



Deep Composition of Tensor-Trains Using Squared Inverse Rosenblatt Transports

Tiangang Cui¹ · Sergey Dolgov²

Received: 15 July 2020 / Revised: 26 February 2021 / Accepted: 28 June 2021 /

Published online: 21 September 2021

© The Author(s) 2021

Abstract

Characterising intractable high-dimensional random variables is one of the fundamental challenges in stochastic computation. The recent surge of transport maps offers a mathematical foundation and new insights for tackling this challenge by coupling intractable random variables with tractable reference random variables. This paper generalises the functional tensor-train approximation of the inverse Rosenblatt transport recently developed by Dolgov et al. (Stat Comput 30:603–625, 2020) to a wide class of high-dimensional non-negative functions, such as unnormalised probability density functions. First, we extend the inverse Rosenblatt transform to enable the transport to general reference measures other than the uniform measure. We develop an efficient procedure to compute this transport from a squared tensor-train decomposition which preserves the monotonicity. More crucially, we integrate the proposed order-preserving functional tensor-train transport into a nested variable transformation framework inspired by the layered structure of deep neural networks. The resulting deep inverse Rosenblatt transport significantly expands the capability of tensor approximations and transport maps to random variables with complicated nonlinear interactions and concentrated density functions. We demonstrate the efficiency of the proposed approach on a range of applications in statistical learning and uncertainty quantification, including parameter estimation for dynamical systems and inverse problems constrained by partial differential equations.

Communicated by Endre Süli.

✉ Sergey Dolgov
s.dolgov@bath.ac.uk

Tiangang Cui
tiangang.cui@monash.edu

¹ School of Mathematics, Monash University, VIC 3800, Australia

² Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, UK

Keywords Tensor-train · Inverse problems · Uncertainty quantification · Rosenblatt transport · Deep transport maps

Mathematics Subject Classification 65D15 · 65D32 · 65C05 · 65C40 · 65C60 · 62F15 · 15A69 · 15A23

1 Introduction

Exploration of high-dimensional probability distributions is a fundamental task in statistical physics, machine learning, uncertainty quantification, econometrics, and beyond. In many practical scenarios, high-dimensional random variables of interest follow *intractable* probability measures that exhibit nonlinear interactions and concentrate in some sub-manifolds. This way, one cannot directly simulate the random variables of interest but may be able to evaluate the unnormalised density function pointwise.

Suppose we have an intractable target probability measure ν_π with the unnormalised density function $\pi(x)$ over a parameter space $\mathcal{X} \subseteq \mathbb{R}^d$, for example, the posterior measure in a Bayesian inference problem. Various approaches have been proposed to characterise ν_π using some reference probability measure μ defined over $\mathcal{U} \subseteq \mathbb{R}^d$, where independent and identically distributed random variables can be drawn from, e.g. a uniform or a Gaussian. For example, Markov chain Monte Carlo (MCMC) methods [36,53] generate a Markov chain of random variables converging to ν_π using μ as the proposal; and importance sampling and/or sequential Monte Carlo [32,48] characterise ν_π using weighted samples drawn from μ . The recently developed *transport map* idea, e.g. [4,17,39,43,50], offers new insights for this task by identifying a measurable mapping, $T : \mathcal{U} \mapsto \mathcal{X}$, such that the pushforward of μ , denoted by $T_\# \mu$, is a close approximation to ν_π . Then, the mapping T can be used to either accelerate classical sampling methods such as MCMC or to improve the efficiency of importance sampling. In this work, we generalise the tensor-train (TT) approach of [17] to offer an order-preserving and multi-layered construction of transport maps that is suitable for high-dimensional random variables with nonlinear interactions and concentrated density functions.

1.1 Outline and Contributions

The TT-based construction of [17] realises the mapping T via a separable TT decomposition [47] of the target density function. Since the separable tensor decomposition enables the marginalisation of the target density at a computational cost scaling linearly in the dimension of the random variables, it offers a computationally viable way to approximating marginal and conditional density functions of the target measure. In turn, the cumulative distribution functions (CDFs) of the marginals and condi-

tionals define the Rosenblatt transport¹ that can couple the target measure with the uniform reference measure. Section 2 presents the relevant background of the Rosenblatt transport, the functional form of the TT decomposition of multivariate functions [2,25,26,28], and the TT-based construction of the inverse Rosenblatt transport.

The TT-based construction faces several challenges. First, the TT decomposition of the non-negative target density function often cannot preserve the non-negativity after rank truncations. This way, the resulting Rosenblatt transport may not preserve the monotonicity. Second, TT decomposition works best when the correlations between random variables are *local*, i.e. the correlation decays with the distance between the indices of the variables. In the extreme case of independent random variables, the joint density factorises into the product of marginal densities. However, high-dimensional random variables of interest often have concentrated density functions and exhibit complicated nonlinear interactions. In such cases, one may need a TT with high ranks to approximate the target probability density with sufficient accuracy, which in turn requires a rather large number of target density evaluations during the TT construction.

In Sect. 3, we overcome the first challenge by proposing a new construction of inverse Rosenblatt transport by approximating the squared root of the target density in the TT format, followed by constructing the marginal and conditional densities from the square of the TT approximation. The resulting *squared inverse Rosenblatt transport* (SIRT) is order- and smoothness-preserving. In addition, utilising the squared structure of the approximation, we also establish error bounds of SIRT in terms of various statistical divergences within the f -divergence family. These bounds are useful for bounding the statistical efficiency of posterior characterisation algorithms such as MCMC and importance sampling.

In Sect. 4, we circumvent the second challenge by introducing a multi-layer *deep inverse Rosenblatt transport* (DIRT) that builds a composition of SIRTs guided by a sequence of bridging measures with increasing complexity. We illustrate this idea in Fig. 1. At each layer of DIRT, we aim to obtain a composition of SIRTs, denoted by $T_0 \circ T_1 \circ \dots \circ T_k$, such that the pushforward of the reference measure under this composition is a close approximation of the k th bridging measure ν_k . The existing composition $T_0 \circ T_1 \circ \dots \circ T_k$ offers a nonlinear transformation of coordinates that can effectively capture the correlations and support of the next bridging measure ν_{k+1} . As a result, the density of the pullback measure, $(T_0 \circ T_1 \circ \dots \circ T_k)_\# \nu_{k+1}$, can have a much simpler structure for building the TT decomposition compared with the density of ν_{k+1} . We can then factorise the density of $(T_0 \circ T_1 \circ \dots \circ T_k)_\# \nu_{k+1}$ to define the incremental mapping T_{k+1} such that $(T_{k+1})_\# \mu = (T_0 \circ T_1 \circ \dots \circ T_k)_\# \nu_{k+1}$. This way, DIRT is capable of characterising random variables with concentrated density functions by factorising a sequence of less complicated density functions in transformed coordinates. To further improve the efficiency, we also present strategies that can embed *general reference measures* rather than the uniform reference measure to avoid complicated boundary layers during DIRT construction. Moreover, we can show that the DIRT construction is robust to TT approximation errors in various statistical divergences, in the sense

¹ This is also referred to as the Knothe–Rosenblatt rearrangement. It was independently proposed by Rosenblatt [55] for statistical purposes and by Knothe [33] for proving the isoperimetric inequality. The setup of the TT-based approach closely follows the work of Rosenblatt.

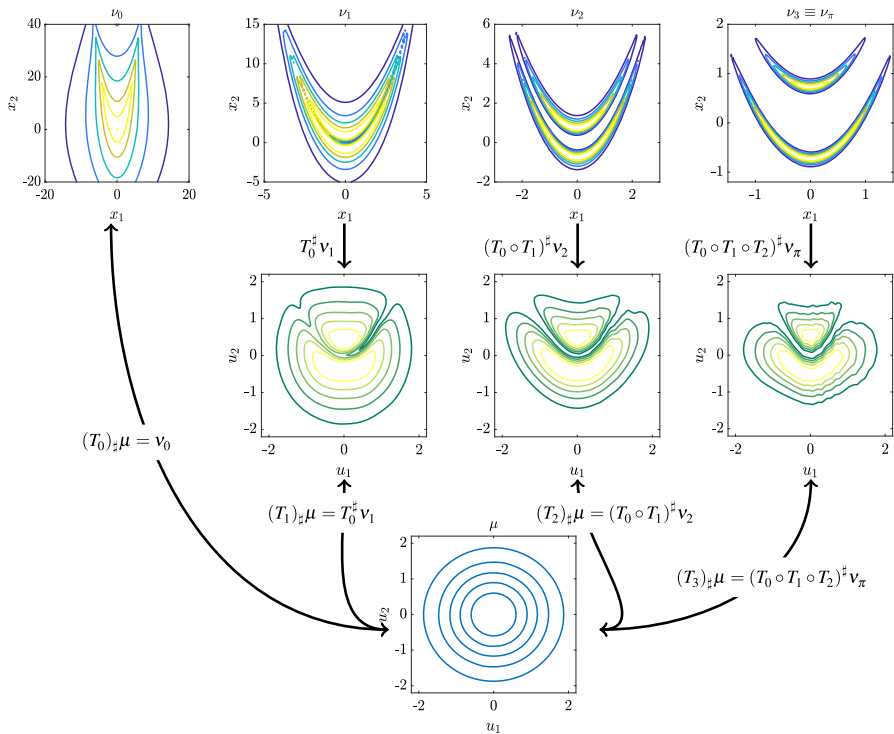


Fig. 1 Illustration of DIRT. Top row shows a sequence of bridging measures towards the target measure ν_π . Each layer of DIRT identifies an incremental mapping that couples the reference measure (bottom row) with the pullback of the bridging measure under existing composition of mappings (middle row), which admits a simpler structure for constructing TT decomposition

that the error bounds on a range of divergences is a linear combination of errors of TT decompositions involved in the DIRT construction process.

In Sect. 5, we integrate SIRT and DIRT into existing MCMC and importance sampling methods to further reduce the estimation and sampling bias due to TT approximation errors. In Sect. 6, we demonstrate the efficiency and various aspects of DIRT on several Bayesian inverse problems governed by ordinary differential equations (ODEs) and partial differential equations (PDEs). Using a predator–prey dynamical system (Sect. 6.1), we benchmark the impact of various tuning parameters of the functional TT decomposition such as the TT rank, the number of collocation points and the choice of the reference measure on the accuracy of the DIRT. Using an elliptic PDE (Sect. 6.3), we are able to compare the single-layered SIRT with DIRT, in which DIRT shows a clear advantage in both the computational efficiency and the accuracy over the single-layered counterpart. In the same example, we also demonstrate the efficiency of TT with the Fourier basis compared to that with the piecewise-linear basis on concentrated posterior measures due to increasing number of measurements and decreasing measurement noises. Furthermore, we can vary the discretisation of the underlying ODE or PDE models from layer to layer to accelerate the DIRT construc-

tion. For an example involving a computationally expensive parabolic PDE (Sect. 6.4), we employ models with increasingly refined grids to construct DIRT that is otherwise computationally infeasible to build.

1.2 Related Work

Apart from building transport maps using TT decompositions, most of other methods approximate the transport map T by solving an optimisation problem such that T minimises some statistical divergence between the target ν_π and the pushforward $T_{\#} \mu$. The mapping T often has a triangular structure, which is computationally efficient for evaluating the Jacobian and the inverse of T and can be represented using polynomials [4,43,50,51], kernel functions [15,37], invertible neural networks [7,9,10,35,49,52], etc. In this setting, the objective function has to be approximated using a Monte Carlo average and minimised by some (stochastic) gradient-based method. Depending on the objective function and how samples are obtained, the resulting methods may have very different structures.

- *Density approximation.* The work of [4,43,51] aims to minimise the Kullback–Leibler (KL) divergence of the pushforward $T_{\#} \mu$ from the target ν_π , in which the pushforward density naturally approximates the target density. In this case, the KL divergence is approximated using the Jacobian of T and the target density function evaluated at samples drawn from the analytically tractable reference measure. The resulting optimisation problem may be highly nonlinear and non-convex. In each optimisation iteration, the target density function has to be re-evaluated as reference samples are transported by the updated map. Our TT-based methods also rely on approximations to the target density. However, TT approximations employ highly efficient deterministic sampling algorithms such as TT-Cross [46], which are free from either gradient or Monte Carlo. This way, TT-based methods may need less number of density evaluations to accurately approximate the target density.
- *Density estimation.* The strategy adopted by normalising flows (e.g. [7,9,10,35,49,52]) and the work of [50,61,63] offer an alternative that can bypass evaluations of the target density. Instead, these methods assume availability of samples drawn from the target measure and construct objective functions using a given set of target samples. Many of these methods, particularly neural networks, were originally designed to approximate high-dimensional distributions of naturally available samples, such as images. However, in our context, the intractable target random variable X cannot be simulated directly. One has to assume that there exists an auxiliary random variable Y such that the density function of X is given by a conditional density $\pi(x|y)$ and the pair of joint random variables (X, Y) can be simulated directly. This way, density estimations can be employed to first build a mapping from some higher dimensional reference measure to the joint measure of (X, Y) and then obtain the mapping T by conditioning on a particular realisation of $Y = y$. We provide a concrete example of normalising flows and its comparisons with DIRT in Sect. 6.1.
- *Greedy methods.* In-between, the fully data-driven density estimation and the function-driven density approximation are the greedy strategy, including the Stein

variational gradient descent method [37], its Newton variant [15], and the lazy maps [4]. While greedy methods build transport maps sharing a similar composition structure with DIRT, they obtain the composition of mappings by iteratively minimising the KL divergence of the pushforward measure under the current composition of maps from the target ν_π . To relax the burden in optimisation, the class of mappings used in each layer of greedy methods is often restricted, for example, to reproducing kernel Hilbert space with Gaussian kernels [15,37] and sparse low-order polynomials [4]. As a result, greedy methods often need a rather large number of layers to accurately approximate concentrated target densities and hence, may lead to a large number of computationally costly target density evaluations. Compared to the greedy strategy, the usage of bridging measures allows DIRT to construct TT decompositions in different layers with arbitrary accuracy. The total error in DIRT is also accumulated linearly with the number of layers. We provide a numerical comparison on the performance of DIRT and the Stein variational Newton methods [15] in Sect. 6.1.

2 Background

In this section, we first introduce some notation and assumptions used throughout the paper. Then, we review the inverse Rosenblatt transport method that offers an algebraically exact transformation from the reference measure to the target measure. We will also discuss the role of the functional TT decomposition in the numerical construction of the (approximate) inverse Rosenblatt transport.

2.1 Notation

We consider probability measures that are absolutely continuous with respect to the Lebesgue measure. Suppose a mapping $S : \mathcal{X} \mapsto \mathcal{U}$ is a diffeomorphism and a probability measure ν has a density $p(x)$, the pushforward of ν under S , denoted by $S_\# \nu$, has the density:

$$S_\# p(u) = (p \circ S^{-1})(u) |\nabla_u S^{-1}(u)|. \quad (1)$$

Similarly, given a probability measure λ with a density $q(u)$, the pullback of λ under S , denoted by $S^\# \lambda$, has the density:

$$S^\# q(x) = (q \circ S)(x) |\nabla_x S(x)|. \quad (2)$$

The short hand $X \sim \nu$ is used to refer a random variable X with the law ν . For a ν -integrable function $q : \mathcal{X} \mapsto \mathbb{R}$, the expectation of q is denoted by $\nu(q) = \int q(x) \nu(dx)$.

