^L \|\tilde{f}_0 - \tilde{f}_0^{\text{Trun,TT}}\|_{L^1(X^\ell)} \leq \sum_{\ell=1}^L C(X^\ell) \|\tilde{f}_0 - \tilde{f}_0^{\text{Trun,TT}}\|_{L^2(X^\ell)}.$$

Moreover, by construction

$$(71) \quad \|\tilde{f}_0 - \tilde{f}_0^{\text{Trun,TT}}\|_{L^2(X^\ell)} = \|\hat{f}_0^\ell - \hat{f}_0^{\ell,\text{TT},N_\ell}\|_{\mathcal{V}(\hat{X}^\ell)}.$$

The claim follows by application of Lemmas 2.7, 3.1 and 3.6 together with (57).

To show (70) for  $\sharp = \text{Hell}$ , note that by Lemma 4.1 and the construction (60) it holds

$$(72) \quad d_{\text{KL}}(Y; f, \tilde{f}^{\text{TT}}) = \sum_{\ell=1}^L \int_{X^\ell} \log \frac{\tilde{f}_0}{\tilde{f}_0^{\ell,\text{TT}}} \tilde{f}_0 d\lambda(x) + \int_{X \setminus K} \log \frac{\tilde{f}_0}{f_{\Sigma,\mu}} \tilde{f}_0 d\lambda(x).$$

Using Lemma 2.7 we can bound the integral over  $X \setminus K$  by the truncation error  $\mathcal{E}_{\text{trun}}$ . Employing the loss function and cost functional of Lemma 3.3 yields

$$(73) \quad \int_{X^\ell} \log \frac{\tilde{f}_0}{\tilde{f}_0^{\ell,\text{TT}}} \tilde{f}_0 d\lambda(x) \leq \mathcal{E}_{\text{app}}^\ell + \mathcal{E}_{\text{gen}}^\ell.$$

The claim follows by application of Lemmas 3.1 and 3.6 together with (57).  $\square$

**4.1. Polynomial approximation in weighted  $L^2$  spaces.** In order to make the error bound (70) in Theorem 4.3 more explicit with respect to  $\mathcal{E}_{\text{best}}$ , we consider the case of a smooth density function with analytic extension. The analysis follows the presentation in [1] and leads to exponential convergence rates by an iterative interpolation argument based on univariate best approximation bounds by interpolation. An analogous analysis for more general regularity classes is possible but not in the scope of this article.

Let  $\hat{X} = \otimes_{i=1}^d \hat{X}_i \subset \mathbb{R}^d$  be bounded and  $w = \otimes_{i=1}^d w_i \in L^\infty(\hat{X})$  a non-negative weight such that  $\mathcal{C}(\hat{X}) \subset \mathcal{V} := L^2(\hat{X}, w) = \otimes_{i=1}^d L^2(\hat{X}_i, w_i)$ .

For a Hilbert space  $H$ , a bounded set  $I \subset \mathbb{R}$  and a function  $f \in \mathcal{C}(I; H) \subset L^2(I, w; H)$  with weight  $w: I \rightarrow \mathbb{R}$ , let  $\mathcal{I}_n: \mathcal{C}(I; H) \rightarrow L^2(I, w; H)$  defined as

$$\mathcal{I}_n f(\cdot) = \sum_{k=1}^{n+1} f(\hat{x}_k) \ell_k(\cdot),$$

denote the continuous Lagrange interpolation operator. The  $\ell_k$  are polynomials of degree  $k$  orthogonal in  $L^2(I, w)$  and  $(\hat{x}_k)_{k=1}^n$  are the roots, respectively.

Assume that  $f \in \mathcal{C}(I; H)$  admits an analytic extension in the region of the complex plane  $\Sigma(I; \tau) := \{z \in \mathbb{C} \mid \text{dist}(z, I) \leq \tau\}$  for some  $\tau > 0$ . Then, referring to [1],

$$(74) \quad \|f - \mathcal{I}_n f\|_{L^2(I, w; H)} \lesssim \sigma(n, \tau) \max_{z \in \Sigma(I; \tau)} \|f(z)\|_H,$$

with  $\sigma(n, \tau) := 2(\rho - 1)^{-1} \exp(-n \log(\rho))$  and  $\rho := 2\tau/|I| + \sqrt{1 + 4\tau^2/|I|^2} > 1$ . By using an iterative argument over  $d$  dimensions, a convergence rate for the interpolation of  $f \in \mathcal{C}(\hat{X}; \mathbb{R}) \subset L^2(\hat{X}, w; \mathbb{R})$  can be derived from the one dimensional convergence. More specifically, let  $\mathcal{I}_\Lambda: \mathcal{C}(\hat{X}) \mapsto L^2(\hat{X}, w)$  denote the continuous interpolation operator written as composition of a 1-dimensional and a  $d-1$ -dimensional interpolation  $\mathcal{I}_\Lambda := \mathcal{I}_{n_1}^1 \circ \mathcal{I}_{n_2:n_d}^{2:d}$  with continuous  $\mathcal{I}_{n_1}^1: \mathcal{C}(\hat{X}_1) \rightarrow$

$L^2(\times_{i=2}^d \hat{X}_i, \otimes_{i=2}^d w_i)$  and  $\mathcal{I}_{n_2, \dots, n_d}^{2, \dots, d} : \mathcal{C}(\times_{i=2}^d \hat{X}_i) \rightarrow H$  with  $H = L^2(\times_{i=2}^d \hat{X}_i, \otimes_{i=2}^d w_i)$ . Then, for  $f \in \mathcal{C}(\hat{X})$  and some  $C > 0$  it follows

$$\begin{aligned} \|f - \mathcal{I}_\Lambda f\| &\leq \|f - \mathcal{I}_{n_1}^1 f\| + \|\mathcal{I}_{n_1}^1 (f - \mathcal{I}_{n_2, \dots, n_d}^{2, \dots, d} f)\| \\ &\lesssim \|f - \mathcal{I}_{n_1}^1 f\| + \sup_{\hat{x}_1 \in \hat{X}_1} \|f(\hat{x}_1) - \mathcal{I}_{n_2, \dots, n_d}^{2, \dots, d} f(\hat{x}_1)\|_H. \end{aligned}$$

The second term of the last bound is a  $d - 1$ -dimensional interpolation and can hence be bounded uniformly over  $\hat{x}_1$  by a similar iterative argument. We summarize the convergence result for  $\mathcal{E}_{\text{best}}^\ell$  in the spirit of [1, Theorem 4.1].

**Lemma 4.4.** *Let  $\hat{f} \in \mathcal{C}(\hat{X}^\ell) \subset L^2(\hat{X}^\ell, w)$  admit an analytic extension in the region*

$$\Sigma(\hat{X}^\ell, (\tau_i^\ell)_{i=1}^d) = \bigtimes_{i=1}^d \Sigma(\hat{X}_i^\ell, \tau_i^\ell)$$

for some  $\tau_i^\ell > 0$ ,  $\ell = 1, \dots, L$ ,  $i = 1, \dots, d$ . Then, with  $\sigma$  from (74),

$$\inf_{v \in \mathcal{V}_\Lambda} \|\hat{f} - v\|_{L^2(\hat{X}^\ell, w)} \lesssim \sum_{i=1}^d \sigma(n_i, \tau_i).$$

In case that  $\underline{c} \leq f(\hat{x})$ ,  $v^*(\hat{x}) \leq \bar{c}$  is satisfied for  $v^* := \operatorname{argmin}_{v \in \mathcal{V}_\Lambda} \|f - v\|_{L^2(\hat{X}^\ell, w)}$ , the decay rate carries over onto the space  $\mathcal{V}_\Lambda^\ell(\underline{c}^\ell, \bar{c}^\ell)$ . If only  $\underline{c} \leq f(\hat{x}) \leq \bar{c}$  holds, the image of  $v^*$  can be restricted to  $[\underline{c}, \bar{c}]$ , see e.g. [8]. This approximation in fact admits a smaller error than  $v^*$ .

**Remark 4.5.** *The interpolation argument on polynomial discrete spaces could be expanded to other orthonormal systems such as trigonometric polynomial, admitting well-known Lebesgue constants as in [9].*

**Remark 4.6.** *Explicit best approximation bounds for appropriate smooth weights  $w$ , as in the case of spherical coordinates, can be obtained using partial integration techniques as in [37]. There the regularity class of  $f$  is based on high-order weighted Sobolev spaces based on derivatives of  $w$  as in the case of classical polynomials.*

## 5. ALGORITHM

Since a variety of techniques are employed in the density discretization, this section provides an exemplary algorithmic workflow to illustrate the required steps in practical applications (see also Figure 1 for a sketch of the components of the method). The general method to obtain a representation of the density (1) by its auxiliary reference (9) is summarized in Algorithm 1. Based on this, the computation of possible quantities of interest such as moments (10) or marginals are considered in Sections 6.3 and 6.4, respectively. In the following we briefly describe the involved algorithmic procedures.