We assume the parameter space \mathcal{X} , and the reference space \mathcal{U} can be expressed as Cartesian products $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \cdots \times \mathcal{X}_d$ and $\mathcal{U} = \mathcal{U}_1 \times \mathcal{U}_2 \times \cdots \times \mathcal{U}_d$, respectively, where $\mathcal{X}_k \subseteq \mathbb{R}$ and $\mathcal{U}_k = [0, 1]$. Using product-form Lebesgue measurable weighting

functions $\lambda(x) = \prod_{i=1}^d \lambda_i(x_i)$ and $\omega(u) = \prod_{i=1}^d \omega_i(x_i)$, the weighted L^p norms on \mathcal{X} and \mathcal{U} can be expressed as

$$\|f\|_{L^p_\lambda(\mathcal{X})} = \left(\int_{\mathcal{X}} |f(x)|^p \lambda(x) dx \right)^{\frac{1}{p}} \quad \text{and} \quad \|g\|_{L^p_\omega(\mathcal{U})} = \left(\int_{\mathcal{U}} |g(u)|^p \omega(u) du \right)^{\frac{1}{p}},$$

respectively. We define constants $\lambda_i(\mathcal{X}_i) = \int_{\mathcal{X}_i} \lambda_i(x_i) dx_i$ for $i = 1, \dots, d$ and $\lambda(\mathcal{X}) = \prod_{i=1}^d \lambda_i(\mathcal{X}_i)$. Likewise, we also define $\omega_i(\mathcal{U}_i)$ for $i = 1, \dots, d$ and $\omega(\mathcal{U})$.

For a vector $x \in \mathbb{R}^d$ and an index $k \in \mathbb{N}$ such that $1 < k < d$, we express the first $k - 1$ coordinates and the last $d - k$ coordinates of x as

$$x_{<k} \equiv [x_1, \dots, x_{k-1}]^\top, \quad \text{and} \quad x_{>k} \equiv [x_{k+1}, \dots, x_d]^\top,$$

respectively. Similarly, we write $x_{\leq k} = (x_{<k}, x_k)$, $x_{\geq k} = (x_k, x_{>k})$, $x_{\leq 1} = x_1$, and $x_{\geq d} = x_d$. For any non-negative function $\pi \in L^1_{\lambda}(\mathcal{X})$, we define its marginal functions as

$$\pi_{\leq k}(x_{\leq k}) \equiv \int \pi(x_{\leq k}, x_{>k}) \left(\prod_{i=k+1}^d \lambda_i(x_i) \right) dx_{>k}, \quad \text{for } 1 \leq k < d, \quad (3)$$

with $\pi_{\leq d}(x_{\leq d}) = \pi(x)$. The marginal functions should not be confounded with $\pi_1(x)$, $\pi_2(x)$, ..., $\pi_k(x)$ where the subscript indexes a sequence of functions on \mathcal{X} .

2.2 Inverse Rosenblatt Transport

We start with a d -dimensional uniform reference probability measure, μ_{uni} , defined in a unit hypercube $\mathcal{U} = [0, 1]^d$, which has the probability density function (PDF) $f_{\text{uni}}(u) = 1$. We aim to characterise a target probability measure ν_π with the PDF

$$f_X(x) = \frac{1}{z} \pi(x) \lambda(x) \quad \text{and} \quad z = \int_{\mathcal{X}} \pi(x) \lambda(x) dx. \quad (4)$$

Here, $\pi \in L^1_{\lambda}(\mathcal{X})$ is the unnormalised density function (with respect to the weight λ) that is non-negative, i.e. $\pi(x) \geq 0, \forall x \in \mathcal{X}$, and z is the normalising constant that is often unknown.

Let $X := (X_1, \dots, X_d)$ be the target d -dimensional random variable with law ν_π and U be the reference d -dimensional random variable with law μ_{uni} . The *Rosenblatt transport* offers a viable way to constructing a map $F : \mathbb{R}^d \mapsto \mathbb{R}^d$ such that $F(X) = U$. As explained in [6,58,64], the principle of the Rosenblatt transport is the following. For $1 \leq k \leq d$, we denote the marginal PDF of the k -dimensional random variable $X_{\leq k} := (X_1, \dots, X_k)$ by

$$f_{X_{\leq k}}(x_{\leq k}) = \frac{1}{z} \pi_{\leq k}(x_{\leq k}) \left(\prod_{i=1}^k \lambda_i(x_i) \right)$$

and the PDF of the conditional random variable $X_k|X_{<k}$ by

$$f_{X_k|X_{<k}}(x_k|x_{<k}) = \frac{f_{X_{\leq k}}(x_{<k}, x_k)}{f_{X_{<k}}(x_{<k})} = \frac{\pi_{<k}(x_{<k}, x_k)}{\pi_{<k}(x_{<k})} \lambda_k(x_k).$$

This way, the CDF of X_1 and the conditional CDF of $X_k|X_{<k}$ can be expressed as

$$F_{X_1}(x_1) = \int_{-\infty}^{x_1} f_{X_1}(x_1) dx_1 \quad \text{and} \quad F_{X_k|X_{<k}}(x_k|x_{<k}) = \int_{-\infty}^{x_k} f_{X_k|X_{<k}}(x_k|x_{<k}) dx_k, \tag{5}$$

respectively. Under mild assumptions [6], the following sequence of transformations

$$\begin{cases} u_1 &= \mathbb{P}[X_1 \leq x_1] &= F_{X_1}(x_1) \\ &\vdots \\ u_k &= \mathbb{P}[X_k \leq x_k | X_{<k} = x_{<k}] &= F_{X_k|X_{<k}}(x_k|x_{<k}) \\ &\vdots \\ u_d &= \mathbb{P}[X_d \leq x_d | X_{<d} = x_{<d}] &= F_{X_d|X_{<d}}(x_d|x_{<d}) \end{cases} \tag{6}$$

defines uniquely a monotonically increasing map $F : \mathcal{X} \mapsto \mathcal{U}$ in the form of

$$F(x) = [F_{X_1}(x_1), \dots, F_{X_k|X_{<k}}(x_k|x_{<k}), \dots, F_{X_d|X_{<d}}(x_d|x_{<d})]^\top, \tag{7}$$

such that the random variable $U = F(X)$ is uniformly distributed in the unit hypercube $[0, 1]^d$. Since the k th component of F is a scalar-valued function depending on only the first k variables, that is, $F_{X_k|X_{<k}} : \mathbb{R}^k \mapsto \mathbb{R}$, the map F has a *lower-triangular* form. Furthermore, the map F (as well as its inverse) is almost surely differentiable and satisfies

$$F^\# f_{\text{uni}}(x) = (f_{\text{uni}} \circ F)(x) |\nabla_x F(x)| = |\nabla_x F(x)| = f_X(x)$$

ν_π -almost surely.

Suppose one can compute the Rosenblatt transport. Then, it provides a viable way to characterising the target measure. One can first generate uniform random variables $U \sim \mu_{\text{uni}}$ and then applying the *inverse Rosenblatt transport* (IRT)

$$X = F^{-1}(U)$$

to obtain a corresponding target random variable $X \sim \nu_\pi$. The inverse Rosenblatt transport $T \equiv F^{-1} : \mathcal{U} \mapsto \mathcal{X}$ is also *lower-triangular* and can be constructed by successively inverting the Rosenblatt transport for $k = 1, \dots, d$:

$$x = T(u) \equiv [F_{X_1}^{-1}(u_1), \dots, F_{X_k|X_{<k}}^{-1}(u_k|x_{<k}), \dots, F_{X_d|X_{<d}}^{-1}(u_d|x_{<d})]^\top. \tag{8}$$

The evaluation of each $F_{X_k|X_{<k}}^{-1}(u_k|x_{<k})$ requires inverting only a scalar-valued monotone function $u_k = F_{X_k|X_{<k}}(x_k|x_{<k})$, where $x_{<k}$ is already determined in the first $k - 1$

steps. Using the change-of-variables formula, the expectation of a function $h : \mathcal{X} \mapsto \mathbb{R}$ can be expressed as

$$v_\pi(h) = \mu_{\text{uni}}(h \circ T).$$

This way, the expectation over the intractable target probability measure can be expressed as the expectation over a reference uniform probability measure, and thus many efficient high-dimensional quadrature methods such as sparse grids [5] and quasi-Monte Carlo [16] may apply.

2.3 Functional Tensor-Train

For high-dimensional target measures, it may be not computationally feasible to compute the marginal densities $\pi_{\leq k}$, and hence, the marginal and the conditional CDFs in (5) for building the inverse Rosenblatt transport. To overcome this challenge, a recent work [17] employed the TT decomposition [47] to factorise the density of the target measure in a separable form, which leads to a computationally scalable method for building the inverse Rosenblatt transport. Here, we first discuss the basics of the TT decomposition of a multivariate function.

Since multivariate functions can be viewed as continuous analogues of tensors [28], one can factorise the unnormalised density function using functional-form of TT [2,25,26]. Given a multivariate function $h : \mathcal{X} \mapsto \mathbb{R}$, where $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2 \times \dots \times \mathcal{X}_d$, TT approximates $h(x)$ as

$$h(x) \approx \tilde{h}(x) \equiv \sum_{\alpha_0=1}^{r_0} \sum_{\alpha_1=1}^{r_1} \dots \sum_{\alpha_d=1}^{r_d} H_1^{(\alpha_0, \alpha_1)}(x_1) \dots H_k^{(\alpha_{k-1}, \alpha_k)}(x_k) \dots H_d^{(\alpha_{d-1}, \alpha_d)}(x_d), \tag{9}$$

with $r_0 = r_d = 1$, where the summation ranges r_0, r_1, \dots, r_d are called *TT ranks*. Each univariate function $H_k^{(\alpha_{k-1}, \alpha_k)}(x_k) : \mathcal{X}_k \mapsto \mathbb{R}$ is represented as a linear combination of a set of n_k basis functions $\{\phi_k^{(1)}(x_k), \dots, \phi_k^{(n_k)}(x_k)\}$. This way, we have

$$H_k^{(\alpha_{k-1}, \alpha_k)}(x_k) = \sum_{i=1}^{n_k} \phi_k^{(i)}(x_k) \mathbf{A}_k[\alpha_{k-1}, i, \alpha_k], \tag{10}$$

where $\mathbf{A}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ is a coefficient tensor. Examples of the basis functions include piecewise polynomials, orthogonal functions, radial basis functions, etc. In general, the TT decomposition $\tilde{h}(x)$ is only an approximation to the original function $h(x)$ because of truncated TT ranks and sets of basis functions used for representing each $H_k^{(\alpha_{k-1}, \alpha_k)}(x_k)$.

Remark 1 For each k , grouping all the univariate functions $H_k^{(\alpha_{k-1}, \alpha_k)}(x_k)$, we have a matrix-valued function $H_k(x_k) : \mathcal{X}_k \mapsto \mathbb{R}^{r_{k-1} \times r_k}$ that is commonly referred to as the *k-th TT core*. This way, the TT decomposition can also be expressed in the matrix

form

$$\tilde{h}(x) = H_1(x_1) \cdots H_k(x_k) \cdots H_d(x_d). \tag{11}$$

We follow the MATLAB notation to denote vector-valued functions consisting of the α_k th column and α_{k-1} th row of $H_k(x_k)$ by $H_k^{(:,\alpha_k)}(x_k) : \mathcal{X}_k \mapsto \mathbb{R}^{r_{k-1} \times 1}$ and $H_k^{(\alpha_{k-1}, :)}(x_k) : \mathcal{X}_k \mapsto \mathbb{R}^{1 \times r_k}$, respectively. In some situations, it is convenient to represent the TT decomposition with grouped coordinates. For example, we can write TT as the functional analogue of the compact singular value decomposition (SVD):

$$\tilde{h}(x) = \sum_{\alpha_k=1}^{r_k} H_{\leq k}^{(\alpha_k)}(x_{\leq k}) H_{> k}^{(\alpha_k)}(x_{> k}), \tag{12}$$

where

$$H_{\leq k}^{(\alpha_k)}(x_{\leq k}) = H_1(x_1) \cdots H_{k-1}(x_{k-1}) H_k^{(:,\alpha_k)}(x_k) : \mathcal{X}_1 \times \cdots \times \mathcal{X}_k \mapsto \mathbb{R}, \tag{13}$$

$$H_{> k}^{(\alpha_k)}(x_{> k}) = H_{k+1}^{(\alpha_k, :)}(x_{k+1}) H_{k+2}(x_{k+2}) \cdots H_d(x_d) : \mathcal{X}_{k+1} \times \cdots \times \mathcal{X}_d \mapsto \mathbb{R}. \tag{14}$$

Given a multivariate function, its TT decomposition can be computed using alternating linear schemes such as the classical alternating least squares method (e.g. [34,46]), density matrix renormalization group methods [30,47,65], and the alternating minimal energy method [18] together with the cross approximation [22–24,38] or the empirical interpolation [1,8]. In Appendix 1, we detail the cross algorithm used for constructing the functional TT decomposition.

2.4 A TT-Based Inverse Rosenblatt Transport

Using the functional TT previously discussed, now we review the TT-based construction of the inverse Rosenblatt transport [17]. Suppose one has the (approximate) TT decomposition $\tilde{\pi}(x)$ of the unnormalised target density $\pi(x)$ in the form of

$$\tilde{\pi}(x_1, x_2, \dots, x_d) = F_1(x_1) \cdots F_k(x_k) \cdots F_d(x_d),$$

where $F_k(x_k) : \mathcal{X}_k \mapsto \mathbb{R}^{r_{k-1} \times r_k}$ is the k th TT core. Then, we can approximate the target PDF by

$$\tilde{f}_X(x) = \frac{1}{\tilde{c}} \tilde{\pi}(x) \lambda(x), \quad \text{where } \tilde{c} = \int_{\mathcal{X}} \tilde{\pi}(x) \lambda(x) dx. \tag{15}$$

Proposition 1 For $k < d$, the k th marginal PDF is given by

$$\tilde{f}_{X_{\leq k}}(x_{\leq k}) = \frac{1}{\tilde{c}} \tilde{\pi}_{\leq k}(x_{\leq k}) \left(\prod_{i=1}^k \lambda_i(x_i) \right),$$

where $\tilde{\pi}_{\leq k}(x_{\leq k}) = F_1(x_1) \cdots F_k(x_k) \bar{F}_{k+1} \cdots \bar{F}_d$, $\tilde{c} = \bar{F}_1 \cdots \bar{F}_d$, and the matrices \bar{F}_k are the integrated TT cores

$$\bar{F}_k = \int_{\mathcal{X}_k} F_k(x_k) \lambda_k(x_k) dx_k \in \mathbb{R}^{r_{k-1} \times r_k}, \quad \text{for } k = 1, \dots, d.$$

Proof The marginal function of $\tilde{\pi}(x)$ can be expressed by

$$\tilde{\pi}_{\leq k}(x_{\leq k}) = \int_{\mathcal{X}_{>k}} F_1(x_1) \cdots F_d(x_d) \left(\prod_{i=k+1}^d \lambda_i(x_i) \right) dx_{>k}.$$

Using the separable form of the tensor–train, the marginal density then has the form

$$\begin{aligned} \tilde{\pi}_{\leq k}(x_{\leq k}) &= F_1(x_1) \cdots F_k(x_k) \left(\int_{\mathcal{X}_{k+1}} F_{k+1}(x_{k+1}) \lambda_{k+1}(x_{k+1}) dx_{k+1} \right) \cdots \left(\int_{\mathcal{X}_d} F_d(x_d) \lambda_d(x_d) dx_d \right) \\ &= F_1(x_1) \cdots F_k(x_k) \bar{F}_{k+1} \cdots \bar{F}_d. \end{aligned}$$

Since $\tilde{c} = \int_{\mathcal{X}} \tilde{f}(x) \lambda(x) dx$, we have $\tilde{c} = \bar{F}_1 \cdots \bar{F}_d$ using a similar argument. □

The above proposition leads to the marginal PDF $\tilde{f}_{X_1} = \frac{1}{\tilde{c}} \tilde{\pi}_{\leq 1}(x_1) \lambda_1(x_1)$ and the sequence of conditional probability densities

$$\tilde{f}_{X_k|X_{<k}}(x_k|x_{<k}) = \frac{\tilde{\pi}_{\leq k}(x_{<k}, x_k)}{\tilde{\pi}_{<k}(x_{<k})} \lambda_k(x_k), \quad k = 2, \dots, d. \tag{16}$$

This leads to the CDF and the sequence of conditional CDFs

$$\tilde{F}_{X_1}(x_1) = \int_{-\infty}^{x_1} \tilde{f}_{X_1} dx'_1 \quad \text{and} \quad \tilde{F}_{X_k|X_{<k}}(x_k|x_{<k}) = \int_{-\infty}^{x_k} \tilde{f}_{X_k|X_{<k}}(x_k|x_{<k}) dx'_k, \tag{17}$$

for $k = 2, \dots, d$, and hence, the Rosenblatt transport $U = \tilde{F}(X)$. This equivalently defines the inverse Rosenblatt transport $\tilde{T} = \tilde{F}^{-1}$. This way, by drawing a reference random variable $U \sim \mu_{\text{uni}}$ and evaluating $X = \tilde{T}(U)$, we obtain an approximate target random variable $X \sim \tilde{T}_\# \mu_{\text{uni}}$. Note that the pushforward measure $\tilde{T}_\# \mu_{\text{uni}}$ has the density $\tilde{f}_X(x)$.