**Computing the transformation.** Obtaining a suitable transport map is a current research topic and examined e.g. in [42, 44, 55, 36]. In Section 2.1, two naive options are introduced. In the numerical applications, we employ an affine transport and also illustrate the capabilities of a quadratic transport in a two-dimensional example. For the affine linear transport we utilize a semi-Newton optimizer to obtain the maximum value of  $f$  and an approximation of the Hessian at the optimal value, see Section 2.1.1. For the construction of a quadratic transport we rely on the library `TransportMaps` [5]. We summarize the task to provide the (possibly inexact) transport map in the function

$$(75) \quad \tilde{T} \leftarrow \text{ComputeTransport}[f].$$

In the following paragraphs we assume  $\Phi^\ell$  to be the multivariate polar transformation as in Example 2.2, defined on the corresponding hyperspherical shells  $\hat{X}^\ell$ . We refer to  $\hat{X}_1^\ell$  as the *radial dimension* and  $\hat{X}_i^\ell$  as the *angular dimensions* for  $1 < i \leq d$ . The computations on each shell  $\hat{X}^\ell$ ,  $\ell = 1, \dots, L$  are fully decoupled and suitable for parallelization. Note that the proposed method is easily adapted to other transformations  $\Phi^\ell$ .

**Generating an orthonormal basis.** To obtain suitable finite dimensional subspaces, one has to introduce spanning sets that allow for an efficient computation of e.g. moments (4) and the optimization of the functional (42). Given a fixed dimension vector  $\mathbf{n}^\ell \in \mathbb{N}^d$  for the current  $\hat{X}^\ell$ ,  $\ell = 1, \dots, L$ , and by the chosen parametrization via  $\Phi^\ell$  introducing the weight  $w^\ell$ , the function

$$(76) \quad \mathcal{P}^\ell = \{\mathcal{P}_i^\ell\}_{i=1}^d \leftarrow \text{GenerateONB}[\hat{X}^\ell, \mathbf{n}^\ell, w^\ell, \tau_{\text{GS}}]$$

can be split into three distinct algorithmic parts as follows.

- *1st coordinate  $\hat{x}_1$ :* The computation of an orthonormal polynomial basis  $\{P_{1,\alpha}^\ell\}_\alpha$  with respect to the weight  $w_1^\ell(\hat{x}_1) = \hat{x}_1^{d-1}$  in the radial dimension by a stabilized Gram-Schmidt method. This is numerically unstable since the involved summations cause cancellation. As a remedy, we define *arbitrary precision polynomials* with a significant digit length  $\tau_{\text{mant}}$  to represent polynomial coefficients. By this, point evaluations of the orthonormal polynomials and computations of integrals of the form

$$(77) \quad \int_{\hat{X}_1^\ell} \hat{x}_1^m P_{1,\alpha}^\ell(\hat{x}_1) \hat{x}_1^{d-1} d\lambda(\hat{x}_1), \quad m \in \mathbb{N},$$

e.g. required for computing moments with polynomial transport, can be realized with high precision. The length  $\tau_{\text{mant}}$  is set to 100 in the numerical examples and the additional run-time is negligible as the respective calculations can be precomputed.

- *2nd coordinate  $\hat{x}_2$ :* Since  $\hat{X}_2^\ell = [0, 2\pi]$  and to preserve periodicity, we employ trigonometric polynomials given by

$$(78) \quad P_{2,j}^\ell(\hat{x}_2) = \begin{cases} \frac{1}{\sqrt{2\pi}}, & j = 1 \\ \frac{\sin(\frac{j}{2}\hat{x}_2)}{\sqrt{\pi}}, & j \text{ even} \\ \frac{\cos(\frac{j-1}{2}\hat{x}_2)}{\sqrt{\pi}}, & j > 1 \text{ odd.} \end{cases}$$

Note that here the weight function is constant, i.e.

$w_2^\ell(\hat{x}_2) \equiv 1$ , and the defined trigonometric polynomials are orthonormal in  $L^2(\hat{X}_2^\ell)$ .

- *coordinate  $\hat{x}_3, \dots, \hat{x}_d$ :* On the remaining angular dimensions  $i = 3, \dots, d$ , we employ the usual Gram-Schmidt orthogonalization algorithm on  $[0, \pi]$  with weight function  $w_i^\ell(\hat{x}_i) = \sin^i(\hat{x}_i)$ , based on polynomials.

Fortunately, the basis for dimensions  $1 < i \leq d$  coincides on every layer  $\ell = 1, \dots, L$ . It hence can be computed just once and passed to the individual process handling the current layer. Only the basis in the radial dimension needs to be adjusted to  $\hat{X}^\ell$ . The parameter  $\tau_{\text{GS}}$  collects all tolerance parameters for the applied numerical quadrature and the significant digit length  $\tau_{\text{mant}}$ .

**Generation of Samples.** To generate samples on  $\hat{X}^\ell$  with respect to the weight function  $w^\ell$ , we employ inverse transform sampling. For this the weight function is rescaled to have unit norm in  $L^1(\hat{X}^\ell)$ . Then, the involved inverse cumulative distribution functions can be computed analytically. We denote the generation process of  $N \in \mathbb{N}$  samples as the function

$$(79) \quad \mathcal{S}^\ell := \left\{ \left( \hat{x}^s, \hat{f}_0^\ell(\hat{x}^s) \right) \right\}_{s=1}^N \leftarrow \text{GenerateSamples}[\hat{f}_0^\ell, \hat{X}^\ell, w^\ell, N].$$

**Reconstruction of a Tensor Train surrogate.** The VMC reconstruction approach of Section 3 is summarized in the function

$$(80) \quad \left\{ \hat{F}_{0,i}^{\ell, \text{TT}} \right\}_{i=1}^d \leftarrow \text{ReconstructTT}[\mathcal{S}^\ell, \mathcal{P}^\ell, \mathbf{r}^\ell, \tau_{\text{Recon}}].$$

The tensor components  $\hat{F}_{0,i}^{\ell, \text{TT}}$  are associated with the corresponding basis  $\mathcal{P}_i^\ell$  to form a rank  $\mathbf{r}^\ell$  extended tensor train as defined in (25) and (29). The additional parameter  $\tau_{\text{Recon}}$  collects all parameters that determine the VMC algorithm.

The method basically involves the optimization of a loss functional over the set of tensor trains with rank (at most)  $\mathbf{r}^\ell$ . In the presented numerical computations we consider a mean-square loss and the respective empirical approximation based on a current sample set  $\mathcal{S}^\ell$ . The tensor optimization, based on a rank adaptive, alternating direction fitting (ADF) algorithm, is implemented in the `xerus` library [32] and wrapped in the `ALEA` framework [16]. Additionally, the machine learning framework `PyTorch` [45] can be utilized in `ALEA` to minimize the empirical cost functional from (52) by a wide class of state-of-the-art stochastic optimizers. The latter enables stochastic gradient methods to compute the tensor coefficients as known from machine learning applications. Having this setting in mind, the actual meaning of the parameter  $\tau_{\text{Recon}}$  depends on the chosen optimizer. In this article we focus on the ADF implementation and initialize e.g. the starting rank, the number of iteration of the ADF and a target residual norm.

## 6. APPLICATIONS

In the preceding sections the creation of surrogate models of quite generic probability density functions were developed. Using this, in the following we focus on actual applications where such a representation is beneficial. We start with the framework of Bayesian inverse problems with target density (1) corresponding to the Lebesgue posterior density. Subsequently, we cover the computation of moments and marginals.

**6.1. Bayesian inversion.** This section is devoted to a brief review of the Bayesian paradigm. We recall the general formalism and highlight the notation with the setup of Section 2 in mind. We closely follow the presentation in [18] and refer to [54, 10, 33] for a comprehensive overview.

Let  $Y$ ,  $V$  and  $\mathcal{Y}$  denote separable Hilbert spaces equipped with norms  $\|\cdot\|_H$  and inner products  $\langle \cdot, \cdot \rangle_H$  for  $H \in \{Y, V, \mathcal{Y}\}$ . The uncertain quantity  $y \in Y$  is tied to the model output  $q \in V$  by the *forward map*

$$(81) \quad G: Y \rightarrow V, \quad \theta \mapsto q(y) := G(y).$$

The usual forward problem reads

$$(82) \quad \text{Given } y \in Y, \text{ find } q \in V.$$

In contrast to this, the inverse problem is defined by

$$(83) \quad \text{Given observations of } q, \text{ find } y \in Y.$$

The term *observations* is determined by a bounded linear operator  $\mathcal{O}: V \rightarrow \mathcal{Y}$  that describes the measurement process of the quantity  $q$ . In practical applications this could be direct observations at sensor points or averaged values from monitoring devices, e.g. with  $\mathcal{Y} = \mathbb{R}^J$  for some  $J \in \mathbb{N}$ .

Classically, the (deterministic) quantification problem (83) is not well-posed. To overcome this, a problem regularization of some kind is required. The chosen probabilistic approach introduces a random

**Algorithm 1** Tensor train surrogate creation of perturbed prior

**Input:** Lebesgue target density  $f: \mathbb{R}^d \rightarrow \mathbb{R}_+$  (1)  
 tensor spaces  $\{\hat{X}^\ell\}_{\ell=1}^L$ , with  $\hat{X}^\ell = \times_{i=1}^d \hat{X}_i^\ell$  (17)  
 coordinate transformations  $\Phi^\ell: \hat{X}^\ell \rightarrow X^\ell \subset \mathbb{R}^d$  (20)  
     with rank-1 Jacobians  $w^\ell := |\det[\mathcal{J}_{\Phi^\ell}]|: \hat{X}^\ell \rightarrow \mathbb{R}$   
 basis dimensions  $(\mathbf{n}^1, \dots, \mathbf{n}^L)$ ,  $\mathbf{n}^\ell \in \mathbb{N}^d$  for  $\ell = 1, \dots, L$  (29)  
 sample size  $N_\ell \in \mathbb{N}$ ,  $\ell = 1, \dots, L$  for level-wise reconstruction  
 tensor train ranks  $(\mathbf{r}^1, \dots, \mathbf{r}^L)$ ,  $\mathbf{r}^\ell \in \mathbb{N}^{d-1}$ , for  $\ell = 1, \dots, L$  (25)  
 Gram-Schmidt tolerance parameter  $\tau_{\text{GS}}$   
 tensor reconstruction parameter  $\tau_{\text{Recon}}$

**Output:** Level-wise low-rank approximation of perturbed prior

Diffeomorphism  $\tilde{T} \leftarrow \text{ComputeTransport}[f]$

**for**  $\ell = 1, \dots, L$ , (in parallel) **do**

- Set transformed *perturbed prior*  $\hat{f}_0^\ell(\hat{x}) := \left( f \circ \tilde{T} \otimes |\det \mathcal{J}_{\tilde{T}}| \right) \circ \Phi^\ell(\hat{x})$ ,  $\hat{x} \in \hat{X}^\ell$

- Build one-dimensional ONB  $\mathcal{P}_i^\ell$  of  $\mathcal{V}_{i, n_i^\ell} \subseteq L^2(\hat{X}_i^\ell, w_i^\ell)$  for  $i = 1, \dots, d$

$$\mathcal{P}^\ell = \{\mathcal{P}_i^\ell\}_{i=1}^d \leftarrow \text{GenerateONB}[\hat{X}^\ell, \mathbf{n}^\ell, w^\ell, \tau_{\text{GS}}]$$

- Generate samples with respect to the weight  $w^\ell$

$$\mathcal{S}^\ell := \left\{ \left( \hat{x}^s, \hat{f}_0^\ell(\hat{x}^s) \right) \right\}_{s=1}^N \leftarrow \text{GenerateSamples}[\hat{f}_0^\ell, \hat{X}^\ell, w^\ell, N]$$

- Reconstruct TT surrogate  $\tilde{f}_0^{\ell, \text{TT}}: \hat{X}^\ell \rightarrow \mathbb{R}$

$$\left\{ \tilde{F}_{0,i}^{\ell, \text{TT}} \right\}_{i=1}^d \leftarrow \text{ReconstructTT}[\mathcal{S}^\ell, \mathcal{P}^\ell, \mathbf{r}^\ell, \tau_{\text{Recon}}]$$

- Equip tensor components with basis

$$\hat{f}_0^{\ell, \text{TT}}(\hat{x}) := \sum_{\mathbf{k}} \prod_{i=1}^d \hat{f}_{0,i}^{\ell, \text{TT}}[k_{i-1}, k_i](\hat{x}_i)$$

$$\text{where } \hat{f}_{0,i}^{\ell, \text{TT}}[k_{i-1}, k_i](\hat{x}_i) := \sum_{j=1}^{n_j^\ell} \tilde{F}_{0,i}^{\ell, \text{TT}}[k_{i-1}, \mu_i, k_i] P_{i,j}^\ell(\hat{x}_i)$$

**end for**

**return**  $\left\{ \tilde{f}_\ell \right\}_{\ell=1}^L$

measurable additive noise  $\eta: (\Omega, \mathcal{U}, \mathbb{P}) \rightarrow (\mathcal{Y}, \mathcal{B}(\mathcal{Y}))$  with law  $\mathcal{N}(0, C_0)$  for some symmetric positive definite covariance operator  $C_0$  on  $\mathcal{Y}$  to define the noisy measurements

$$(84) \quad \delta = (\mathcal{O} \circ G)(y) + \eta =: \mathcal{G}(y) + \eta \quad \text{where } \mathcal{G}: Y \rightarrow \mathcal{Y}.$$

As a consequence, the quantities  $y$ ,  $q$  and  $\delta$  become random variables over a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$  with values in  $Y$ ,  $V$  and  $\mathcal{Y}$ , respectively. In [54] mild conditions on the forward operator are derived to show a continuous version of Bayes formula which yields the existence and uniqueness of the Radon-Nikodym derivative of the (posterior) measure  $\pi_\delta$  of the conditional random variable  $y|\delta$  with respect to a prior measure  $\pi_0$  of  $y$ . More precisely, by assuming Gaussian  $\eta$  and independence

with respect to  $y$ , both measures  $\pi_0$  and  $\pi_\delta$  on  $Y$  are related by the *Bayesian potential*

$$(85) \quad \Psi(y, \delta) := \frac{1}{2} \langle C_0^{-1}(\delta - \mathcal{G}(y)), \delta - \mathcal{G}(y) \rangle_Y$$

in the sense that

$$(86) \quad \frac{d\pi_\delta}{d\pi_0}(y) = Z^{-1} \exp(-\Psi(y, \delta)),$$

with normalization constant  $Z := \mathbb{E}_{\pi_0} [\exp(-\Psi(y, \delta))]$ . Note that we interchangeably write  $y$  as an element of  $Y$  and the corresponding random variable with values in  $Y$ .

**6.2. Bayesian inversion for parametric PDEs.** Random partial differential equations (PDEs), i.e. PDEs with correlated random data, play an important role in the popular field of Uncertainty Quantification (UQ). As a prominent benchmark example, we consider the ground water flow model, also called the Darcy problem, as e.g. examined in [15, 20, 17]. In this linear second order PDE model, the forward operator  $G$  in (81) on some domain  $D \subset \mathbb{R}^d$ ,  $d = 1, 2, 3$  is determined by a forcing term  $g \in L^2(D)$  and the random quantity  $a(y) \in L^\infty(D)$ , which for almost every  $y \in Y$  models a conductivity or permeability coefficient. The physical system is described by

$$(87) \quad -\operatorname{div}(a(y)\nabla q(y)) = g \quad \text{in } D, \quad q(y)|_{\partial D} = 0,$$

and the solution  $q(y) \in V := H_0^1(D)$  corresponds to the system response  $G(y) = q(y)$ . Pointwise solvability of (87) for almost every  $y \in Y$  is guaranteed by a Lax-Milgram argument. For details we refer to [53].

For the applications in this article we employ a truncated log-normal coefficient field

$$(88) \quad a(y) = \exp\left(\sum_{k=1}^d a_k y_k\right)$$

for some fixed  $(a_k)_{k=1}^d$  with  $a_k \in L^2(D)$  and the image of some random variable with law  $\mathcal{N}(0, I)$  denoted by  $y = (y_k)_{k=1}^d \in Y$ . Assume point observations (84) of  $q$  at nodes  $\delta = (\delta_1, \dots, \delta_J)$  in  $D$  corresponding to some unknown  $q(y^*)$ ,  $y^* \in Y$ . We consider the Bayesian posterior density (86) and set

$$(89) \quad f(y) = Z^{-1} d\pi_\delta(y) d\pi_0(y)$$

as the Lebesgue density of the target measure  $\pi$  on  $Y$  according to (1).