To estimate the numerical complexity, let us introduce the maximal number of basis functions $n = \max_{k=1, \dots, d} n_k$, TT rank $r = \max_{k=0, \dots, d} r_k$, and suppose we need to draw N samples from $\tilde{\pi}(x)$. Note that we can precompute $\bar{F}_{k+1} \cdots \bar{F}_d$ with the total cost of $\mathcal{O}(dnr^2)$ operations, before any sampling starts. Similarly, the conditioning requires the interpolation of $F_1(x_1) \cdots F_{k-1}(x_{k-1})$ at the current sample coordinates, which can be built up sequentially. Each univariate interpolation needs $\mathcal{O}(nr^2)$ operations in general, but for a piecewise interpolation this can be reduced to $\mathcal{O}(r^2)$ operations per sample per coordinate. Finally, the assembling of the conditional density requires the multiplication of N vectors $F_1(x_1) \cdots F_{k-1}(x_{k-1}) \in \mathbb{R}^{r_{k-1}}$ with

a vector-valued function $F_k(x_k)\bar{F}_{k+1} \cdots \bar{F}_d \in \mathbb{R}^{r_k-1}$. The total complexity is therefore $\mathcal{O}(dnr^2 + Ndr^2 + Ndnr)$ [17].

Constructing the inverse Rosenblatt transport using the TT decomposition of the target density faces several challenges. First, the density function $\pi(x)$ is non-negative; however, its truncated TT decomposition $\tilde{\pi}(x)$ can have negative values—a discrete analogue is that the truncated SVD of a matrix filled with non-negative entries can be negative. This leads to a critical issue: if the set $\{x \in \mathcal{X} \mid \tilde{\pi}(x) < 0\}$ has nonzero measure under ν_π , then the Rosenblatt transport constructed from $\tilde{\pi}(x)$ loses monotonicity. A simple way to circumvent this is to take the modulus of each univariate conditional density $\tilde{f}_{X_k|X_{<k}}(x_k|x_{<k})$ and then renormalise the modulus before computing the CDF [17]. However, the use of moduli and renormalisations may degrade the smoothness of marginal PDFs and conditional PDFs. This way, the resulting inverse Rosenblatt transport and its induced PDF can lose accuracy and smoothness. More importantly, the construction of the TT decomposition (see Sect. 1 for details) requires evaluating the target density at parameter points where the target density is significant. In practice, the high probability region of a high-dimensional target density, e.g. the posterior in the Bayesian inference context, can be hard to characterise. Thus, it can be challenging to construct the TT decomposition for approximating the target density directly. In the next section, we generalise the TT-based construction of the inverse Rosenblatt transport by tackling the aforementioned challenges.

3 Squared Inverse Rosenblatt Transport

We first introduce the SIRT to overcome the negativity issue outlined above. Instead of directly decomposing the unnormalised target density $\pi(x)$, we first obtain the (approximate) functional TT decomposition $\tilde{g}(x)$ of the square root of $\pi(x)$ in the form of

$$\sqrt{\pi}(x) \approx \tilde{g}(x) = G_1(x_1) \cdots G_k(x_k) \cdots G_d(x_d), \tag{18}$$

where $G_k(x_k) : \mathcal{X}_k \mapsto \mathbb{R}^{r_{k-1} \times r_k}$ is the k th TT core. This leads to an alternative approximation to the target PDF:

$$\hat{f}_{\hat{X}}(x) = \frac{1}{\hat{z}} \hat{\pi}(x) \lambda(x), \text{ with } \hat{\pi}(x) = \gamma + \tilde{g}(x)^2 \text{ and } \hat{z} = \gamma \lambda(\mathcal{X}) + \int_{\mathcal{X}} \tilde{g}(x)^2 \lambda(x) dx, \tag{19}$$

where $\gamma > 0$ is a constant chosen according to the L^2 error of $\tilde{g}(x)$. Similar to the process discussed in Sect. 2.4, we can obtain the SIRT, $\hat{X} = \hat{T}(U)$, by constructing the sequence of marginal functions $\hat{\pi}_{\leq k}(x_{\leq k}) = \int_{\mathcal{X}} \hat{\pi}(x) \prod_{i=k+1}^d \lambda_i(x_i) dx_{>k}$ for $k = 1, \dots, d - 1$ and computing the normalising constant \hat{z} . Given a reference random variable $U \sim \mu_{\text{uni}}$, we can evaluate $\hat{X} = \hat{T}(U)$ to obtain an approximate target random variable $\hat{X} \sim \hat{T}_{\hat{z}} \mu_{\text{uni}}$, which has exactly the PDF $\hat{f}_{\hat{X}}(x)$. Since the function $\hat{\pi}(x)$ is positive by construction, we can *preserve the smoothness and monotonicity* in the resulting SIRT \hat{T} .

Remark 2 For a target density $\pi(x)$ satisfying $\sup_{x \in \mathcal{X}} \pi(x) < \infty$, the ratio between $\pi(x)$ and the approximate density $\hat{\pi}(x)$ satisfies

$$\sup_{x \in \mathcal{X}} \frac{\pi(x)}{\hat{\pi}(x)} = \sup_{x \in \mathcal{X}} \frac{\pi(x)}{\gamma + \tilde{g}(x)^2} = \hat{c} < \infty. \tag{20}$$

This bound is essential to ensure the uniform ergodicity of the Metropolis independent algorithm and the rate of convergence of importance sampling schemes defined by SIRT. See Sect. 5 for further details.

3.1 Marginal Functions and Conditional PDFs

We represent each TT core of the decomposition in (18) as

$$\mathbf{G}_k^{(\alpha_{k-1}, \alpha_k)}(x_k) = \sum_{i=1}^{n_k} \phi_k^{(i)}(x_k) \mathbf{A}_k[\alpha_{k-1}, i, \alpha_k], \quad k = 1, \dots, d, \tag{21}$$

where $\{\phi_k^{(i)}(x_k)\}_{i=1}^{n_k}$ is the set of basis functions for the k th coordinate and $\mathbf{A}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ is the associated k th coefficient tensor. For the k th set of basis functions, we define the mass matrix $\mathbf{M}_k \in \mathbb{R}^{n_k \times n_k}$ by

$$\mathbf{M}_k[i, j] = \int_{\mathcal{X}_k} \phi_k^{(i)}(x_k) \phi_k^{(j)}(x_k) \lambda(x_k) dx_k, \text{ for } i = 1, \dots, n_k, j = 1, \dots, n_k. \tag{22}$$

Then, we can represent the marginal functions by

$$\hat{\pi}_1(x_1) = \gamma \prod_{i=2}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_1=1}^{r_1} \left(\mathbf{P}_1^{(\alpha_0, \ell_1)}(x_1) \right)^2, \tag{23}$$

$$\hat{\pi}_{\leq k}(x_{\leq k}) = \gamma \prod_{i=k+1}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_k=1}^{r_k} \left(\sum_{\alpha_{k-1}=1}^{r_{k-1}} \mathbf{G}_{< k}^{(\alpha_{k-1})}(x_{< k}) \mathbf{P}_k^{(\alpha_{k-1}, \ell_k)}(x_k) \right)^2, \quad k = 2, \dots, d, \tag{24}$$

where $\alpha_0 = 1$ and

$$\mathbf{G}_{< k}^{(\alpha_{k-1})}(x_{< k}) = \mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-1}^{(\cdot, \alpha_{k-1})}(x_{k-1}) : \mathcal{X}_{< k} \mapsto \mathbb{R}, \tag{25}$$

$$\mathbf{P}_k^{(\alpha_{k-1}, \ell_k)}(x_k) = \sum_{i=1}^{n_k} \phi_k^{(i)}(x_k) \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] : \mathcal{X}_k \mapsto \mathbb{R}, \tag{26}$$

for a coefficient tensor $\mathbf{B}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ that is recursively defined as follows.

Proposition 2 Starting with the last coordinate $k = d$, we set $\mathbf{B}_d = \mathbf{A}_d$. Suppose for the first k dimensions ($k > 1$), we have a coefficient tensor $\mathbf{B}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ that defines a marginal function $\hat{\pi}_{\leq k}(x_{\leq k})$ as in (24). The following procedure can be used to obtain the coefficient tensor $\mathbf{B}_{k-1} \in \mathbb{R}^{r_{k-2} \times n_{k-1} \times r_{k-1}}$ for defining the next marginal function $\hat{\pi}_{< k}(x_{< k})$:

1. Use the Cholesky decomposition of the mass matrix, $L_k L_k^\top = M_k \in \mathbb{R}^{n_k \times n_k}$, to construct a tensor $\mathbf{C}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$:

$$\mathbf{C}_k[\alpha_{k-1}, \tau, \ell_k] = \sum_{i=1}^{n_k} \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] L_k[i, \tau]. \tag{27}$$

2. Unfold \mathbf{C}_k along the first coordinate [34] to obtain a matrix $\mathbf{C}_k^{(R)} \in \mathbb{R}^{r_{k-1} \times (n_k r_k)}$ and compute the thin QR decomposition

$$Q_k R_k = (\mathbf{C}_k^{(R)})^\top, \tag{28}$$

where $Q_k \in \mathbb{R}^{(n_k r_k) \times r_{k-1}}$ is semi-orthogonal and $R_k \in \mathbb{R}^{r_{k-1} \times r_{k-1}}$ is upper-triangular.

3. Compute the new coefficient tensor

$$\mathbf{B}_{k-1}[\alpha_{k-2}, i, \ell_{k-1}] = \sum_{\alpha_{k-1}=1}^{r_{k-1}} \mathbf{A}_{k-1}[\alpha_{k-2}, i, \alpha_{k-1}] R_k[\ell_{k-1}, \alpha_{k-1}]. \tag{29}$$

Furthermore, at index $k = 1$, the unfolded \mathbf{C}_1 along the first coordinate is a row vector $\mathbf{C}_1^{(R)} \in \mathbb{R}^{1 \times (n_1 r_1)}$. Thus, the thin QR decomposition $Q_1 R_1 = (\mathbf{C}_1^{(R)})^\top$ produces a scalar R_1 such that $R_1^2 = \|\mathbf{C}_1^{(R)}\|^2$, and then the normalising constant $\hat{z} = \int_{\mathcal{X}_1} \hat{\pi}_{\leq 1}(x_1) \lambda_1(x_1) dx_1$ can be obtained by $\hat{z} = \gamma \prod_{i=1}^d \lambda_i(\mathcal{X}_i) + R_1^2$.

Proof See Appendix 2. □

Proposition 3 The marginal PDF of \hat{X}_1 can be expressed as

$$\hat{f}_{\hat{X}_1}(x_1) = \frac{1}{\hat{z}} \left(\gamma \prod_{i=2}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_1=1}^{r_1} \left(\sum_{i=1}^{n_1} \phi_1^{(i)}(x_1) D_1[i, \ell_1] \right)^2 \right) \lambda_1(x_1), \tag{30}$$

where $D_1[i, \ell_1] = \mathbf{B}_1[\alpha_0, i, \ell_1]$ for $i = 1, \dots, n_1$ and $\alpha_0 = 1$. For $k > 1$ and a given $x_{<k}$, the conditional PDF of $\hat{X}_k | \hat{X}_{<k}$ can be expressed as

$$\hat{f}_{\hat{X}_k | \hat{X}_{<k}}(x_k | x_{<k}) = \frac{1}{\hat{\pi}_{<k}(x_{<k})} \left(\gamma \prod_{i=k+1}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_k=1}^{r_k} \left(\sum_{i=1}^{n_k} \phi_k^{(i)}(x_k) D_k[i, \ell_k] \right)^2 \right) \lambda_k(x_k), \tag{31}$$

where $D_k \in \mathbb{R}^{n_k \times r_k}$ is given by

$$D_k[i, \ell_k] = \sum_{\alpha_{k-1}=1}^{r_{k-1}} G_{<k}^{(\alpha_{k-1})}(x_{<k}) \mathbf{B}_k[\alpha_{k-1}, i, \ell_k].$$

Proof The above results directly follow from the definition of conditional PDF and the marginal functions in (23) and (24). □

Note that the product $G_1(x_1) \cdots G_{k-1}(x_{k-1})$ requires $k - 1$ univariate interpolations and $k - 2$ products of matrices per sample, that is the same operations as in the standard inverse Rosenblatt transport. The QR decomposition (28) and the construction of the coefficient tensors (29) need $\mathcal{O}(dnr^3)$ operations, but these are pre-processing steps that are independent of the number of samples. However, in contrast to the vector-valued function $\mathcal{F}_k(x_k)\bar{\mathcal{F}}_{k+1} \cdots \bar{\mathcal{F}}_d \in \mathbb{R}^{r_{k-1}}$, in evaluating the PDF $\hat{f}_{\hat{X}}$, we need to multiply the matrix-valued function $P_k(x_k) \in \mathbb{R}^{r_{k-1} \times r_k}$ for each sample. Thus, the leading term of the complexity becomes $\mathcal{O}(Ndnr^2)$, one order of r or n higher than the complexity of the standard inverse Rosenblatt transport. However, for small r and n this is well compensated by a smoother map, which will be crucial in Sect. 4.

3.2 Implementation of CDFs

To evaluate SIRT, one has to first construct the marginal CDF of \hat{X}_1 and the conditional CDFs of $\hat{X}_k | \hat{X}_{<k}$ for $k > 1$ and then, inverts the CDFs (see (8)). Here, we discuss the computation and the inversion of CDFs, which are based on methods, for problems with bounded domains and extensions to problems with unbounded domains. We refer the readers to [3,57,62] and references therein for a more details.

3.2.1 Bounded Domain with Polynomial Basis

For a bounded parameter space $\mathcal{X} \subset \mathbb{R}^d$, we consider the weighting function $\lambda(x) = 1$. Since \mathcal{X} can be expressed as a Cartesian product, without loss of generality, here we discuss the CDF of a one-dimensional random variable Z with the PDF

$$\hat{f}_Z(\zeta) = C + \sum_{\ell=1}^r \left(\sum_{i=1}^n \phi^{(i)}(\zeta) D[i, \ell] \right)^2, \tag{32}$$

where $C > 0$ is some constant, $\{\phi^{(i)}(\zeta)\}_{i=1}^n$ are the basis functions, $D \in \mathbb{R}^{n \times r}$ is a coefficient matrix, and $\zeta \in [-1, 1]$. Here, $\hat{f}_Z(\zeta)$ can be either the marginal PDF or the conditional PDFs defined in Proposition 3 with a suitable linear change of coordinate.