**6.3. Moment computation.** In this section we discuss the computation of moments for the presented layer-based format with low-rank tensor train approximations. In particular we are interested in an efficient generation of the moment map

$$(90) \quad \alpha \mapsto \int_Y y^\alpha f(y) d\lambda(y), \quad \alpha = (\alpha_k)_k \in \mathbb{N}_0^d.$$

Given some transport  $\tilde{T}: X \rightarrow Y$  with an associated perturbed prior  $\tilde{f}_0 = (f \circ \tilde{T}) \otimes |\det \mathcal{J}_{\tilde{T}}|$ , by an integral transformation it holds

$$(91) \quad \int_Y y^\alpha f(y) d\lambda(y) = \int_X \tilde{T}(x)^\alpha \tilde{f}_0(x) d\lambda(x).$$

We fix  $1 \leq \ell \leq L$  and assume tensor spaces  $\hat{X}^\ell, X^\ell$  such that a layer based splitting can be employed to obtain integrals over  $X^\ell$  of the form

$$(92) \quad \int_Y y^\alpha f(y) d\lambda(y) = \sum_{\ell=1}^L \int_{X^\ell} \tilde{T}(x)^\alpha \tilde{f}_0(x) dx.$$

Note that we neglect the remaining unbounded layer  $X^{L+1}$  since for moderate  $|\alpha|$  and  $\text{vol}(\bigcup_{\ell=1}^L X^\ell)$  sufficiently large, the contribution to the considered moment does not have a significant influence on the overall approximation. Additionally, a rank-1 stable diffeomorphism  $\Phi^\ell: \hat{X}^\ell \mapsto X^\ell$  is assumed for which there exist univariate functions  $\Phi_{:,j}^\ell: \hat{X}_j^\ell \rightarrow X^\ell$  with  $\Phi_{:,j}^\ell = (\Phi_{i,j}^\ell)_{i=1}^d$  and  $h_j: \hat{X}_j^\ell \rightarrow \mathbb{R}$  for every  $j = 1, \dots, d$ , such that

$$(93) \quad \Phi^\ell(\hat{x}) = \prod_{j=1}^d \Phi_{:,j}^\ell(\hat{x}_j) \quad \text{and} \quad |\det[\mathcal{J}_{\Phi^\ell}](\hat{x})| = \prod_{j=1}^d h_j(\hat{x}_j).$$

**6.3.1. Moments under affine transport.** Let  $H = [h_{ki}]_{k,i=1}^d = [h_1, h_2, \dots, h_d] \in \mathbb{R}^{d,d}$  be a symmetric positive definite matrix and  $M = (M_i)_{i=1}^d \in \mathbb{R}^d$  such that the considered transport map takes the form

$$(94) \quad \tilde{T}(\cdot) = H \cdot + M.$$

With the multinomial coefficient for  $j \in \mathbb{N}$ ,  $\beta \in \mathbb{N}_0^d$  with  $j = |\beta|$  given by

$$\binom{j}{\beta} := \frac{j!}{\beta_1! \cdot \dots \cdot \beta_d!},$$

the computation of moments corresponds to the multinomial theorem as seen in the next lemma.

**Lemma 6.1.** *Let  $k \in \mathbb{N}$  with  $1 \leq k \leq d$  and  $\alpha_k \in \mathbb{N}_0$ . It holds*

$$(95) \quad [H\Phi^\ell(\hat{x}) + M]_k^{\alpha_k} = \sum_{j_k=0}^{\alpha_k} \sum_{|\beta_k|=j_k} C_k^H[j_k, \alpha_k, \beta_k] \prod_{j=1}^d \Phi_j^{\beta_k}(\hat{x}_j),$$

where the high-dimensional coefficient  $C_k^H$  is given by

$$(96) \quad C_k^H[j_k, \alpha_k, \beta_k] := \binom{\alpha_k}{j_k} c_k^{\alpha_k - j_k} \binom{j_k}{\beta_k} h_k^{\beta_k},$$

with  $c_k := \sum_{i=1}^d h_{ki} M_i$  and

$$(97) \quad \Phi_j^{\beta_k} := [\Phi_{1,j}^\ell(\hat{x}_j), \dots, \Phi_{d,j}^\ell(\hat{x}_j)]^{\beta_k}.$$

*Proof.* Note that

$$[H\Phi^\ell(\hat{x}) + M]_k^{\alpha_k} = \sum_{j_k=0}^{\alpha_k} \binom{\alpha_k}{j_k} c_k^{\alpha_k - j_k} \left( \sum_{i=1}^d h_{ki} \prod_{j=1}^d \Phi_{i,j}^\ell(\hat{x}_j) \right)^{j_k}.$$

The statement follows by the multinomial theorem since

$$\left( \sum_{i=1}^d h_{ki} \prod_{j=1}^d \Phi_{i,j}^\ell(\hat{x}_j) \right)^{j_k} = \sum_{|\beta_k|=j_k} \binom{j_k}{\beta_k} \left( \prod_{i=1}^d h_{ki}^{\beta_k(i)} \right) \left( \prod_{j=1}^d \prod_{i=1}^d \Phi_{i,j}^\ell(\hat{x}_j)^{\beta_k(i)} \right).$$

□

Generalizing Lemma 6.1 to multiindices  $\alpha \in \mathbb{N}_0^d$  yields

$$(98) \quad [H\Phi^\ell(\hat{x}) + M]^\alpha = \sum_{j=0}^{\alpha} \sum_{(|\beta_k|)_k=j} \left( \prod_{k=1}^d C_k^H[j_k, \alpha_k, \beta_k] \right) \prod_{j=1}^d \Phi_j^{\sum_{k=1}^d \beta_k}(\hat{x}_j),$$

where  $\sum_{(|\beta_k|)_k=j} := \sum_{|\beta_1|=j_1} \dots \sum_{|\beta_d|=j_d}$  is used.

Exploiting the layerwise TT representation of  $\hat{f}_\ell$  from (60) and using the rank-1 stable map (93), the high-dimensional integral over  $X^\ell$  reduces to

$$(99) \quad \int_{X_\ell} \tilde{T}(x)^\alpha \tilde{f}_0(x) d\lambda(x) = \sum_{j=0}^{\alpha} \sum_{(|\beta_k|)_k=j} \left( \prod_{k=1}^d C_k^H[j_k, \alpha_k, \beta_k] \right) \times \\ \times \sum_{k=0}^{r_\ell} \prod_{i=1}^d \int_{\hat{X}_i} \left[ \hat{f}_{\ell,i}[k_{i-1}, k_i] \otimes \Phi_i^{\sum_{k=1}^d \beta_k} \otimes h_i \right] (\hat{x}_i) d\hat{x}_i.$$

Note that the right-hand side is composed via decoupled one dimensional integrals only. We point out that while the structure is simplified, the definition of  $\Phi_j$  in (97) a priori results in several integrals (indexed by  $\sum_{k=1}^d \beta_k$ ). These integrals, whose number depends on the cardinality of  $\alpha$ , have to be computed. This simplifies further in several cases, e.g. when  $\Phi^\ell$  transforms the spherical coordinate system to Cartesian coordinates.

**Moment computation using spherical coordinates.** In the special case that  $\Phi^\ell$  is the multivariate polar transformation of Example 2.2, the number of distinct computation of integrals from (99) reduces significantly. Recall that  $\hat{x}_1 = \rho$ ,  $\hat{x}_{2:d} = \theta = (\theta_0, \dots, \theta_{d-2})$  and let  $\beta_i^k := (\beta_k)_i$  be the  $i$ -th entry of  $\beta_k$ . We find that

$$(100) \quad \Phi_1^{\sum_{k=1}^d \beta_k}(\rho) = \rho^{|\alpha|},$$

$$(101) \quad \Phi_2^{\sum_{k=1}^d \beta_k}(\theta_0) = \cos\left(\sum_{k=1}^d \beta_1^k\right)(\theta_0) \sin\left(\sum_{k=1}^d \beta_2^k\right)(\theta_0),$$

$$(102) \quad \Phi_{i+2}^{\sum_{k=1}^d \beta_k}(\theta_i) = \sin\left(\sum_{l=1}^i \sum_{k=1}^d \beta_l^k\right)(\theta_i) \cos\left(\sum_{k=1}^d \beta_{i+1}^k\right)(\theta_i).$$

for  $1 \leq i \leq d-2$ .

The exponential complexity due to the indexing by  $\sum_{k=1}^d \beta_k$  reduces to linear complexity in  $|\alpha|$ . More precisely, the amount of exponents in (100) - (102) is linear in the dimensions since the sums only depend on  $|\alpha|$ , leading to  $\mathcal{O}(|\alpha|d)$  different integrals that may be precomputed for each tuple  $(k_{i-1}, k_i)$ . This exponential complexity in the rank vanishes in the presence of an approximation basis associated with each coordinate dimension as defined in Section 3.

**6.4. Computation of marginals.** In probability theory and statistics, marginal distributions and especially marginal probability density functions provide insights into an underlying joint density by means of lower dimensional functions that can be visualized. The computation of marginal densities is a frequent problem encountered e.g. in parameter estimation and when using sampling techniques since histograms and corner plots provide easy access to (in general high-dimensional) integral quantities.

In contrast to the Markov chain Monte Carlo algorithm, the previously presented method of a layer based surrogate for the Lebesgue density function  $f: Y = \mathbb{R}^d \rightarrow \mathbb{R}$  allows for a functional representation and approximation of marginal densities without additional evaluations of  $f$ .

For simplicity, for  $y \in Y$  and  $i = 1, \dots, d$  define  $y_{-i} = (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_d)$  as the marginalized variable where the  $i$ -th component is left out and  $f(y_{-i}, y_i) := f(y)$ . Then, for given  $i = 1, \dots, d$ , the  $i$ -th marginal density reads

$$(103) \quad df_i(y_i) := \int_{\mathbb{R}^{d-1}} f(y_{-i}, y_i) d\lambda(y_{-i}).$$

Computing this high-dimensional integral by quadrature or sampling is usually infeasible and the transport map approach as given by (4) fails since the map  $T: X \rightarrow Y$  cannot be used directly in (103). Alternatively, we can represent  $df_i: \mathbb{R} \rightarrow \mathbb{R}$  in a given orthonormal basis  $\{\varphi_j\}_{j=1}^{N_\varphi}$  and consider

$$(104) \quad df_i(y_i) = \sum_{j=1}^{N_\varphi} \beta_j \varphi_j(y_i),$$

where  $\beta_j, j = 1, \dots, N_\varphi$  denotes the  $L^2(\mathbb{R})$  projection coefficient

$$(105) \quad \beta_j := \int_{\mathbb{R}} \varphi_j(y_i) df_i(y_i) d\lambda(y_i).$$

With this the marginalisation can be carried out similar to the computations in Section 6.3.

A convenient basis is given by monomials since (105) then simplifies to

$$(106) \quad \beta_j = \int_{\mathbb{R}^d} y_i^j f(y) d\lambda(y).$$

This is the moment corresponding to the multiindex  $\alpha = (\alpha_k)_{k=1}^d \in \mathbb{N}^d$  with  $\alpha_k = \delta_{k,j}$ . Alternatively, indicator functions may be considered in the spirit of histograms.

**6.5. More general quantities of interest.** One is frequently concerned with efficiently computing the expectation of some quantity of interest (QoI)  $Q: Y \rightarrow \mathbb{R}$

$$(107) \quad \mathbb{E}[Q] = \int_Y Q(y) f(y) d\lambda(y).$$

We discussed this issue for moments in Section 6.3 and basis representations of marginals in Section 6.4. In those cases the structure of  $Q$  allows for direct computations of the integrals via tensor contractions. For more involved choices of the QoI we suggest a universal sampling approach by repeated evaluation of the low-rank surrogate. More precisely, by application of the integral transformation we can approximate

$$(108) \quad \mathbb{E}[Q] \approx \sum_{\ell=1}^L \int_{\hat{X}^\ell} Q \circ \tilde{T} \circ \Phi^\ell(\hat{x}) \tilde{f}_0^{\ell, \text{TT}}(\hat{x}) |\det[\mathcal{J}_{\Phi^\ell}](\hat{x})| d\lambda(\hat{x})$$

and replace the integrals over  $\hat{X}^\ell$  by Monte Carlo estimates with samples according to the (normalized) weight  $|\det[\mathcal{J}_{\Phi^\ell}]|$ . Those samples can be obtained by uniform sampling on the tensor spaces  $\hat{X}^\ell$  and the inverse transform approach as mentioned in the paragraph **Generating Samples** of Section 5. Alternatively, efficient MCMC sampling by marginalization can be employed [57].

## 7. NUMERICAL VALIDATION AND APPLICATIONS

This section is devoted to a numerical validation of the proposed Algorithm 1 using various types of transformations  $T$  while employing it with practical applications. We focus on three example settings. The first consists of an artificial Gaussian posterior density, which could be translated to a linear forward model and Gaussian prior assumptions in the Bayesian setting. Second, we study the approximation under non-exact transport and conclude as a third setting with an actual Bayesian inversion application governed by the log-normal Darcy flow problem of Section 6.2.

**7.1. Validation experiment 1: Gaussian density.** In this experiment we confirm the theoretical results from Section 4 and verify the numerical algorithm. Even though the examined approximation of a Gaussian density is not a challenging task for the proposed algorithm, it can be seen as the most basic illustration revealing the possible rank-1 structure of the perturbed prior under optimal transport.

We consider the posterior density determined by a Gaussian density with covariance matrix  $\Sigma \in \mathbb{R}^{d,d}$  and mean  $\mu \in \mathbb{R}^d$  as

$$(109) \quad \frac{d\pi}{d\lambda}(x) = f(x) = C \exp\left(-\frac{1}{2}\|x - \mu\|_{\Sigma^{-1}}^2\right),$$

where  $C = (2\pi)^{-d/2} \det \Sigma^{-1/2}$  is the normalizing factor of the multivariate Gaussian. We set the covariance operator such that the Gaussian density belongs to uncorrelated random variables, i.e.  $\Sigma$  exhibits a diagonal structure, and it holds for some  $0 < \sigma \ll 1$  that  $\Sigma = \sigma^2 I$ . This Gaussian setting has several benefits as a validation setting. On the one hand, we have explicit access to the quantities that are usually of interest in Bayesian inference like the mean, covariance, normalization constant and marginals. On the other hand, the optimal transport to a standard normal density

$$(110) \quad f_0(x) = (2\pi)^{-d/2} \exp\left(-\frac{1}{2}\|x\|^2\right)$$

is given by an affine linear function, defined via mean  $\mu$  and covariance  $\Sigma$  as proposed in Remark 2.6. We subsequently employ the multivariate polar transformation from Example 2.2 and expect a rank-1 structure in the reconstruction of the local approximations of the (perturbed) prior.

The remainder of this section considers different choices of  $\sigma \in \mathbb{R}$  and  $d \in \mathbb{N}$  and highlights the stability of our method under decreasing variance (i.e. with higher density concentration) and increasing dimension. The approximations are compared with their exact counterparts. More specifically, the error of the normalization constant is observed, namely

$$(111) \quad \text{err}_Z := |1 - Z_h|,$$

the relative error of the mean and covariance in the Euclidean and Frobenius norms  $|\cdot|_2$  and  $|\cdot|_F$ ,

$$(112) \quad \text{err}_\mu := |\mu - \mu_h|_2 |\mu|_2^{-1}, \quad \text{err}_\Sigma := |\Sigma - \Sigma_h|_F |\Sigma|_F^{-1},$$

and the deviation in terms of the Kullback-Leibler divergence (66). Computing the Kullback-Leibler divergence is accomplished by Monte Carlo samples  $(x_i)_{i=1}^{N_{\text{KL}}}$  of the posterior (i.e. in this case the multivariate Gaussian posterior) to compute the empirical approximation

$$(113) \quad d_{\text{KL}}(\pi, \pi_h) = \int_{\mathbb{R}^d} \log\left(\frac{f(x)}{f_h(x)}\right) f(x) d\lambda(x) \approx \frac{1}{N_{\text{KL}}} \sum_{i=1}^{N_{\text{KL}}} \log\left(\frac{f(x_i)}{f_h(x_i)}\right).$$

The index  $h$  generically denotes the employed approximation (60). In the numerical experiments the convergence of these error measures is depicted with respect to the number of calls to the forward model (i.e. the Gaussian posterior density), the discretization of the radial component  $\rho \in [0, \infty)$  in

dimension	$\sigma^2 = 10^{-2}$	$\sigma^2 = 10^{-4}$	$\sigma^2 = 10^{-6}$	$\sigma^2 = 10^{-8}$
2	$5.24 \cdot 10^{-11}$	$1.09 \cdot 10^{-10}$	$2.8 \cdot 10^{-11}$	$9.3 \cdot 10^{-11}$
4	$2.21 \cdot 10^{-10}$	$4.57 \cdot 10^{-10}$	$5.48 \cdot 10^{-10}$	$3.4 \cdot 10^{-10}$
6	$5.01 \cdot 10^{-11}$	$9.5 \cdot 10^{-11}$	$7.49 \cdot 10^{-11}$	$6.19 \cdot 10^{-10}$
8	$1.48 \cdot 10^{-11}$	$8.21 \cdot 10^{-10}$	$2.99 \cdot 10^{-10}$	$2.1 \cdot 10^{-10}$
10	$2.91 \cdot 10^{-9}$	$9.61 \cdot 10^{-10}$	$4.43 \cdot 10^{-11}$	$2.46 \cdot 10^{-9}$

TABLE 1. Numerical approximation of  $Z$  in the Gaussian example. Error of the normalization constant computed via a TT surrogate to  $Z = 1$ .

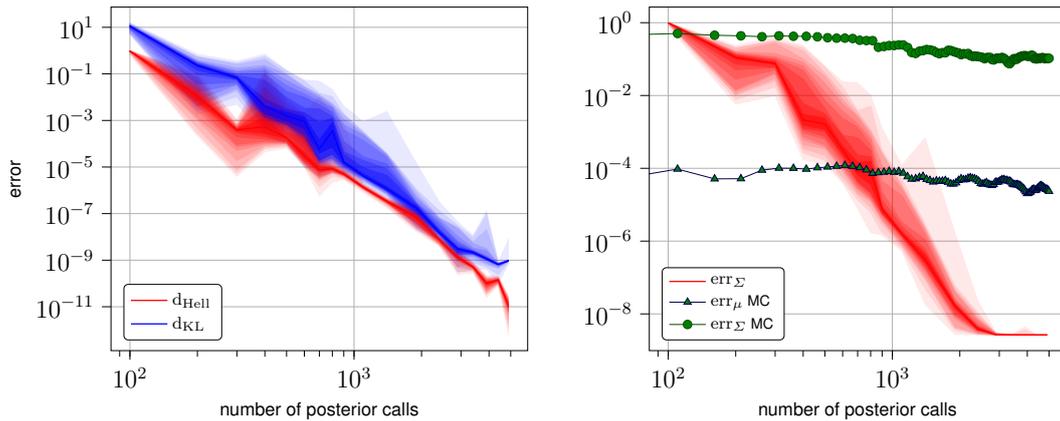


FIGURE 4. Gaussian density example with  $d = 10$ , mean  $\mu = 1$  and noise level  $\sigma = 10^{-7}$ . Tensor reconstructions are repeated with 50 random sample sets to show quantile range from 5% – 95% (pastel) to the median (bold). Hellinger distance and Kullback-Leibler divergence are shown (left) and the relative covariance error together with MCMC results for mean and covariance are given (right).

the polar coordinate system and the number of samples on each layer  $X^\ell$ ,  $\ell = 1, \dots, L$ , for fixed  $L \in \mathbb{N}$ .