We first consider a polynomial basis, $\phi^{(i)}(z) \in \mathbb{P}_{n-1}$ for $i = 1, \dots, n$, where \mathbb{P}_{n-1} is a vector space of polynomials of degree at most $n - 1$ defined on $[-1, 1]$. Thus, the PDF $\hat{f}_Z(\zeta)$ can be represented exactly in \mathbb{P}_{2n-2} . To enable fast computation of the CDF, we choose the Chebyshev polynomials of the second kind

$$p_m(\zeta) = \frac{\sin((m + 1) \cos^{-1}(\zeta))}{\sin(\cos^{-1}(\zeta))}, \quad m = 0, 1, \dots, 2n - 2,$$

as the basis of \mathbb{P}_{2n-2} . Using the roots of $p_{2n-1}(\zeta)$, we can define the set of collocation points

$$\{\zeta_m\}_{m=1}^{2n-1}, \quad \text{where } \zeta_m = \cos\left(\frac{m\pi}{2n}\right).$$

This way, by evaluating $\hat{f}_Z(\zeta)$ on the collocation points, which needs $\mathcal{O}(nr)$ operations, one can apply the collocation method ([3], Chapter 4) to represent $\hat{f}_Z(\zeta)$ using the Chebyshev basis:

$$\hat{f}_Z(\zeta) = \sum_{m=0}^{2n-2} a_m p_m(\zeta), \quad (33)$$

where the coefficients $\{a_m\}_{m=1}^{2n-2}$ can be computed by the fast Fourier transform with $\mathcal{O}(n \log(n))$ operations. Then, one can express the CDF of Z as

$$F_Z(\zeta) = \int_{-1}^{\zeta} \hat{f}_Z(\zeta') d\zeta' = \sum_{m=0}^{2n-2} \frac{a_m}{m+1} (t_{m+1}(\zeta) - t_{m+1}(-1)), \quad (34)$$

where $t_m(\zeta) = \cos(m \cos^{-1}(\zeta))$ is the Chebyshev polynomial of the first kind of degree m . A random variable Z can be generated by drawing a uniform random variable U and evaluating $Z = F_Z^{-1}(U)$ by solving the root finding problem $F_Z(Z) = U$.

Remark 3 The PDF in (32) is positive for all $\zeta \in [-1, 1]$ by construction and can be represented exactly in \mathbb{P}_{2n-2} with the polynomial basis. Thus, its Chebyshev representation in (33) is also positive. This way, the resulting CDF in (34) is monotone, and thus, the solution to the inverse CDF equation, $F_Z(Z) = U$, admits a unique solution.

Remark 4 One can also employ piecewise Lagrange polynomials as a basis to enable hp-adaptivity. With piecewise Lagrange polynomials, the above-mentioned technique can also be used to obtain the piecewise definition of the CDF.

Since $F_Z(Z) = U$ has a unique solution and F_Z is monotone and bounded between $[0, 1]$, it requires usually only a few iterations to apply the root finding methods, such as the regula falsi method and the Newton's method, to solve $F_Z(Z) = U$ with an accuracy close to machine precision. Overall, the construction of the CDF needs $\mathcal{O}(nr + n \log(n))$ operations, and the inversion of the CDF function needs $\mathcal{O}(cn)$ operations, where $\mathcal{O}(n)$ is the cost of evaluating the CDF and c is the number of iterations required by the root finding method. In comparison, building the matrix D requires $\mathcal{O}(nr^2)$ operations (cf. Proposition 3).

3.2.2 Bounded Domain with Fourier Basis

If the Fourier transform of the PDF of Z , which is the characteristic function, is band-limited in the frequency domain, then one may choose the sine and cosine Fourier series as the basis for representing the PDF in (32). In this case, the above strategy can also be applied. Recall the Fourier basis with an even cardinality n ,

$$\{1, \dots, \sin(m\pi\zeta), \cos(m\pi\zeta), \dots, \cos(n\pi\zeta/2)\}, \quad m = 1, \dots, n/2 - 1,$$

which consists of $n/2 - 1$ sine functions and $n/2 + 1$ cosine functions. The PDF $\hat{f}_Z(\zeta)$ defined in (32) yields an exact representation using the Fourier basis with cardinality $2n$. This way, one can represent $\hat{f}_Z(\zeta)$ as

$$\hat{f}_Z(\zeta) = a_0 + \sum_{m=1}^n a_m \cos(m\pi \zeta) + \sum_{m=1}^{n-1} b_m \sin(m\pi \zeta),$$

where the coefficients, a_m and b_m , are obtained by evaluating $\hat{f}_Z(\zeta)$ on the collocation points

$$\{\zeta_m\}_{m=1}^{2n}, \quad \text{where } \zeta_m = \frac{m}{n} - 1,$$

and applying the rectangular rule. This leads to the CDF

$$\begin{aligned} F_Z(\zeta) &= \int_{-1}^{\zeta} \hat{f}_Z(\zeta') d\zeta' \\ &= a_0(\zeta + 1) + \sum_{m=1}^n \frac{a_m}{m\pi} \sin(m\pi \zeta) - \sum_{m=1}^{n-1} \frac{b_m}{m\pi} (\cos(m\pi \zeta) - \cos(m\pi)). \end{aligned}$$

The construction and the inversion of the CDF using the Fourier basis cost a similar amount of operations compared to the polynomial basis.

3.2.3 Unbounded Domain

Given an unbounded domain, the simplest approach is to truncate the domain at the tail of the PDF. With the domain truncation, the above-mentioned implementations based on Chebyshev and Fourier basis can be applied directly. Although the function approximation error induced by the domain truncation can be bounded, using the resulting SIRT for computing expectations may lead to a biased estimator.

One can also consider basis functions that are intrinsic to an unbounded domain. For the domain $\mathcal{X}_k = (0, \infty)$, one can employ the Laguerre polynomials as the basis. This equips \mathcal{X}_k with a natural exponential weighting function $\lambda_k(x_k) = \exp(-x_k)$. The collocation method using higher order Laguerre polynomials can be applied again to obtain the exact representation of the CDF. Similarly, for $\mathcal{X}_k = (-\infty, \infty)$, the Hermite polynomials can be used as a basis, which equips \mathcal{X}_k with a Gaussian weighting function $\lambda_k(x_k) = \exp(-\frac{1}{2}x_k^2)$. Although one can apply the collocation method to obtain an algebraically exact representation of the CDF, the resulting CDF involves error functions, complementary error functions, and imaginary error functions. Those functions have to be approximated numerically. Thus, the computational cost of computing the CDF can be high, and it may be hard to guarantee the monotonicity and uniqueness of the inverse CDF solution at the tails. Using other bases such as the Whittaker cardinal functions for $\mathcal{X}_k = (-\infty, \infty)$ may face a similar challenge.

Remark 5 In a situation where the squared form of the PDF in (32) can be computed, but it is challenging to invert the CDF function, one can employ the rejection sampling [53] to generate random variables. In this situation, our TT approximation can still be used to draw conditional samples. However, this approach may not lead to the deterministic inverse Rosenblatt transport.

3.2.4 Change of Coordinate

One can also apply a diffeomorphic mapping to change the coordinate of an unbounded domain \mathcal{X} to a bounded one, e.g. $\mathcal{Z} = [-1, 1]^d$, followed by application of the Chebyshev polynomials or Fourier series. Given a PDF $f_X(x)$ of a random variable $X \in \mathcal{X}$, suppose we have a diffeomorphic mapping $R : \mathcal{X} \mapsto \mathcal{Z}$ and let $q(x) = |\nabla R(x)| \geq 0$. For any Borel set $\mathcal{B}_X \subseteq \mathcal{X}$, we have

$$\begin{aligned} \mathbb{P}[X \in \mathcal{B}_X] &= \int_{\mathcal{B}_X} f_X(x) dx \\ &= \int_{\mathcal{B}_Z} (f_X \circ R^{-1})(\zeta) |\nabla R^{-1}(\zeta)| d\zeta = \int_{\mathcal{B}_Z} \frac{(f_X \circ R^{-1})(\zeta)}{(q \circ R^{-1})(\zeta)} d\zeta, \end{aligned}$$

where $\mathcal{B}_Z = R(\mathcal{B}_X)$. Thus, one can draw a random variable Z with PDF

$$f_Z(\zeta) = \frac{(f_X \circ R^{-1})(\zeta)}{(q \circ R^{-1})(\zeta)},$$

and apply the mapping $X = R(Z)$ to obtain a random variable X . With the change of the coordinate $\zeta = R(x)$, one needs to build a TT to approximate $\sqrt{f_Z(\zeta)}$ and construct the corresponding SIRT to simulate the random variable Z . To avoid singularities at the boundary of \mathcal{Z} , one can choose a mapping R such that the function $q(x)$ decays slower than $f_X(x)$.

3.3 SIRT Error

Since SIRT enables us to generate i.i.d. samples from the probability measure $\hat{T}_\# \mu_{\text{uni}}$, it can be used to define either Metropolis independence samplers or importance sampling schemes. Based on certain assumptions on the TT approximation \tilde{g} , here we establish error bounds for the TV distance, the Hellinger distance, and the χ^2 -divergence of the target measure ν_π from $\hat{T}_\# \mu_{\text{uni}}$. These divergences play a vital role in analysing the convergence of Metropolis–Hastings methods and the efficiency of importance sampling. The error analysis also provides a heuristic for choosing the constant γ in the approximate target density (19).

Proposition 4 Recall the approximate target density $\hat{\pi}(x) = \gamma + \tilde{g}(x)^2$, where $\tilde{g}(x)$ is a TT approximation to the square root of the unnormalised target density $\sqrt{\pi}$. Suppose

the error of \tilde{g} and the constant γ satisfy

$$\|\tilde{g} - \sqrt{\pi}\|_{L^2_\lambda(\mathcal{X})} \leq \epsilon \quad \text{and} \quad \gamma \leq \frac{1}{\lambda(\mathcal{X})} \|\tilde{g} - \sqrt{\pi}\|_{L^2_\lambda(\mathcal{X})}^2, \tag{35}$$

respectively. Then, the error of $\sqrt{\hat{\pi}}$ satisfies $\|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} \leq \sqrt{2}\epsilon$.

Proof Applying the identity $(\sqrt{\pi} - \sqrt{\gamma + \tilde{g}^2})^2 \leq (\sqrt{\pi} - \tilde{g})^2 + \gamma$, we have

$$\|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} \leq \left(\int_{\mathcal{X}} (\sqrt{\pi(x)} - \tilde{g}(x))^2 \lambda(x) dx + \gamma \lambda(\mathcal{X}) \right)^{\frac{1}{2}} \leq \sqrt{2}\epsilon.$$

□

Proposition 5 Suppose the conditions in (35) hold. Then, the approximate normalising constant \hat{z} satisfies $|\sqrt{z} - \sqrt{\hat{z}}| \leq \sqrt{2}\epsilon$.

Proof The normalising constants z and \hat{z} satisfy

$$|z - \hat{z}| = \left| \int_{\mathcal{X}} (\pi(x) - \hat{\pi}(x)) \lambda(x) dx \right| \leq \|\pi - \hat{\pi}\|_{L^1_\lambda(\mathcal{X})}. \tag{36}$$

Applying the Hölder’s inequality (with $p = q = 2$) and the Minkowski inequality, the right hand side of the above inequality also satisfies

$$\begin{aligned} \|\pi - \hat{\pi}\|_{L^1_\lambda(\mathcal{X})} &= \|(\sqrt{\pi} - \sqrt{\hat{\pi}})(\sqrt{\pi} + \sqrt{\hat{\pi}})\|_{L^1_\lambda(\mathcal{X})} \\ &\leq \|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} \|\sqrt{\pi} + \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} \\ &\leq \|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} \left(\|\sqrt{\pi}\|_{L^2_\lambda(\mathcal{X})} + \|\sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} \right) \\ &= \|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})} (\sqrt{z} + \sqrt{\hat{z}}). \end{aligned} \tag{37}$$

Since both z and \hat{z} are positive, we have $|z - \hat{z}| = |\sqrt{\hat{z}} - \sqrt{z}| (\sqrt{z} + \sqrt{\hat{z}})$. Substituting this identity and the inequality in (37) into (36), we have $|\sqrt{z} - \sqrt{\hat{z}}| \leq \|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L^2_\lambda(\mathcal{X})}$. Thus, the result follows from Proposition 4. □

Theorem 1 Suppose the conditions in (35) hold. The Hellinger distance between ν_π and $\hat{T}_\# \mu_{\text{uni}}$ satisfies $D_H(\nu_\pi \| \hat{T}_\# \mu_{\text{uni}}) \leq 2\epsilon / \sqrt{z}$.

Proof Since the target measure ν_π and the approximate measure $\hat{T}_\# \mu_{\text{uni}}$, respectively, have the densities $\frac{1}{z} \pi(x) \lambda(x)$ and $\frac{1}{\hat{z}} \hat{\pi}(x) \lambda(x)$, the squared Hellinger distance satisfies

$$D_H^2(\nu_\pi \| \hat{T}_\# \mu_{\text{uni}}) = \frac{1}{2} \int_{\mathcal{X}} \left(\sqrt{\frac{\pi(x)}{z}} - \sqrt{\frac{\hat{\pi}(x)}{\hat{z}}} \right)^2 \lambda(x) dx.$$

This leads to the inequality

$$D_H^2(v_\pi \| \hat{T}_\# \mu_{\text{uni}}) = \frac{1}{2z} \|\sqrt{\pi} - \sqrt{\hat{\pi}} + \sqrt{\hat{\pi}} - \sqrt{\pi} \sqrt{z/\hat{z}}\|_{L_\lambda^2(\mathcal{X})}^2 \leq \frac{1}{2z} \left(\|\sqrt{\pi} - \sqrt{\hat{\pi}}\|_{L_\lambda^2(\mathcal{X})} + \|\sqrt{\hat{\pi}}\|_{L_\lambda^2(\mathcal{X})} |1 - \sqrt{z/\hat{z}}| \right)^2.$$

Applying $\|\sqrt{\hat{\pi}}\|_{L_\lambda^2(\mathcal{X})}^2 = \hat{z}$ and Propositions 4 and 5, the above inequality can be further reduced to

$$D_H^2(v_\pi \| \hat{T}_\# \mu_{\text{uni}}) \leq \frac{1}{2z} \left(\sqrt{2}\epsilon + |\sqrt{\hat{z}} - \sqrt{z}| \right)^2 \leq \frac{4\epsilon^2}{z}.$$

Thus, we have $D_H(v_\pi \| \hat{T}_\# \mu_{\text{uni}}) \leq 2\epsilon/\sqrt{z}$. □

Corollary 1 Suppose the conditions in (35) hold. The total variation distance between v_π and $\hat{T}_\# \mu_{\text{uni}}$ satisfies $D_{\text{TV}}(v_\pi \| \hat{T}_\# \mu_{\text{uni}}) \leq 2\sqrt{2}\epsilon/\sqrt{z}$.

Proof The result directly follows from the inequality $D_{\text{TV}} \leq \sqrt{2}D_H$ and Theorem 1. □

Proposition 6 Given two probability measures v_π and $\hat{v}_{\hat{\pi}}$ and a function h with finite second moments with respect to v_π and $\hat{v}_{\hat{\pi}}$. Then,

$$|v_\pi(h) - \hat{v}_{\hat{\pi}}(h)| \leq \sqrt{2} \left(v_\pi(h^2)^{\frac{1}{2}} + \hat{v}_{\hat{\pi}}(h^2)^{\frac{1}{2}} \right) D_H(v_\pi \| \hat{v}_{\hat{\pi}}).$$

Proof Suppose v_π and $\hat{v}_{\hat{\pi}}$, respectively, have density functions $f_X(x)$ and $\hat{f}_{\hat{X}}(x)$ with respect to the Lebesgue measure. We have the following inequality

$$\begin{aligned} & |v_\pi(h) - \hat{v}_{\hat{\pi}}(h)| \\ &= \left| \int_{\mathcal{X}} h(x)(f_X(x) - \hat{f}_{\hat{X}}(x))dx \right| \\ &= \left| \int_{\mathcal{X}} h(x) \left(f_X(x)^{\frac{1}{2}} + \hat{f}_{\hat{X}}(x)^{\frac{1}{2}} \right) \left(f_X(x)^{\frac{1}{2}} - \hat{f}_{\hat{X}}(x)^{\frac{1}{2}} \right) dx \right| \\ &\leq \left(\int_{\mathcal{X}} \left(h(x)f_X(x)^{\frac{1}{2}} + h(x)\hat{f}_{\hat{X}}(x)^{\frac{1}{2}} \right)^2 dx \right)^{\frac{1}{2}} \left(\int_{\mathcal{X}} \left(f_X(x)^{\frac{1}{2}} - \hat{f}_{\hat{X}}(x)^{\frac{1}{2}} \right)^2 dx \right)^{\frac{1}{2}} \\ &\leq \sqrt{2} \left(\left(\int_{\mathcal{X}} h(x)^2 f_X(x) dx \right)^{\frac{1}{2}} + \left(\int_{\mathcal{X}} h(x)^2 \hat{f}_{\hat{X}}(x) dx \right)^{\frac{1}{2}} \right) D_H(v_\pi \| \hat{v}_{\hat{\pi}}). \end{aligned}$$

Thus, the result follows. □

Corollary 2 Suppose the conditions in (35) hold. Suppose further the bound in (20) holds. Then, the χ^2 -divergence of v_π from $\hat{T}_\# \mu_{\text{uni}}$ satisfies

$$D_{\chi^2}(v_\pi \| \hat{T}_\# \mu_{\text{uni}}) \leq \left(v_\pi(\pi^2/\hat{\pi}^2)^{\frac{1}{2}} + (\hat{T}_\# \mu_{\text{uni}})(\pi^2/\hat{\pi}^2)^{\frac{1}{2}} \right) \frac{2\sqrt{2}\hat{z}}{z\sqrt{z}} \epsilon.$$

Proof Given the bound in (20), the ratio between $f_X(x)$ and $\hat{f}_{\hat{X}}(x)$ satisfies

$$\sup_{x \in \mathcal{X}} \frac{f_X(x)}{\hat{f}_{\hat{X}}(x)} = \frac{\hat{z}}{z} \sup_{x \in \mathcal{X}} \frac{\pi(x)}{\hat{\pi}(x)} = \frac{\hat{z}}{z} \hat{c} < \infty.$$

This way, ν_π is absolutely continuous with respect to $\hat{T}_\# \mu_{\text{uni}}$. Thus, the χ^2 -divergence of ν_π from $\hat{T}_\# \mu_{\text{uni}}$ can be expressed as

$$\begin{aligned} D_{\chi^2}(\nu_\pi \| \hat{T}_\# \mu_{\text{uni}}) &= \int_{\mathcal{X}} \left(\frac{f_X(x)}{\hat{f}_{\hat{X}}(x)} \right)^2 \hat{f}_{\hat{X}}(x) dx - 1 \\ &= \nu_\pi(f_X / \hat{f}_{\hat{X}}) - (\hat{T}_\# \mu_{\text{uni}})(f_X / \hat{f}_{\hat{X}}) \\ &= \frac{\hat{z}}{z} \left(\nu_\pi(\pi / \hat{\pi}) - (\hat{T}_\# \mu_{\text{uni}})(\pi / \hat{\pi}) \right). \end{aligned}$$

The bound in (20) also implies that the expectations $\nu_\pi(\pi^2 / \hat{\pi}^2)$ and $(\hat{T}_\# \mu_{\text{uni}})(\pi^2 / \hat{\pi}^2)$ are finite. Therefore, the result follows from Proposition 6. \square

Remark 6 The TV distance, Hellinger distance, and χ^2 -divergence of the target measure ν_π from the pushforward of the reference μ_{uni} under the SIRT \hat{T} are linear in the approximation error of the TT. Note that $\sqrt{z} = \|\sqrt{\pi}\|_{L^2_\lambda(\mathcal{X})}$. Suppose the relative error of \tilde{g} satisfies

$$\|\tilde{g} - \sqrt{\pi}\|_{L^2_\lambda(\mathcal{X})} / \|\sqrt{\pi}\|_{L^2_\lambda(\mathcal{X})} \leq \tau,$$

the TV distance, Hellinger distance, and χ^2 -divergence of ν_π from $\hat{T}_\# \mu_{\text{uni}}$ are bounded by $\mathcal{O}(\tau)$.

4 Deep Inverse Rosenblatt Transport

In many practical applications, probability densities can be concentrated to a small region of the parameter space or have complicated correlation structures. For example, posterior densities in Bayesian inference problems with informative data often occupy a relatively small region of the parameter space and demonstrate complicated nonlinear interactions in some sub-manifold, see [11,12,21,50] for detailed examples. In this situation, straightforward approximation of a complicated density function in a TT decomposition may require rather large ranks. As the number of function evaluations needed in constructing TT decompositions grows quadratically with the ranks, such direct factorisation of the target densities with complicated structures may become infeasible.

Example 1 Consider a d -dimensional multivariate normal distribution with the unnormalised density $\pi(x) = \exp(-\frac{1}{2}x^\top C^{-1}x)$. If the covariance matrix $C \in \mathbb{R}^{d \times d}$ is diagonal, the joint density factorises into a product of marginal densities, that is a

TT decomposition with ranks 1. This corresponds to zero-rank *off-diagonal* blocks $C[1 : k, (k + 1) : d], k = 1, \dots, d - 1$. However, the TT ranks of a correlated normal density $\pi(x)$ may grow *exponentially* in the rank of the off-diagonal blocks of C [54].

We design a DIRT framework to construct a *composition* of order-preserving mappings in the SIRT format that can characterise concentrated probability densities with complicated correlation structures. The construction of DIRT is guided by a sequence of bridging probability measures $\nu_0, \nu_1, \dots, \nu_L$, where $\nu_L = \nu_\pi$ is the target measure. Each bridging measure ν_k has the corresponding PDF

$$f_{X^k}(x) = \frac{1}{z_k} \pi_k(x) \lambda(x), \quad \text{where } z_k = \int_{\mathcal{X}} \pi_k(x) \lambda(x) dx. \tag{38}$$

Here, $\pi_0(x)$ is the unnormalised initial density such that $\sup_{x \in \mathcal{X}} \pi_0(x) < \infty, \pi_L(x) = \pi(x)$ is the unnormalised target density, and the superscript k indexes the random variable $X^k \in \mathcal{X}, X^k \sim \nu_k$. Our goal is to construct a composition of mappings $T_0 \circ T_1 \circ \dots \circ T_k$ such that the pushforward of the reference measure under this composition matches the k th bridging measure, i.e. $(T_0 \circ T_1 \circ \dots \circ T_k)_\# \mu = \nu_k$. This way, by gradually increasing the complexity in the geometry and/or the computational cost of the densities of the bridging measures, it becomes computationally feasible to construct TT and the corresponding SIRT at each layer of the composition.

Assumption 1 Denoting the ratio between two unnormalised densities by

$$r_{k,j}(x) = \frac{\pi_k(x)}{\pi_j(x)}, \tag{39}$$

we assume that for each pair of $j < k$, the ratio $r_{k,j}(x)$ is finite such that

$$\sup_{x \in \mathcal{X}} r_{k,j}(x) = c_{k,j} < \infty, \quad \forall j < k. \tag{40}$$

In practice, there are many ways to choose the bridging measures. For example, one can consider *tempered* distributions [20,31,40,44,60] where $\pi_k(x) = \pi(x)^{\beta_k}$ for a suitable chosen set of powers (*reciprocal temperatures*) $0 \leq \beta_0 < \dots < \beta_L = 1$; and for problems involving computationally expensive PDE models, one can employ a hierarchy of models with different grid resolutions to reduce the computational cost for building the DIRT. In the rest of this section, we will present the recursive construction of DIRT and provide error analysis.

4.1 Recursive Construction

In the initial step ($k = 0$), we compute a TT $\tilde{g}_0(x)$ that approximates $\sqrt{\pi_0}$ and construct the corresponding SIRT $\hat{X}^0 = \hat{T}_0(U)$ so that the reference random variable $U \sim \mu_{\text{uni}}$ and $\hat{X}^0 \sim (\hat{T}_0)_\# \mu_{\text{uni}}$ with the PDF

$$\hat{f}_{\hat{X}^0}(x) = \frac{1}{\hat{z}_0} (\gamma_0 + \tilde{g}_0(x)^2) \lambda(x), \quad \text{with } \hat{z}_0 = \gamma_0 \lambda(\mathcal{X}) + \int_{\mathcal{X}} \tilde{g}_0(x)^2 \lambda(x) dx, \tag{41}$$

where the constant γ_0 is chosen such that $0 < \gamma_0 \leq \lambda(\mathcal{X})^{-1} \|\sqrt{\pi_0} - \tilde{g}_0\|_{L^2_\lambda(\mathcal{X})}^2$. Note that $\hat{f}_{\hat{\mathcal{X}}^0}(x)$ is an approximation to $f_{X^0}(x)$.

Remark 7 We can replace the uniform reference measure μ_{uni} with a general product-form probability measure μ that has the PDF $f_U(u) = \prod_{k=1}^d f_{U_k}(u_k)$ with support in $\mathcal{U} = \mathcal{U}_1 \times \mathcal{U}_2 \times \dots \times \mathcal{U}_d$. One can construct a mapping

$$R(u) = [F_{U_1}(u_1), \dots, F_{U_k}(u_k), \dots, F_{U_d}(u_d)]^\top,$$

where $F_{U_k}(u_k)$ is the CDF of U_k , such that $R_\# \mu = \mu_{\text{uni}}$. Then, the composition of mappings $\hat{T}_0 \circ R : \mathcal{U} \mapsto \mathcal{X}$ is lower-triangular and $(\hat{T}_0 \circ R)_\# \mu$ has the density $\hat{f}_{\hat{\mathcal{X}}^0}(x)$. We initialise the DIRT by $\bar{T}_0 = \hat{T}_0 \circ R$, where R is the identity map if $\mu = \mu_{\text{uni}}$.

After $k > 1$ steps, suppose we have the k th DIRT given as a composition of mappings

$$\bar{T}_k = (\hat{T}_0 \circ R) \circ (\hat{T}_1 \circ R) \circ \dots \circ (\hat{T}_k \circ R),$$

where each \hat{T}_j is a SIRT. Denoting the pushforward of the reference probability measure μ under \bar{T}_k by $\hat{\nu}_k$, i.e. $\hat{\nu}_k \equiv (\bar{T}_k)_\# \mu$, and the density function of $\hat{\nu}_k$ by $\hat{f}_{\hat{\mathcal{X}}^k}(x)$, the pullback density of $\hat{\nu}_k$ under \bar{T}_k satisfies

$$\bar{T}_k^\# \hat{f}_{\hat{\mathcal{X}}^k}(u) = (\hat{f}_{\hat{\mathcal{X}}^k} \circ \bar{T}_k)(u) |\nabla_u \bar{T}_k(u)| = f_U(u). \tag{42}$$

The density function of the pullback probability measure $\bar{T}_k^\# \hat{\nu}_k$ is the reference product density $f_U(u)$. Suppose the corresponding approximate PDF $\hat{f}_{\hat{\mathcal{X}}^k}(x)$ can capture the range of variation and the correlation structure of the next PDF $f_{X^{k+1}}(x)$, then the density function of the pullback probability measure $\bar{T}_k^\# \nu_{k+1}$,

$$\bar{T}_k^\# f_{X^{k+1}}(u) = (f_{X^{k+1}} \circ \bar{T}_k)(u) |\nabla_u \bar{T}_k(u)|, \tag{43}$$

may become easier to factorise in the TT format compared to the direct factorisation of the original target density function $f_{X^{k+1}}(x)$. This way, for step $k + 1$, the existing composition \bar{T}_k can be used to precondition the construction of the coupling between μ and ν_{k+1} : by building a coupling between the pullback measure $\bar{T}_k^\# \nu_{k+1}$ and the reference measure

$$U^{k+1} = (T_{k+1} \circ R)(U), \quad \text{where } U \sim \mu \quad \text{and} \quad U^{k+1} \sim \bar{T}_k^\# \nu_{k+1},$$

one can obtain a new composition of maps $(\bar{T}_k \circ T_{k+1} \circ R)$ such that

$$(\bar{T}_k \circ T_{k+1} \circ R)^\# \nu_{k+1} = \mu \quad \text{or} \quad (\bar{T}_k \circ T_{k+1} \circ R)_\# \mu = \nu_{k+1}.$$

We use SIRT to approximate T_{k+1} . Using (42), we have $|\nabla_u \bar{T}_k(u)| = ((\hat{f}_{\hat{X}^k} \circ \bar{T}_k)(u))^{-1} f_U(u)$. Thus, the pullback density in (43) can be expressed as a *ratio function*

$$(\bar{T}_k)^\# f_{X^{k+1}}(u) = \frac{(f_{X^{k+1}} \circ \bar{T}_k)(u)}{(\hat{f}_{\hat{X}^k} \circ \bar{T}_k)(u)} f_U(u). \tag{44}$$

This way, we can compute a TT $\tilde{g}_{k+1}(u)$ to approximate the function

$$q_{k+1}(u) \propto \left(\frac{(f_{X^{k+1}} \circ \bar{T}_k)(u)}{(\hat{f}_{\hat{X}^k} \circ \bar{T}_k)(u)} \frac{f_U(u)}{\omega(u)} \right)^{\frac{1}{2}}, \tag{45}$$

where $\omega(u)$ is the weighting function associated with the reference domain \mathcal{U} and $a \propto b$ denotes that a is proportional to b . Since z_k , the normalising constant of $f_{X^{k+1}}$, is unknown, here we only need to factorise an unnormalised version of $q_{k+1}(u)$ into a TT. The normalising constant is computed automatically during the marginalisation process of SIRT (see Proposition 2).

The SIRT $U^{k+1} = \hat{T}_{k+1}(U')$ built from the TT $\tilde{g}_{k+1}(u)$ couples the uniform reference random variable $U' \sim \mu_{\text{uni}}$ with $U^{k+1} \sim (\hat{T}_{k+1})^\# \mu_{\text{uni}}$. Thus, the composition of transformations $U^{k+1} = (\hat{T}_{k+1} \circ R)(U)$ couples the general reference random variable $U \sim \mu$ with $U^{k+1} \sim (\hat{T}_{k+1} \circ R)^\# \mu$, where $(\hat{T}_{k+1} \circ R)^\# \mu$ is an approximation to the pullback measure $\bar{T}_k^\# \nu_{k+1}$. Thus, we have

$$(\hat{T}_{k+1} \circ R)^\# (\bar{T}_k^\# \nu_{k+1}) = (\bar{T}_k \circ \hat{T}_{k+1} \circ R)^\# \nu_{k+1} \approx \mu \quad \text{or} \quad (\bar{T}_k \circ \hat{T}_{k+1} \circ R)^\# \mu \approx \nu_{k+1}.$$

The next DIRT is therefore defined by the new composition of mappings

$$\bar{T}_{k+1} := \bar{T}_k \circ (\hat{T}_{k+1} \circ R).$$

The recursion is completed by obtaining \hat{T}_L and \bar{T}_L .

Proposition 7 At the j th DIRT step, the Jacobian of the incremental mapping \hat{T}_j is given by

$$|\nabla_u \hat{T}_j^{-1}(u)| = \hat{p}_{U^j}(u) = \frac{1}{\hat{z}_j} (\gamma_j + \tilde{g}_j(u)^2) \omega(u), \tag{46}$$

with

$$\hat{z}_j = \gamma_j \omega(\mathcal{U}) + \int_{\mathcal{U}} \tilde{g}_j(u)^2 \omega(u) du, \tag{47}$$

where the constant $\gamma_j > 0$ is chosen according to the L^2 error of \tilde{g}_j .