In Table 7.1  $\text{err}_Z$  is depicted for different choices of  $\sigma$  and  $d$ . The experiment comprises radial discretizations  $0 = \rho_0 < \rho_1 < \dots < \rho_L = 10$  with  $L = 19$  equidistantly chosen layers and 1000 samples of  $f_0$  on each resulting subdomain  $X^\ell$ . The generated basis (76) contains polynomials of maximal degree 7 in  $\rho_\ell$ ,  $\ell = 0, \dots, L$ , and constant functions in every angular direction. The choice of constant functions relies on the assumption that the perturbed prior that has to be approximated corresponds to the polar transformation of (110), which is a function in  $\rho$  only. Additional numerical test show that even much fewer samples and a larger basis lead to the assumed rank-1 structure. It can be observed that the approximation quality of  $Z$  is invariant under the choice of  $\sigma$  and fairly robust with the dimension  $d$ , which is expected since the transformation is exact and the function to reconstruct is a rank-1 object.

In Figure 4 we compare the number of calls of the posterior density  $f$  explicitly. Here, the presented low-rank surrogate is again constructed on an increasing number of layers, whereas the Monte Carlo estimates are computed using a Markov Chain Monte Carlo algorithm and subsequent empirical integration of the error quantity. By taking 100 samples for each added layer, we observe fast convergence in comparison to the slow MC approach<sup>4</sup>. To further analyse the reconstruction stability we repeat the

<sup>4</sup>We emphasize that we just use a baseline MCMC algorithm for comparison. Although more sophisticated MCMC methods could show a more favorable convergence behavior, the fundamental qualitative difference due to entirely different approximation approaches would still persist.

experiment 50 times and show empirical quantiles. The light area represents the 90% quantile of the distribution and the bold line is the median. We observe a larger variance for the Kullback-Leibler divergence in contrast to the Hellinger distance.

Note that we do not show the tensor approximation result for  $\text{err}_\mu$  since already for the first case of only 100 evaluations of the posterior (which corresponds to a single layer) we obtain results close to machine precision. This is due to the choice of an exact transport, already containing the correct mean, and how the mean is computed in the presented format, see Section 6.3.1. In short, the approximation cancels due to the normalization and only the correct mean of the transport formula is left. Concerning the stagnation of  $\text{err}_\Sigma$  we suspect a precision problem in the computation, which is confirmed by the small variance. Nevertheless, an approximation of around seven magnitudes smaller than MCMC for the covariance is achieved.

**7.2. Validation experiment 2: Perturbation of exact transport.** In the following experiment we consider a so-called “banana example” as posterior density, see e.g. [36]. Let  $f_0$  be the density of a standard normal Gaussian measure and let  $T_\Sigma$  be the affine transport of  $\mathcal{N}(0, I)$  to the Gaussian measure  $\mathcal{N}(0, \Sigma)$ . Furthermore, set

$$(114) \quad T_2(x) = \begin{pmatrix} x_1 \\ x_2 - (x_1^2 + 1) \end{pmatrix}.$$

The exact transport  $T$  from  $\mathcal{N}(0, I)$  to the curved and concentrated banana distribution with density  $f$  is then given by

$$(115) \quad T(x) = T_2 \circ T_\Sigma(x), \quad \Sigma = \begin{pmatrix} 1 & 0.9 \\ 0.9 & 1 \end{pmatrix}.$$

Note that the employed density can be transformed into a Gaussian using a quadratic transport function. For this experiment, we employ transport maps  $\tilde{T}$  of varying accuracy for the pull-back of the posterior density to a standard Gaussian. In particular we use an approximation  $\tilde{T}_1$  (obtained with [5]) of the optimal affine transport  $T_1$ , and the quadratic transport  $T$  to build an approximation  $\tilde{T}$  given as convex combination

$$(116) \quad \tilde{T}(x) = (1 - t) \tilde{T}_1(x) + t T(x), \quad t \in [0, 1].$$

For  $t = 1$ , the transport map is optimal since it generates the desired reference density. For  $0 \leq t < 1$  a perturbed prior density is obtained with strength of perturbation determined by  $t$ . The impact of the perturbed transport is visualized in Figure 5. It can be observed that the transformed perturbed prior is not of rank-1 as long as the transformation is inexact. Furthermore, the difference between the target prior and the perturbed prior is eminent, which implies that e.g. a Laplace approximation to the considered banana density would neglect possible important features of the distribution.

In Figure 6 we illustrate the impact of an inexact transport on the approximation results in terms of  $\text{err}_\mu$  and  $\text{err}_\Sigma$ . For the considered target density, mean and covariance are known analytically and hence no reference sampling has to be carried out. We additionally employ an MCMC sampling to show the improvement due to the additional low-rank reconstruction. For the optimal transport map one observes that the surrogate reconstruction reduces to the approximation of a rank-1 Gaussian density, which can be done efficiently with few evaluations of  $f$ . If the transport is only linear and inaccurate, results comparable to MCMC are achieved. For a more accurate transport, the low-rank reconstruction leads to drastically improved estimates.

**7.3. Bayesian inversion with log-normal Darcy forward model.** Revisiting the example of Section 6.2, we consider the elliptic diffusion problem with a log-normal random parametric permeability

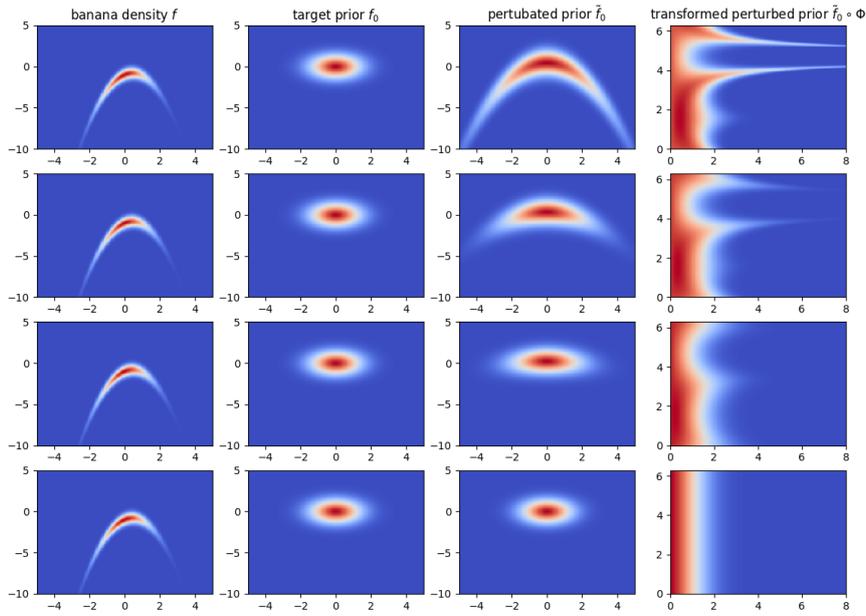


FIGURE 5. Illustration of the effect of different transports in (116) for  $t = 0, 0.25, 0.5, 1$ . (top to bottom).

coefficient. The considered field in  $L^2(Y, L^\infty(D))$  takes the form

$$(117) \quad a(x, y) = \exp\left(\sum_{i=1}^d a_i(x)y_i\right),$$

where the  $y_i$  correspond to random variables with law  $\mathcal{N}(0, 1)$  and  $L^2(D)$  orthonormal functions  $a_i$  being planar Fourier cosine modes. A detailed description and an adaptive Galerkin approach to solve the forward problem can be found in [17]. For the inverse problem, the observation operator is modelled by  $J = 144$  equidistantly distributed observations in  $D = [0, 1]^2$  of the solution  $q(y^*) \in H_0^1(D)$  for some  $y^* \in Y = \mathbb{R}^d$ , which is drawn from a standard normal distribution. Additionally, the observations are perturbed by a centered Gaussian noise with covariance  $\sigma I$  with  $\sigma = 10^{-7}$ .

To obtain the desired relative error quantities, we employ reference computations that involve adaptive quadrature for the two dimensional example in Figure 7 and Markov Chain Monte Carlo integration with  $10^6$  steps of the chain and a burn-in time of 1000 samples for the experiment in Figure 8. For the reconstruction algorithm an affine linear transport is estimated by Hessian information of the log-likelihood and on every layer we employ 100 samples. The respective relative errors are displayed in Figure 8.

The stagnation of the graphs in Figure 7 is on the one hand governed by the observation noise and on the other hand explicable by a non-optimal reference solutions since the TT approximation yields results equivalent to an adaptive quadrature when taking  $L = 5$  layers of refinement and thus a total of 500 samples.

The improvement of the mean and covariance estimate by the low-rank approach can already be observed for a low sample number. We note that the Monte Carlo estimate did not allow for an adequate computation of the empirical covariance, which therefore is left out of the comparison.

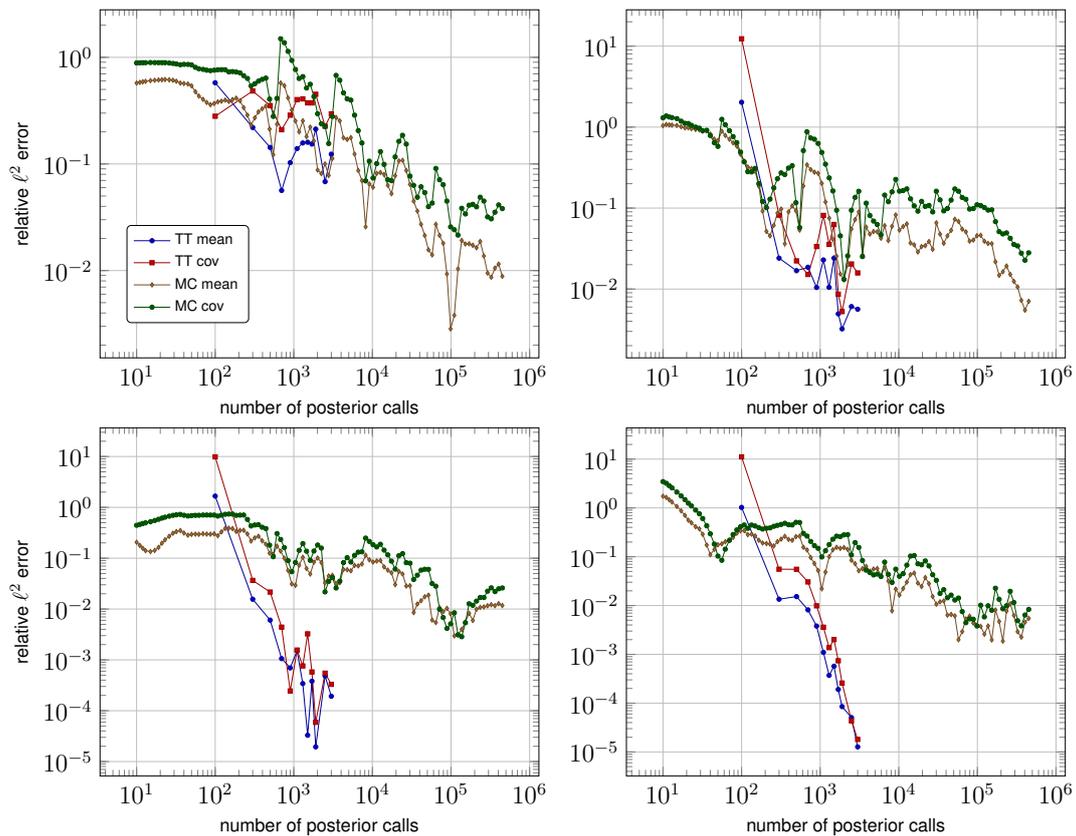


FIGURE 6. Convex combination of affine and quadratic transport for the banana posterior. Affine linear map ( $t = 0$  top left), transport with  $t = 0.25$  (top right),  $t = 0.5$  (bottom left) and exact quadratic map ( $t = 1$ , bottom right). Error quantities  $\text{err}_\mu$  and  $\text{err}_\Sigma$  for the employed tensor train surrogate and a MCMC approximation in terms of the number of calls to the posterior function. The surrogate is reconstructed from 100 samples per layer yielding a tensor with radial basis up to polynomial degree 9 and Fourier modes up to degree 20.

## 8. CONCLUSION

We developed a novel approach to approximate probability densities with high accuracy, combining the notion of *transport maps* and *low-rank functional representations* of auxiliary (perturbed) reference densities. Based on a suitable class of transformations, an approximation with respect to a finite tensorized basis can be carried out in extended hierarchical tensor formats. This yields a compressed representation for an efficient computation of statistical quantities (e.g. moments or marginals) of interest in a sampling free manner. In this work the multivariate polar transformation is used as a particular rank 1 stable transformation. The method requires point evaluations of the perturbed reference density (up to a multiplicative constant). The approach can hence be applied to not normalized posterior densities in the context of Bayesian inversion.

We presented the application of the method to an inverse problem with a log-normal Darcy forward model. A comparison with classical MCMC illustrates the superior convergence in terms of the moment accuracy relative to the number of posterior evaluations. Future research will be concerned with

- application: usage of the approximated densities for subsequent computations e.g. with SGFEM,
- analysis: Given a function  $f_0$  it has to be examined which rank 1 stable transformations  $\Phi$  lead to a low-rank function  $\tilde{f}_0 \circ \Phi$ .

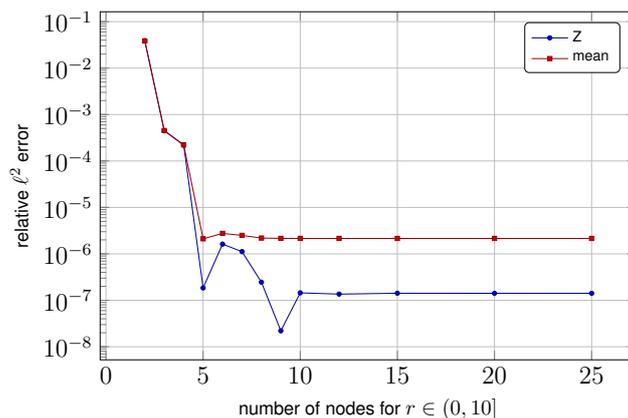


FIGURE 7. Comparison of the computed reference and the low-rank surrogate of (1) normalization constant ( $\text{err}_Z$ ), and (2) mean ( $\text{err}_\mu$ ). For the Darcy setting with  $d = 2$  we observe 144 nodes in the physical domain. The measurements are perturbed by Gaussian noise with deviation  $\eta = 1e - 7$ . We employ an adaptive quadrature in the two dimensional space to obtain the reference quantities. The stagnation of the graphs are due to non-optimal reference solutions. More precisely, the TT approximation yields equivalent results to adaptive quadrature when taking 5 nodes of refinement.

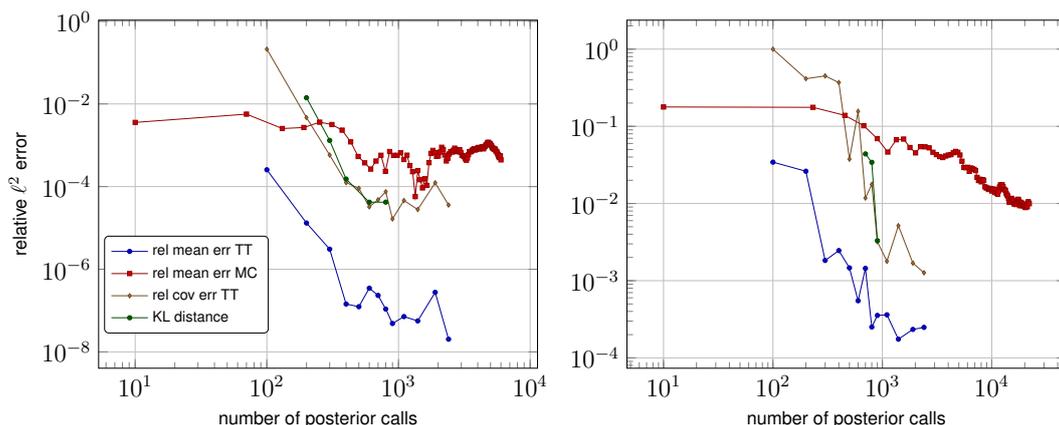


FIGURE 8. Darcy example with  $d = 2$  (left) and  $d = 10$  (right). Comparison of an MCMC method and the low-rank surrogate for computing the mean error ( $\text{err}_\mu$ ) with respect to the number of calls to the solution of the forward problem. The reference mean is computed with  $10^6$  MCMC samples. Additionally the KL divergence is shown, which is computed using empirical integration.

ACKNOWLEDGMENTS

The authors would like to thank Reinhold Schneider for fruitful discussions.

REFERENCES

[1] Ivo Babuška, Fabio Nobile, and Raúl Tempone. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM review*, 52(2):317–355, 2010.  
 [2] Markus Bachmayr, Albert Cohen, and Wolfgang Dahmen. Parametric pdes: sparse or low-rank approximations? *IMA Journal of Numerical Analysis*, 38(4):1661–1708, 2017.