Proof The SIRT $U^j = \hat{T}_j(U')$, which is constructed by integrating $(\gamma_j + \tilde{g}_j(u)^2)\omega(u)$, maps the uniform random variable $U' \sim \mu_{\text{uni}}$ to $U^j \sim (\hat{T}_j)_\# \mu_{\text{uni}}$. Thus, the pushforward measure $(\hat{T}_j)_\# \mu_{\text{uni}}$ has the density function $\hat{p}_{U^j}(u)$ defined in (46), which yields

$$\hat{p}_{U^j}(u) \equiv (\hat{T}_j)_\# f_{\text{uni}}(u) = |\nabla_u \hat{T}_j^{-1}(u)|.$$

□

Lemma 1 *At step k of the DIRT construction, suppose we have the DIRTs*

$$\bar{T}_j = (\hat{T}_0 \circ R) \circ (\hat{T}_1 \circ R) \circ \dots \circ (\hat{T}_j \circ R), \quad \text{for } j \leq k.$$

Suppose further we have a normalised density function over the domain \mathcal{X} defined in (41) for $j = 0$, and normalised density functions over the reference domain \mathcal{U} defined in (46) for $j = 1, \dots, k$. Then, the pushforward measure $(\bar{T}_k)_\# \mu$ has the PDF

$$\hat{f}_{\hat{\mathcal{X}}^k}(x) \equiv (\bar{T}_k)_\# f_U(x) = \hat{f}_{\hat{\mathcal{X}}^0}(x) \prod_{j=1}^k \frac{(\hat{p}_{U^j} \circ \bar{T}_{j-1}^{-1})(x)}{(f_U \circ \bar{T}_{j-1}^{-1})(x)}. \tag{48}$$

Proof The result can be shown using induction. For the case $k = 0$, the result follows directly from (41). Suppose (48) holds for $k > 0$. We define the composition of mappings

$$\tilde{T}_k = \bar{T}_k \circ R^{-1} = \hat{T}_0 \circ (R \circ \hat{T}_1) \circ \dots \circ (R \circ \hat{T}_k).$$

Since $R_\# \mu = \mu_{\text{uni}}$, we have the identity $(\bar{T}_k)_\# \mu = (\tilde{T}_k)_\# \mu_{\text{uni}}$, which leads to

$$\hat{f}_{\hat{\mathcal{X}}^k}(x) = (\bar{T}_k)_\# f_U(x) = (\tilde{T}_k)_\# f_{\text{uni}}(x).$$

At step $k + 1$, we have the new composition of mappings $\tilde{T}_{k+1} = \tilde{T}_k \circ R \circ \hat{T}_{k+1}$, the pushforward measures, $(\tilde{T}_{k+1})_\# \mu$ and $(\tilde{T}_k)_\# \mu$, have the density functions

$$\begin{aligned} \hat{f}_{\hat{\mathcal{X}}^{k+1}}(x) &= (\tilde{T}_{k+1})_\# f_{\text{uni}}(x) = (f_{\text{uni}} \circ \tilde{T}_{k+1}^{-1})(x) |\nabla_x \tilde{T}_{k+1}^{-1}(x)| = |\nabla_x \tilde{T}_{k+1}^{-1}(x)|, \\ \hat{f}_{\hat{\mathcal{X}}^k}(x) &= (\tilde{T}_k)_\# f_{\text{uni}}(x) = (f_{\text{uni}} \circ \tilde{T}_k^{-1})(x) |\nabla_x \tilde{T}_k^{-1}(x)| = |\nabla_x \tilde{T}_k^{-1}(x)|, \end{aligned}$$

respectively. Applying the change of variables $u' = \tilde{T}_k^{-1}(x)$ and $u = R^{-1}(u')$, the determinant of $\nabla_x \tilde{T}_{k+1}^{-1}(x)$ can be expressed as

$$|\nabla_x \tilde{T}_{k+1}^{-1}(x)| = |\nabla_u \hat{T}_{k+1}^{-1}(u)| |\nabla_{u'} R^{-1}(u')| |\nabla_x \tilde{T}_k^{-1}(x)|. \tag{49}$$

Note that the above change of variables implies also that $|\nabla_{u'} R^{-1}(u')| = f_U(u)^{-1}$ and $u = (\tilde{T}_k \circ R)^{-1}(x) = \bar{T}_k^{-1}(x)$. Together with Proposition 7, we have

$$|\nabla_x \tilde{T}_{k+1}^{-1}(x)| = \frac{(\hat{p}_{U^{k+1}} \circ \bar{T}_k^{-1})(x)}{(f_U \circ \bar{T}_k^{-1})(x)} |\nabla_x \tilde{T}_k^{-1}(x)|.$$

Thus, the result follows. □

Corollary 3 *At step k of the DIRT construction, the composition of mappings, \bar{T}_k , satisfies*

$$|\nabla_x \bar{T}_k^{-1}(x)| = \frac{\hat{f}_{\hat{X}^k}(x)}{(f_U \circ \bar{T}_k^{-1})(x)}.$$

Proof The result direct follows from $\bar{T}_k^{-1} = R^{-1} \circ \tilde{T}_k^{-1}$ and the proof of Lemma 1. □

Remark 8 The normalised PDFs of the k th DIRT step can be expressed as

$$\hat{f}_{\hat{X}^k}(x) = \frac{1}{\bar{z}_k} \hat{\pi}_k(x) \lambda(x), \tag{50}$$

where $\bar{z}_k = \prod_{j=0}^k \hat{z}_j$ is an approximation to the normalising constant z_k , and

$$\hat{\pi}_k(x) = (\gamma_0 + \tilde{g}_0(x)^2) \prod_{j=1}^k \frac{(\gamma_j + (\tilde{g}_j \circ \bar{T}_{j-1}^{-1})(x)^2) (\omega \circ \bar{T}_{j-1}^{-1})(x)}{(f_U \circ \bar{T}_{j-1}^{-1})(x)} \tag{51}$$

is an approximation to the unnormalised bridging density $\pi_k(x)$.

4.2 Ratio Functions and Error Analysis

We will first discuss the ratio function (45) and its approximation and then present the corresponding error analysis.

4.2.1 Ratio Functions

Given the unnormalised PDF in (51), the pullback density in (44) can be expressed as

$$(\bar{T}_k)^\# f_{X^{k+1}}(u) = \frac{(f_{X^{k+1}} \circ \bar{T}_k)(u)}{(\hat{f}_{\hat{X}^k} \circ \bar{T}_k)(u)} f_U(u) \propto \frac{(\pi_{k+1} \circ \bar{T}_k)(u)}{(\hat{\pi}_k \circ \bar{T}_k)(u)} f_U(u).$$

This way, we need to compute a TT $\tilde{g}_{k+1}(u)$ to approximate the function

$$q_{k+1}(u) = \left(\frac{(\pi_{k+1} \circ \bar{T}_k)(u)}{(\hat{\pi}_k \circ \bar{T}_k)(u)} \frac{f_U(u)}{\omega(u)} \right)^{\frac{1}{2}}, \tag{52}$$

to build the SIRT \hat{T}_{k+1} . We call this strategy the *exact ratio approach*.

Alternatively, the pullback density in (44) can be expressed as

$$(\bar{T}_k)^\# f_{X^{k+1}}(u) \propto \frac{(\pi_{k+1} \circ \bar{T}_k)(u)}{(\pi_k \circ \bar{T}_k)(u)} \frac{(\pi_k \circ \bar{T}_k)(u)}{(\hat{\pi}_k \circ \bar{T}_k)(u)} f_U(u) \tag{53}$$

Since the DIRT density function $\hat{\pi}_k$ approximates the k th unnormalised bridging density function π_k , the pullback density in (53) can be approximated as

$$(\bar{T}_k)^\# f_{X^{k+1}}(u) \propto (r_{k+1,k} \circ \bar{T}_k)(u) \frac{(\pi_k \circ \bar{T}_k)(u)}{(\hat{\pi}_k \circ \bar{T}_k)(u)} f_U(u) \approx (r_{k+1,k} \circ \bar{T}_k)(u) f_U(u).$$

This way, we need to compute a TT $\tilde{g}_{k+1}(u)$ that approximates the function

$$\tilde{q}_{k+1}(u) = \left(\frac{(r_{k+1,k} \circ \bar{T}_k)(u) f_U(u)}{\omega(u)} \right)^{\frac{1}{2}}, \tag{54}$$

to build an alternative SIRT \hat{T}_{k+1} . We call this strategy the *approximate ratio approach*.

Remark 9 For all $k \geq 0$, we want the ratio $\pi_k/\hat{\pi}_k$ to be finite in \mathcal{U} . Otherwise, it may cause large errors in the TT decomposition and may deteriorate the convergence of the resulting sampling schemes for characterising ν_π . Given Assumption 1 and $\gamma_j > 0$ for all $j = 0, \dots, k$, we have

$$\sup_{x \in \mathcal{X}} \frac{\pi_0(x)}{\gamma_0 + \tilde{g}_0(x)^2} < \infty \quad \text{and} \quad \sup_{u \in \mathcal{U}} \frac{(r_{k,k-1} \circ \bar{T}_{k-1})(u) f_U(u)}{(\gamma_k + \tilde{g}_k(u)^2) \omega(u)} < \infty, \tag{55}$$

and thus, it can be shown (using induction) that the ratio $\pi_k/\hat{\pi}_k$ is bounded.

Remark 10 In some situations, the ratio function $(r_{k+1,k} \circ \bar{T}_k)(u)$ may exhibit sharp boundary layers if the uniform reference measure μ_{uni} (with $R = I$) is used. This can increase the complexity of the resulting TT decompositions. Apart from carefully choosing the bridging measures, a partial remedy to the boundary layer is to use a reference measure with the density $f_U(u)$ decaying towards the boundary, such as the normal density truncated on a sufficiently large hypercube $[-\sigma, \sigma]^d$. The function $f_U(u)$ in (52) and (54) smoothens the previous approximation errors, which can improve the accuracy of TT approximations. With a reference measure defined on a hypercube, the collocation techniques based on Chebyshev and Fourier bases (see Sect. 3.2) can be applied to construct and evaluate functional TT decompositions in DIRT.

4.2.2 DIRT Error

Based on assumptions on the TT error at each layer of the DIRT construction, here we first establish bounds on approximation errors of the DIRT-induced approximate

density $\hat{\pi}_k$ for both the exact ratio approach and the approximate ratio approach. The error analysis also provides heuristics for choosing the constants γ_j . Then, the error bounds of $\sqrt{\hat{\pi}_k}$ leads to bounds on the TV distance, the Hellinger distance, and the χ^2 -divergence of the k th bridging measure ν_k from the pushforward measure $(\bar{T}_k)_\# \mu$.

Theorem 2 (Exact ratio approach) *At the k th DIRT step ($k > 0$), suppose the TT decomposition $\tilde{g}_k \approx q_k$ and the constant γ_k , respectively, satisfy*

$$\|\tilde{g}_k - q_k\|_{L^2_\omega(\mathcal{U})} \leq \epsilon_k \quad \text{and} \quad \gamma_k \leq \frac{1}{\omega(\mathcal{U})} \|\tilde{g}_k - q_k\|_{L^2_\omega(\mathcal{U})}^2,$$

where q_k is defined in (52). Then, the unnormalised PDF $\hat{\pi}_k$ (51) approximates the unnormalised density function of the k th bridging measure π_k with the error

$$\|\sqrt{\pi_k} - \sqrt{\hat{\pi}_k}\|_{L^2_\lambda(\mathcal{X})} \leq \sqrt{2\bar{z}_{k-1}}\epsilon_k,$$

where $\bar{z}_{k-1} = \prod_{j=0}^{k-1} \hat{z}_j$ is the normalising constant of the unnormalised PDF $\hat{\pi}_{k-1}$.

Proof We start with the identity

$$\begin{aligned} \|\sqrt{\pi_k} - \sqrt{\hat{\pi}_k}\|_{L^2_\lambda(\mathcal{X})} &= \left(\int_{\mathcal{X}} \left(\left(\frac{\pi_k(x)}{\hat{\pi}_{k-1}(x)} \right)^{\frac{1}{2}} - \left(\frac{\hat{\pi}_k(x)}{\hat{\pi}_{k-1}(x)} \right)^{\frac{1}{2}} \right)^2 \hat{\pi}_{k-1}(x) \lambda(x) \, dx \right)^{\frac{1}{2}} \\ &= \sqrt{\bar{z}_{k-1}} \left(\int_{\mathcal{X}} \left(\left(\frac{\pi_k(x)}{\hat{\pi}_{k-1}(x)} \right)^{\frac{1}{2}} - \left(\frac{\hat{\pi}_k(x)}{\hat{\pi}_{k-1}(x)} \right)^{\frac{1}{2}} \right)^2 \hat{f}_{\hat{X}^{k-1}}(x) \, dx \right)^{\frac{1}{2}}. \end{aligned}$$

Applying the change of variables $u = \bar{T}_{k-1}^{-1}(x)$ and Equations (51) and (52), we have

$$\begin{aligned} \frac{\pi_k(x)}{\hat{\pi}_{k-1}(x)} &= \frac{(\pi_k \circ \bar{T}_{k-1})(u)}{(\hat{\pi}_{k-1} \circ \bar{T}_{k-1})(u)} = \frac{q_k(u)^2 \omega(u)}{f_U(u)}, \\ \frac{\hat{\pi}_k(x)}{\hat{\pi}_{k-1}(x)} &= \frac{(\hat{\pi}_k \circ \bar{T}_{k-1})(u)}{(\hat{\pi}_{k-1} \circ \bar{T}_{k-1})(u)} = \frac{(\gamma_k + \tilde{g}_k(u)^2) \omega(u)}{f_U(u)}. \end{aligned}$$

In addition, Corollary 3 implies $\hat{f}_{\hat{X}^{k-1}}(x) \, dx = f_U(u) \, du$. Thus, we have

$$\begin{aligned} \|\sqrt{\pi_k} - \sqrt{\hat{\pi}_k}\|_{L^2_\lambda(\mathcal{X})} &= \sqrt{\bar{z}_{k-1}} \left(\int_{\mathcal{U}} \left(q_k(u) - (\gamma_k + \tilde{g}_k(u)^2)^{\frac{1}{2}} \right)^2 \omega(u) \, du \right)^{\frac{1}{2}} \\ &\leq \sqrt{\bar{z}_{k-1}} \left(\int_{\mathcal{U}} \left(q_k(u) - \tilde{g}_k(u) \right)^2 \omega(u) \, du + \gamma_k \omega(\mathcal{U}) \right)^{\frac{1}{2}} \\ &\leq \sqrt{2\bar{z}_{k-1}}\epsilon_k, \end{aligned}$$

where the second last inequality follows from the same proof of Proposition 4. □

Theorem 3 (Approximate ratio approach) *Suppose the initial TT decomposition $\tilde{g}_0 \approx \sqrt{\pi_0}$ and the constant γ_0 satisfy*

$$\|\tilde{g}_0 - \sqrt{\pi_0}\|_{L^2_\omega(\mathcal{U})} \leq \epsilon_0 \quad \text{and} \quad \gamma_0 \leq \frac{1}{\lambda(\mathcal{X})} \|\tilde{g}_0 - \sqrt{\pi_0}\|_{L^2_\omega(\mathcal{U})}^2,$$

respectively. At the k th DIRT step ($k > 0$), suppose further the TT decompositions $\tilde{g}_j \approx \tilde{q}_j$ and the constants γ_j satisfy

$$\|\tilde{g}_j - \tilde{q}_j\|_{L^2_\omega(\mathcal{U})} \leq \epsilon_j \quad \text{and} \quad \gamma_j \leq \frac{1}{\omega(\mathcal{U})} \|\tilde{g}_j - \tilde{q}_j\|_{L^2_\omega(\mathcal{U})}^2, \quad \text{for } j = 1, \dots, k, \quad (56)$$

respectively, where \tilde{q}_j is defined in (54). Then, the unnormalised PDF of DIRT defined by (51) approximates the k th unnormalised bridging density function with the error

$$\|\sqrt{\pi_k} - \sqrt{\hat{\pi}_k}\|_{L^2_\lambda(\mathcal{X})} \leq \sqrt{2} c_{k,0} \epsilon_0 + \sum_{j=1}^k \sqrt{2} c_{k,j} \bar{z}_{j-1} \epsilon_j,$$

where $c_{k,j} = \sup_{x \in \mathcal{X}} r_{k,j}(x)$ is given in Assumption 1, and $\bar{z}_k = \prod_{j=0}^k \hat{z}_j$.