- [3] Markus Bachmayr, Reinhold Schneider, and André Uschmajew. Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations. *Foundations of Computational Mathematics*, 16(6):1423–1472, 2016.
- [4] Jonas Ballani, Lars Grasedyck, and Melanie Kluge. Black box approximation of tensors in hierarchical tucker format. *Linear algebra and its applications*, 438(2):639–657, 2013.
- [5] Ricardo M. Baptista, Daniele Bigoni, Rebecca Morrison, and Alessio Spantini. TransportMaps, <http://transportmaps.mit.edu/docs/>, MIT Uncertainty Quantification Group, 2015-2018.
- [6] Daniele Bigoni, Olivier Zahm, Alessio Spantini, and Youssef Marzouk. Greedy inference with layers of lazy maps. *arXiv preprint arXiv:1906.00031*, 2019.
- [7] Peng Chen and Christoph Schwab. Sparse-grid, reduced-basis bayesian inversion: Nonaffine-parametric nonlinear equations. *Journal of Computational Physics*, 316:470–503, 2016.
- [8] Albert Cohen and Giovanni Migliorati. Optimal weighted least-squares methods. *The SMAI journal of computational mathematics*, 3:181–203, 2017.
- [9] Gaspare Da Fies and Marco Vianello. On the lebesgue constant of subperiodic trigonometric interpolation. *Journal of Approximation Theory*, 167:59–64, 2013.
- [10] Masoumeh Dashti and Andrew M Stuart. The bayesian approach to inverse problems. *Handbook of uncertainty quantification*, pages 1–118, 2016.
- [11] Gianluca Detommaso, Tiangang Cui, Youssef Marzouk, Alessio Spantini, and Robert Scheichl. A stein variational newton method. In *Advances in Neural Information Processing Systems*, pages 9169–9179, 2018.
- [12] Gianluca Detommaso, Jakob Kruse, Lynton Ardizzone, Carsten Rother, Ullrich Köthe, and Robert Scheichl. Hint: Hierarchical invertible neural transport for general and sequential bayesian inference. *arXiv preprint arXiv:1905.10687*, 2019.
- [13] TJ Dodwell, C Ketelsen, R Scheichl, and AL Teckentrup. Multilevel markov chain monte carlo. *Siam Review*, 61(3):509–545, 2019.
- [14] Sergey Dolgov, Karim Anaya-Izquierdo, Colin Fox, and Robert Scheichl. Approximation and sampling of multivariate probability distributions in the tensor train decomposition. *arXiv preprint arXiv:1810.01212*, 2018.
- [15] Martin Eigel, Claude Jeffrey Gittelsohn, Christoph Schwab, and Elmar Zander. Adaptive stochastic galerkin fem. *Computer Methods in Applied Mechanics and Engineering*, 270:247–269, 2014.
- [16] Martin Eigel, Robert Gruhlke, Manuel Marschall, and Elmar Zander. `alea` - A Python Framework for Spectral Methods and Low-Rank Approximations in Uncertainty Quantification, <https://bitbucket.org/aleadev/alea>.
- [17] Martin Eigel, Manuel Marschall, Max Pfeffer, and Reinhold Schneider. Adaptive stochastic galerkin fem for lognormal coefficients in hierarchical tensor representations. *arXiv preprint arXiv:1811.00319*, 2018.
- [18] Martin Eigel, Manuel Marschall, and Reinhold Schneider. Sampling-free bayesian inversion with adaptive hierarchical tensor representations. *Inverse Problems*, 34(3):035010, 2018.
- [19] Martin Eigel, Johannes Neumann, Reinhold Schneider, and Sebastian Wolf. Non-intrusive tensor reconstruction for high-dimensional random pdes. *Computational Methods in Applied Mathematics*, 19(1):39–53, 2019.
- [20] Martin Eigel, Max Pfeffer, and Reinhold Schneider. Adaptive stochastic galerkin fem with hierarchical tensor representations. *Numerische Mathematik*, 136(3):765–803, 2017.
- [21] Martin Eigel, Reinhold Schneider, Philipp Trunschke, and Sebastian Wolf. Variational monte carlo—bridging concepts of machine learning and high-dimensional partial differential equations. *Advances in Computational Mathematics*, 10 2019.
- [22] Tarek A El Moselhy and Youssef M Marzouk. Bayesian inference with optimal maps. *Journal of Computational Physics*, 231(23):7815–7850, 2012.
- [23] Oliver G Ernst, Björn Sprungk, and Lorenzo Tamellini. On expansions and nodes for sparse grid collocation of lognormal elliptic pdes. *arXiv preprint arXiv:1906.01252*, 2019.
- [24] Mike Espig, Lars Grasedyck, and Wolfgang Hackbusch. Black box low tensor-rank approximation using fiber-crosses. *Constructive approximation*, 30(3):557, 2009.
- [25] Jasmine Foo and George Em Karniadakis. Multi-element probabilistic collocation method in high dimensions. *Journal of Computational Physics*, 229(5):1536–1557, 2010.
- [26] Jochen Garcke and Michael Griebel. *Sparse grids and applications*, volume 88. Springer Science & Business Media, 2012.
- [27] Walter R Gilks, Sylvia Richardson, and David Spiegelhalter. *Markov chain Monte Carlo in practice*. Chapman and Hall/CRC, 1995.
- [28] Alex A Gorodetsky, Sertac Karaman, and Youssef M Marzouk. Function-train: A continuous analogue of the tensor-train decomposition. *arXiv preprint arXiv:1510.09088*, 2015.
- [29] Michael Griebel and Helmut Harbrecht. On the construction of sparse tensor product spaces. *Mathematics of computation*, 82(282):975–994, 2013.

- [30] Wolfgang Hackbusch. *Tensor spaces and numerical tensor calculus*, volume 42. Springer Science & Business Media, 2012.
- [31] Sebastian Holtz, Thorsten Rohwedder, and Reinhold Schneider. The alternating linear scheme for tensor optimization in the tensor train format. *SIAM Journal on Scientific Computing*, 34(2):A683–A713, 2012.
- [32] Benjamin Huber and Sebastian Wolf. Xerus - a general purpose tensor library. <https://libxerus.org/>, 2014–2017.
- [33] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160. Springer Science & Business Media, 2006.
- [34] Jinglai Li and Youssef M Marzouk. Adaptive construction of surrogates for the bayesian solution of inverse problems. *SIAM Journal on Scientific Computing*, 36(3):A1163–A1186, 2014.
- [35] Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances in neural information processing systems*, pages 2378–2386, 2016.
- [36] Youssef Marzouk, Tarek Moselhy, Matthew Parno, and Alessio Spantini. An introduction to sampling via measure transport. *arXiv preprint arXiv:1602.05023*, 2016.
- [37] KO Mead and LM Delves. On the convergence rate of generalized fourier expansions. *IMA Journal of Applied Mathematics*, 12(3):247–259, 1973.
- [38] Radford M Neal. Annealed importance sampling. *Statistics and computing*, 11(2):125–139, 2001.
- [39] Fabio Nobile, Raúl Tempone, and Clayton G Webster. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 46(5):2309–2345, 2008.
- [40] Ivan Oseledets and Eugene Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications*, 432(1):70–88, 2010.
- [41] Ivan V Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.
- [42] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- [43] Matthew Parno, Tarek Moselhy, and Youssef Marzouk. A multiscale strategy for bayesian inference using transport maps. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1160–1190, 2016.
- [44] Matthew D Parno and Youssef M Marzouk. Transport map accelerated markov chain monte carlo. *SIAM/ASA Journal on Uncertainty Quantification*, 6(2):645–682, 2018.
- [45] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [46] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- [47] Paul B Rohrbach, Sergey Dolgov, Lars Grasedyck, and Robert Scheichl. Rank bounds for approximating gaussian densities in the tensor-train format. *arXiv preprint arXiv:2001.08187*, 2020.
- [48] Daniel Rudolf and Björn Sprungk. Metropolis-hastings importance sampling estimator. *PAMM*, 17(1):731–734, 2017.
- [49] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser*, NY, 55:58–63, 2015.
- [50] Claudia Schillings and Christoph Schwab. Scaling limits in computational bayesian inversion. *ESAIM: Mathematical Modelling and Numerical Analysis*, 50(6):1825–1856, 2016.
- [51] Claudia Schillings, Björn Sprungk, and Philipp Wacker. On the convergence of the laplace approximation and noise-level-robustness of laplace-based monte carlo methods for bayesian inverse problems. *arXiv preprint arXiv:1901.03958*, 2019.
- [52] Reinhold Schneider and André Uschmajew. Approximation rates for the hierarchical tensor format in periodic sobolev spaces. *Journal of Complexity*, 30(2):56–71, 2014.
- [53] Christoph Schwab and Claude Jeffrey Gittelson. Sparse tensor discretizations of high-dimensional parametric and stochastic pdes. *Acta Numerica*, 20:291–467, 2011.
- [54] Andrew M Stuart. Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559, 2010.
- [55] Dustin Tran, Keyon Vafa, Kumar Krishna Agrawal, Laurent Dinh, and Ben Poole. Discrete flows: Invertible generative models of discrete data. *arXiv preprint arXiv:1905.10347*, 2019.
- [56] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- [57] Jonathan Weare. Efficient Monte Carlo sampling by parallel marginalization. *Proceedings of the National Academy of Sciences*, 104(31):12657–12662, 2007.