Proof The difference between $\sqrt{\pi_k}$ and $\sqrt{\hat{\pi}_k}$ can be written as

$$\begin{aligned} \sqrt{\pi_k} - \sqrt{\hat{\pi}_k} &= \left(\sqrt{\frac{\pi_k}{\pi_0} \pi_0} - \sqrt{\frac{\pi_k}{\pi_0} \hat{\pi}_0} \right) + \sum_{j=1}^k \left(\sqrt{\frac{\pi_k}{\pi_j} \frac{\pi_j}{\pi_{j-1}} \hat{\pi}_{j-1}} - \sqrt{\frac{\pi_k}{\pi_j} \frac{\hat{\pi}_j}{\hat{\pi}_{j-1}} \hat{\pi}_{j-1}} \right) \\ &= \sqrt{r_{k,0}} (\sqrt{\pi_0} - \sqrt{\hat{\pi}_0}) + \sum_{j=1}^k \sqrt{r_{k,j}} \left(\sqrt{r_{j,j-1}} - \sqrt{\frac{\hat{\pi}_j}{\hat{\pi}_{j-1}}} \right) \sqrt{\hat{\pi}_{j-1}}. \end{aligned}$$

Then, we have $\|\sqrt{\pi_k} - \sqrt{\hat{\pi}_k}\|_{L^2_\lambda(\mathcal{X})} \leq I_0 + \sum_{j=1}^k I_j$, where

$$I_0 = \|\sqrt{r_{k,0}} (\sqrt{\pi_0} - \sqrt{\hat{\pi}_0})\|_{L^2_\lambda(\mathcal{X})} \quad \text{and} \quad I_j = \|\sqrt{r_{k,j}} (\sqrt{r_{j,j-1}} - \sqrt{\frac{\hat{\pi}_j}{\hat{\pi}_{j-1}}}) \sqrt{\hat{\pi}_{j-1}}\|_{L^2_\lambda(\mathcal{X})}.$$

Recalling that $c_{k,j} = \sup_{x \in \mathcal{X}} r_{k,j}(x)$ and applying Proposition 4, we have

$$I_0 \leq \sqrt{c_{k,0}} \|\sqrt{\pi_0} - \sqrt{\gamma_0 + \tilde{g}_0^2}\|_{L^2_\lambda(\mathcal{X})} \leq \sqrt{2} c_{k,0} \epsilon_0.$$

Applying the change of variables $u = \tilde{T}_{j-1}^{-1}(x)$ and Corollary 3 for each $j > 0$, we have

$$r_{j,j-1}(x) = (r_{j,j-1} \circ \tilde{T}_{j-1})(u) = \frac{\tilde{q}_j(u)^2 \omega(u)}{f_U(u)}, \quad \frac{\hat{\pi}_j(x)}{\hat{\pi}_{j-1}(x)} = \frac{(\gamma_j + \tilde{g}_j(u)^2) \omega(u)}{f_U(u)}$$

and $\int \hat{f}_{\hat{\mathcal{X}}_{j-1}}(x) dx = \int f_U(u) du$. This leads to

$$\begin{aligned} I_j &\leq \sqrt{c_{k,j}} \left(\int_{\mathcal{X}} \left(r_{j,j-1}(x)^{\frac{1}{2}} - \left(\frac{\hat{\pi}_j(x)}{\hat{\pi}_{j-1}(x)} \right)^{\frac{1}{2}} \right)^2 \hat{\pi}_{j-1}(x) \lambda(x) dx \right)^{\frac{1}{2}} \\ &= \sqrt{c_{k,j} \bar{z}_{j-1}} \left(\int_{\mathcal{X}} \left(r_{j,j-1}(x)^{\frac{1}{2}} - \left(\frac{\hat{\pi}_j(x)}{\hat{\pi}_{j-1}(x)} \right)^{\frac{1}{2}} \right)^2 \hat{f}_{\hat{\mathcal{X}}_{j-1}}(x) dx \right)^{\frac{1}{2}} \\ &= \sqrt{c_{k,j} \bar{z}_{j-1}} \left(\int_{\mathcal{U}} \left(\tilde{q}_j(u) - (\gamma_j + \tilde{g}_j(u)^2)^{\frac{1}{2}} \right)^2 \omega(u) du \right)^{\frac{1}{2}} \\ &\leq \sqrt{c_{k,j} \bar{z}_{j-1}} \left(\int_{\mathcal{U}} \left(\tilde{q}_j(u) - \tilde{g}_j(u) \right)^2 \omega(u) du + \gamma_j \omega(\mathcal{U}) \right)^{\frac{1}{2}} \\ &\leq \sqrt{2 c_{k,j} \bar{z}_{j-1}} \epsilon_j. \end{aligned}$$

Thus, the result follows. □

Remark 11 At first glance, it appears that Theorem 2 gives smaller errors than Theorem 3. However, this assumes that the two ratio functions in (52) and (54) are approximated with the *same* TT error ϵ_k . Ideally, this should also require the same number of degrees of freedom in TT cores. In practice, this may not be the case: the exact ratio (53) carries the previous approximation errors in the term $(\pi_k \circ \bar{T}_k)(u) / (\hat{\pi}_k \circ \bar{T}_k)(u)$, which can have a complicated structure that is difficult to approximate in TT. In contrast, the approximate ratio involves only the target densities. For example, if the bridging densities π_k were introduced by tempering, the ratio $r_{k+1,k} = \pi^{\beta_{k+1} - \beta_k}$ is just another tempered density. For this reason, DIRT built by the approximate ratio approach may be more accurate in practice.

Corollary 4 Given π_k and $\hat{\pi}_k$ constructed using either the exact or the approximate ratio functions, we suppose the error of $\sqrt{\hat{\pi}_k}$ satisfies

$$\| \sqrt{\pi_k} - \sqrt{\hat{\pi}_k} \|_{L^2_{\lambda}(\mathcal{X})} \leq e_k.$$

Then, the Hellinger distance and the total variation distance between the k th bridging measure ν_k and the pushforward measure $(\bar{T}_k)_{\#} \mu$ satisfy

$$D_H(\nu_k \| (\bar{T}_k)_{\#} \mu) \leq \frac{\sqrt{2} e_k}{\sqrt{z_k}} \quad \text{and} \quad D_{TV}(\nu_k \| (\bar{T}_k)_{\#} \mu) \leq \frac{2 e_k}{\sqrt{z_k}},$$

respectively. The χ^2 -divergence of ν_k from $(\bar{T}_k)_{\#} \mu$ satisfies

$$D_{\chi^2}(\nu_k \| (\bar{T}_k)_{\#} \mu) \leq \left(\nu_k \left(\frac{\pi_k^2}{\hat{\pi}_k^2} \right)^{\frac{1}{2}} + ((\bar{T}_k)_{\#} \mu) \left(\frac{\pi_k^2}{\hat{\pi}_k^2} \right)^{\frac{1}{2}} \right) \frac{2 \bar{z}_k}{z_k \sqrt{z_k}} e_k.$$

Proof The results can be obtained by applying the same proofs of Theorem 1, Corollary 1, and Corollary 2, respectively. Note that for the χ^2 -divergence, the condition $\sup_{x \in \mathcal{X}} \pi_k(x) / \hat{\pi}_k(x) < \infty$ holds as discussed in Remark 9. □

5 Debiasing

Applying either SIRT or DIRT, one can obtain an approximate map $T : \mathcal{U} \mapsto \mathcal{X}$ that enables the simulation of a random variable $\hat{X} \sim T_{\#} \mu$ approximating the target random variable $X \sim \nu_{\pi}$. In a situation where SIRT and DIRT have high accuracy in approximating the target measure, one can approximate the expectation $\nu_{\pi}(h)$ of a function of interest $h(x)$ directly, using the expectation of $h(x)$ over $T_{\#} \mu$, i.e. $(T_{\#} \mu)(h) \equiv \mu(h \circ T)$. The bias of the approximated expectation is proportional to the Hellinger distance $D_H(\nu_{\pi} \| T_{\#} \mu)$ (see Proposition 6). In addition, we can apply the approximate inverse Rosenblatt transport T within the Metropolis–Hastings method and importance sampling to reduce the bias in computing $\nu_{\pi}(h)$. For the sake of completeness, here we discuss some debiasing strategies based on existing work.

Algorithm 1 IRT-MCMC

- 1: Choose an initial state $X^{(0)} = x^*$ for the Markov chain.
- 2: **for** $j = 1, 2, \dots, N$ **do**
- 3: Draw $U \sim \mu$ and compute the proposal candidate $\hat{X} = T(U)$.
- 4: Given the previous state of the Markov chain $X^{(j-1)} = x$, with probability

$$\alpha(x, \hat{X}) = \min \left[1, \frac{f_X(\hat{X}) \hat{f}_{\hat{X}}(x)}{f_X(x) \hat{f}_{\hat{X}}(\hat{X})} \right], \tag{57}$$

- accept \hat{X} by setting $X^{(j)} = \hat{X}$, otherwise set $X^{(j)} = x$.
 - 5: Evaluate the function $H_j = h(X^{(j)})$.
 - 6: **end for**
 - 7: Estimate $\nu_{\pi}(h)$ by the sample average $\frac{1}{N} \sum_{j=1}^N H_j$.
-

We first consider the IRT-MCMC (Algorithm 1), in which the approximate IRT is used as a proposal mechanism in the Metropolised independent sampler for constructing a Markov chain of random variables that converges to the target measure. In the acceptance probability (57), $f_X(\cdot)$ is the PDF of the target measure ν_{π} and $\hat{f}_{\hat{X}}(\cdot)$ is the PDF of $T_{\#} \mu$ that is defined by either the SIRT (19) or the DIRT (50). Following the result of Mengersen and Tweedie [41], the bounds discussed in Remarks 2 and 9 can guarantee the uniform ergodicity of the Markov chain constructed by Algorithm 1. In addition, the average rejection probability is bounded by $2 D_{TV}(\nu_{\pi} \| T_{\#} \mu)$, see Lemma 1 of [17]. This provides an indicator on the performance of the Metropolised independent sampler. However, our bound on $D_{TV}(\nu_{\pi} \| T_{\#} \mu)$ does not directly connect to the bound on the convergence rate of the Metropolised independent sampler, in which the use of acceptance/rejection may require a more precise control on the pointwise error, e.g. $\|\tilde{g} - \sqrt{\pi}\|_{L^{\infty}(\mathcal{X})}$, to assess the convergence rate of the sampler.

One can also employ the approximate IRT built by either the SIRT or the DIRT as the biasing distribution in importance sampling, which leads to the IRT-IS algorithm (Algorithm 2). Compared to IRT-MCMC, IRT-IS generates random variables $\hat{X}^{(j)}$ from the approximate IRT and correct the bias using the weights W_j . By avoiding the Markov chain, importance sampling offers several advantages over the Metropolised

Algorithm 2 IRT-IS

-
- 1: **for** $j = 1, 2, \dots, N$ **do**
 - 2: Draw $U^{(j)} \sim \mu$ and compute the approximate target random variable $\hat{X}^{(j)} = T(U^{(j)})$.
 - 3: Evaluate the unnormalised weight

$$W_j = w(\hat{X}^{(j)}) := \frac{\pi(\hat{X}^{(j)})\lambda(\hat{X}^{(j)})}{\hat{f}_{\hat{X}}(\hat{X}^{(j)})}$$

and the function $\hat{H}_j = h(\hat{X}^{(j)})$.

- 4: **end for**
 - 5: Estimate the normalising constant using $\bar{z}_N = \frac{1}{N} \sum_{j=1}^N W_j$.
 - 6: Compute the sample average $\bar{h}_N = \frac{1}{N} \sum_{j=1}^N W_j \hat{H}_j$.
 - 7: Estimate $v_\pi(h)$ by the ratio estimator $I_N = \bar{h}_N / \bar{z}_N$.
-

independent sampler: (i) it can be easily parallelised; and (ii) variance reduction techniques such as antithetic variable and control variates (see [53, Chapter 4] and references therein) and efficient high-dimensional quadrature methods such as quasi-Monte Carlo [16] can be naturally applied within importance sampling.

The error bounds established in Sects. 3 and 4 offer insights into the efficiency of IRT-IS. As discussed in [48, Chapter 9], for N approximate target random variables, one can use the effective sample size (ESS)

$$\text{ESS}(N) = N \frac{(T_{\#}\mu)(w)^2}{(T_{\#}\mu)(w^2)},$$

where $w(x) = \pi(x)/\hat{\pi}(x)$ for SIRT and $w(x) = \pi(x)/\hat{\pi}_L(x)$ for DIRT, to measure the efficiency of importance sampling for representing the target measure v_π .

Lemma 2 *Given the χ^2 -divergence of v_π from $T_{\#}\mu$, the ESS of Algorithm 2 satisfies*

$$\text{ESS}(N) = \frac{N}{1 + D_{\chi^2}(v_\pi \| T_{\#}\mu)}.$$

Proof Since we have $(T_{\#}\mu)(\pi\lambda/\hat{f}_{\hat{X}}) = z$, where z is the normalising constant of the target density, the χ^2 -divergence of v_π from $T_{\#}\mu$ satisfies

$$D_{\chi^2}(v_\pi \| T_{\#}\mu) = \frac{(T_{\#}\mu)((\pi\lambda/\hat{f}_{\hat{X}})^2)}{((T_{\#}\mu)(\pi\lambda/\hat{f}_{\hat{X}}))^2} - 1 = \frac{(T_{\#}\mu)(w^2)}{(T_{\#}\mu)(w)^2} - 1,$$

in which one can choose w to be the ratio $\pi\lambda/\hat{f}_{\hat{X}}$ multiplied by any nonzero constant. Thus, the result follows. \square

The bounds discussed in Remarks 2 and 9 imply that \bar{z}_N and \bar{h}_N computed by Algorithm 2 are unbiased estimators for the normalising constant z and the expectation $(T_{\#}\mu)(wh)$, respectively. However, the ratio estimator I_N is only asymptotically

unbiased such that

$$\mathbb{P}\left[\lim_{N \rightarrow \infty} I_N = v_\pi(h)\right] = 1.$$

For a finite sample size, $N < \infty$, the ratio estimator I_N is a biased estimator of $v_\pi(h)$. However, for sufficiently large sample size N , one can apply the Delta method (see [48, Chapter 2] and references therein) to show that the mean square error (MSE) of I_N yields the approximation

$$\text{MSE}(I_N) \equiv \mathbb{E}[(I_N - v_\pi(h))^2] \approx \frac{1}{N} \frac{(T_\# \mu)(w^2(h - v_\pi(h))^2)}{(T_\# \mu)(w)^2}. \tag{58}$$

Thus, for a sufficiently regular function h , the MSE of the ratio estimator I_N can also be controlled by the χ^2 -divergence $D_{\chi^2}(v_\pi \| T_\# \mu)$.

6 Numerical Examples

We demonstrate the efficiency and various aspects of DIRT, which employs SIRT within each layer, using four Bayesian inference problems arising in dynamical systems and PDEs. In all numerical examples, as efficiency measures of DIRT (or SIRT in the single layer case), we employ the estimated integrated autocorrelation time (IACT) for IRT-MCMC (Algorithm 1) and the ratio between the total number of samples and the estimated ESS, $N/\text{ESS}(N)$, for IRT-IS (Algorithm 2). For both IACT and $N/\text{ESS}(N)$, a lower value indicates a better sampling efficiency. The minimum value of both IACT and $N/\text{ESS}(N)$ is 1. Since $N/\text{ESS}(N) = 1 + D_{\chi^2}(v_k \| T_\# \mu)$, it also measures directly the accuracy of DIRT for approximating the posteriors. Matlab implementation of IRT methods and numerical examples is available at <https://github.com/dolgov/TT-IRT>.

6.1 Predator and Prey

The predator–prey model is a system of coupled ODEs frequently used to describe the dynamics of biological systems. The populations of predator (denoted by Q) and prey (denoted by P) change over time according to a pair of ODEs

$$\begin{cases} \frac{dP}{dt} = rP\left(1 - \frac{P}{K}\right) - s\left(\frac{PQ}{\alpha + P}\right), \\ \frac{dQ}{dt} = u\left(\frac{PQ}{\alpha + P}\right) - vQ, \end{cases} \tag{59}$$

with initial conditions $P(t = 0) = P_0$ and $Q(t = 0) = Q_0$. The dynamical system is controlled by several parameters. In the absence of the predator, the population of the prey evolves according to the logistic equation characterised by r and K . In the absence of the prey, the population of the predator decreases exponentially with a rate

v . In addition, the two populations have a nonlinear interaction characterised by α , s , and u . We often do not know the initial populations and the parameters r , K , α , s , u , and v . This way, we need to estimate unknowns

$$x = [P_0, Q_0, r, K, \alpha, s, u, v]^\top,$$

from observed populations of the predator and prey at time instances t_i for $i = 1, \dots, n_T$.

6.1.1 Posterior Density

Let $y \in \mathbb{R}^{2n_T}$ denote the observed populations of the predator and prey. We define a forward model $G : \mathcal{X} \mapsto \mathbb{R}^{2n_T}$ in the form of $G(x) = [P(t_i), Q(t_i)]_{i=1}^{n_T}$ to represent the populations of the predator and prey computed at $\{t_i\}_{i=1}^{n_T}$ for a given parameters x . Assuming independent and identically distributed (i.i.d.) normal noise in the observed data and assigning a prior density $\pi_0(x)$ to the unknown parameter, one can define the unnormalized posterior density

$$\pi(x) \propto \exp\left(-\frac{1}{2\sigma^2}\|G(x) - y\|_2^2\right) \pi_0(x),$$

where σ is the standard deviation of the normally distributed noise. Synthetic observed data are used in this example. With $n_T = 13$ time instances $t_i = (i - 1) \times 25/6$ and a given parameter $x_{\text{true}} = [50, 5, 0.6, 100, 1.2, 25, 0.5, 0.3]^\top$, we generate synthetic noisy data $y = y_{\text{true}} + \eta$, where η is a realization of the i.i.d. zero mean normally distributed noise with the standard derivation $\sigma = \sqrt{2}$. A uniform prior density $\pi_0(x) = \prod_{k=1}^8 \mathbb{I}_{[a_k, b_k]}(x_k)$ is specified to restrict the support of x_k to the interval $[a_k, b_k]$, where $a = [30, 3, 0.36, 60, 0.72, 15, 0.3, 0.18]^\top$ and $b = [80, 8, 0.96, 160, 1.92, 40, 0.8, 0.48]^\top$. To illustrate the behaviour of the posterior density, we plot the kernel density estimates of the marginal posterior densities in Fig. 2. Note that some of the parameters are significantly correlated, which makes the posterior density function difficult to explore by both MCMC and a straightforward TT approximation.

6.1.2 Numerical Results

We use $L = 8$ bridging measures in the construction of DIRT by tempering the unnormalised posterior density with $\pi_k(x) = \pi(x)^{\beta_k}$, starting from $\beta_0 = 10^{-4}$ and following by $\beta_{k+1} = \sqrt{10} \cdot \beta_k$. This way, $\beta_L = 1$ gives the target probability density. We consider two reference measures: the uniform reference measure μ_{uni} and the truncated normal reference measure μ_{TG} with the density $f_U(u) \propto \prod_{k=1}^8 \mathbb{I}_{[-4, 4]}(u_k) \exp(-\|u_k\|_2^2/2)$. Note that at layer 0, the ratio function is just the tempered density $\pi_0(x)$ in the original domain $x_k \in [a_k, b_k]$. We employ the piecewise-linear basis functions with n equally spaced interior collocation points for both reference measures. In addition, we tune TT-cross (Algorithm 4) using three parameters: the initial TT rank R_0 , enrichment TT ranks $\rho_1 = \dots = \rho_{d-1} = \text{Rh}_0$, and the




Fig. 2 Marginal posterior densities of the predator–prey model estimated from 10^6 posterior samples

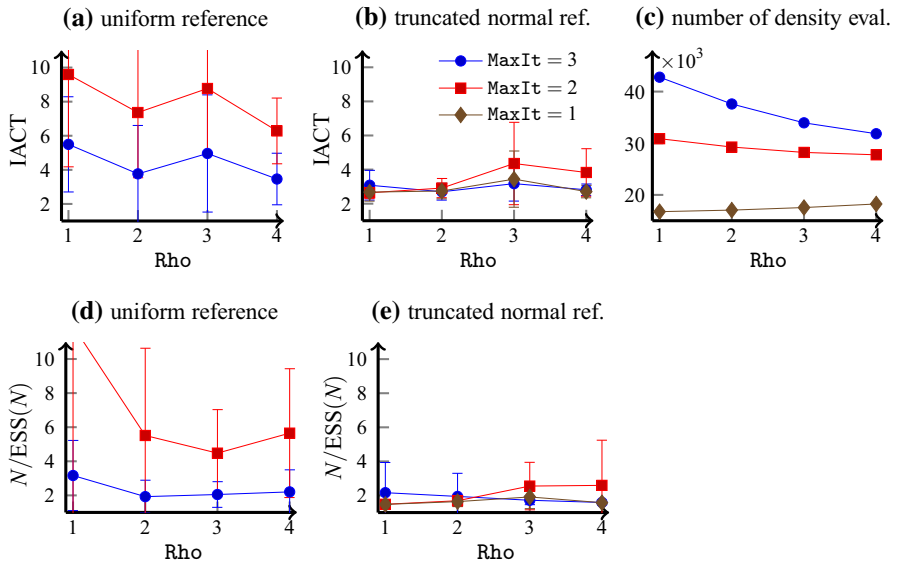


Fig. 3 $IAC\pm$ standard deviation over 10 runs with the uniform reference measure (a). $IAC\pm$ standard deviation with the truncated normal reference measure (b). Number of density evaluations in TT-cross at each layer (c). $N/ESS(N)\pm$ standard deviation over 10 runs with the uniform reference measure (d). $N/ESS(N)\pm$ standard deviation with the truncated normal reference measure (e). Initial TT rank is adjusted such that the maximal TT rank is 13 in all tests, whereas enrichment ranks Rho and numbers of TT-cross iterations $MaxIt$ are varied

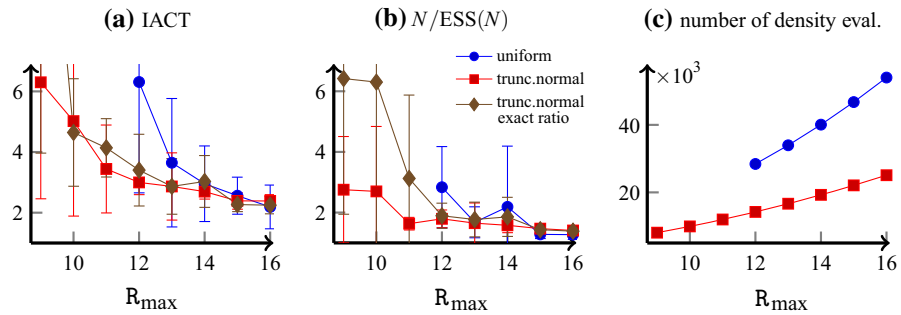


Fig. 4 IACT \pm standard deviation over 10 runs (a), $N/ESS(N) \pm$ standard deviation over 10 runs (b), and number of density evaluations in TT-cross at each layer (c) with varying maximum TT ranks R_{\max} and different reference measures. For the uniform reference, $\text{MaxIt} = 3$ and $\text{Rho} = 3$ are used. For the truncated normal reference, $\text{MaxIt} = 1$ and $\text{Rho} = 0$ are used




Fig. 5 IACT \pm standard deviation over 10 runs (a), $N/ESS(N) \pm$ standard deviation over 10 runs (b), and number of density evaluations in TT-cross at each layer (c) for varying numbers of collocation points n and different reference measures. For the uniform reference, $\text{MaxIt} = 3$ and $\text{Rho} = 3$ are used. For the truncated normal reference, $\text{MaxIt} = 1$ and $\text{Rho} = 0$ are used

maximum number of TT-cross iterations MaxIt . Those define uniquely the maximum TT rank $R_{\max} = R_0 + \text{Rho} \cdot \text{MaxIt}$.

Firstly, we vary one tuning variable at a time and investigate its impact on the efficiency and computational cost of the DIRT. We take the number of posterior density function evaluations in TT-cross in each DIRT layer to measure the computational cost for building DIRT.

In Fig. 3, we vary the enrichment rank Rho and the number of TT-cross iterations MaxIt . The initial TT rank R_0 is adjusted such that the maximum TT rank is 13 in all cases. We set the number of collocation points to be $n = 16$. All the DIRTs are constructed using the approximate ratio (54). With each Rho and MaxIt , we repeat the IRT-MCMC and IRT-IS for 10 experiments and report the estimated mean and standard deviation of the efficiency indicators. For the uniform reference (Fig. 3a, d), carrying out $\text{MaxIt} = 1$ iteration gives very inaccurate results with $\text{IACT} > 10$ and $N/ESS(N) > 10$. Increasing the number of TT-cross iterations for the uniform reference measure significantly improves the DIRT accuracy. Since the ratio function varies considerably from layer to layer, TT-cross needs at least 3 iterations and a nontrivial enrichment to adapt the approximation to the new function. This comes

at the expense of tripling the number of density evaluations, in addition to those needed to compute the enrichment, as shown in Fig. 3c. In contrast, using the truncated normal reference measure (Fig. 3b, e) can significantly improve the efficiency in this example. With only one TT-cross iteration, it can reduce the final IACT to below 4 and $N/\text{ESS}(N)$ to below 3.

Remark 12 At levels $k > 0$ of the DIRT construction, the ratio functions may have a similar shape (see Fig. 1). Thus, one can take the TT of the ratio function at the previous level $k > 0$ as the initial guess for building TT at level $k + 1$. This initialization provides good index sets in TT-cross, such that only one TT-cross iteration is sufficient with the truncated normal reference measure.

In Fig. 4, we vary the maximum TT rank R_{\max} . With the uniform reference, we set $R_{\text{ho}} = 3$ and $\text{MaxIt} = 3$. With the truncated normal reference, we set $R_{\text{ho}} = 0$ and $\text{MaxIt} = 1$, which makes the number of density evaluations equal to the number of degrees of freedom in the TT decomposition, $n_1 r_1 + \sum_{k=2}^{d-1} n_k r_{k-1} r_k + r_{d-1} n_d = (d - 2)nR_{\max}^2 + 2nR_{\max}$. We observe that the two reference measures give eventually comparable IACTs and ESSs with increasing R_{\max} . However, the truncated normal reference achieves this with much fewer density evaluations.

In Fig. 4, we compare also the approximate ratio (54) used in all experiments with the exact ratio (52). The diamond-shaped markers in Fig. 4a, b show IACTs and ESSs obtained by the exact ratio approach. In this example, it gives worse results with a larger IACT and $N/\text{ESS}(N)$ for the truncated normal reference measure with lower R_{\max} values and does not lead to any meaningful results for the uniform reference measure.

In Fig. 5, we vary the number of collocation points n used in each dimension. The truncated normal reference starts with a larger error since $n = 10$ points cannot resolve the rather large reference domain $[-4, 4]$. With increasing n , the IACT obtained using the truncated normal reference decays rapidly. In comparison, the IACT obtained using the uniform reference exhibits a spike and does not show rapid decay with increasing n . This may be caused by the boundary layers in the ratio function. Similar trends are observed in the reported $N/\text{ESS}(N)$. Again, the truncated normal reference requires significantly fewer density evaluations to achieve the same level of accuracy compared to the uniform reference in this experiment.

Next, we benchmark DIRT with the truncated normal reference, $\text{MaxIt} = 1$, $R_{\text{ho}} = 0$, $n = 16$, and $R_0 = R_{\max} = 13$ against other sampling algorithms, including the Delayed Rejection Adaptive Metropolis (DRAM) [27], the Stein variational Newton (SVN) [15], and the Hierarchical Invertible Neural Transport (HINT) [35].

DRAM is initialized with the covariance matrix $5I$, adaptation scale $2.4/\sqrt{d}$, adaptation interval 10 and delayed rejection scale 2. These parameters are commonly recommended in general case.

For this example, SVN is sensitive to the choice of the step size and to the initial distribution of particles. We choose the step size to be $2 \cdot 10^{-2}$ and generate the initial particle set from the normal distribution $\mathcal{N}(x_{\text{true}}, (2 \cdot 10^{-2} x_{\text{true}})^2)$, which gives a reasonable balance between the stability and the rate of convergence. We carry out 23 Newton iterations in SVN to approach stationarity.

HINT is an autoregressive normalising flow estimator for the joint probability density $\pi(y, x)$. Since both the prior random variable X with the density $\pi_0(x)$ and the noise random variable η can be directly simulated, drawing samples from the joint distribution is easy. One can first draw a sample X from the prior and then simulate the corresponding data sample $Y = G(X) + \eta$ by generating a noise random variable η . Drawn a set of independent and identically distributed samples $(Y^{(i)}, X^{(i)})_{i=1}^N$ from the joint measure, HINT computes a triangular invertible map $(u_Y, u_X) = S_\theta(y, x)$ from the (joint) target measure to the reference measure by minimizing the maximum likelihood

$$L(\theta) = \sum_{i=1}^N \frac{1}{2} \|S_\theta(Y^{(i)}, X^{(i)})\|_2^2 - \log |\nabla S_\theta(Y^{(i)}, X^{(i)})|$$

over the parameter θ that defines the neural networks S_θ . Given observed data y , this allows one to define a conditional map $x = T_\theta(u_X; y) := (S_\theta^X)^{-1}(y, u_X)$ that maps from the reference measure of U_X to the posterior measure conditioned on data y .

We simulate each method $M = 10$ times with N samples produced in each simulation, denoted by $\{x^{(\ell,j)}\}_{j=1}^N$, where $\ell = 1, \dots, M$ indexes the simulations. For each simulation, we compute the empirical posterior covariance matrix $C^\ell = \frac{1}{N} \sum_{j=1}^N (x^{(\ell,j)} - \bar{x}^\ell)(x^{(\ell,j)} - \bar{x}^\ell)^\top$, where $\bar{x}^\ell = \frac{1}{N} \sum_{j=1}^N x^{(\ell,j)}$ is the empirical posterior mean. Then, we use the average deviation of covariance matrices to benchmark the sampling performance of different sampling algorithms. Here, we employ the Förstner–Moonen distance [19] over the cone of symmetric and positive definite (SPD) matrices,

$$d_{\text{FM}}(A, B) = \sum_{i=1}^d \ln^2(\lambda_i(A, B)),$$

where $\lambda_i(A, B)$ denotes the i th generalised eigenvalue of the pair of SPD matrices (A, B) , to measure the deviation. This way, averaging the Förstner–Moonen distance between the ℓ th empirical covariance matrix and the average covariance matrix over all M simulations,

$$\mathcal{E}_C = \frac{1}{M} \sum_{\ell=1}^M d_{\text{FM}}(C^\ell, \bar{C}), \quad \text{where } \bar{C} = \frac{1}{M} \sum_{\ell=1}^M C^\ell, \tag{60}$$

provides an estimated deviation of empirical covariance matrices computed by a given algorithm.

In Fig. 6, we plot the covariance deviations (60) obtained by IRT-MCMC, DRAM, SVN and HINT versus the total number of target density function evaluations and the total CPU time needed by each algorithm. Here, the reported total numbers of density evaluations and CPU times include the construction of DIRT in each simulation experiment. The 10 independent simulations are run in parallel on a workstation with an Intel Xeon E5-2640v4 CPU at 2.4GHz. We can notice that DIRT produces estimated

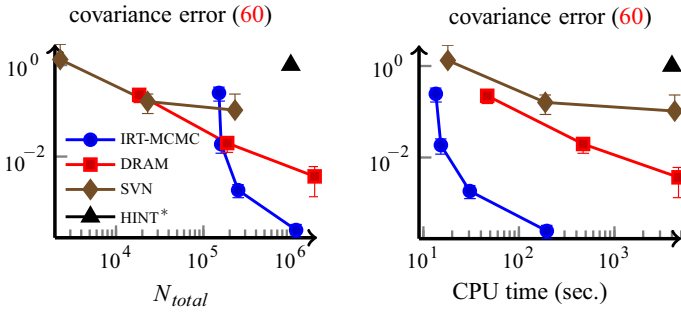


Fig. 6 Estimated deviation of empirical covariance matrices (60) computed by IRT-MCMC, DRAM, SVN and HINT for different total numbers of density evaluations (left) and CPU times (right) (* the marker in the figure is only indicative, and the actual results are larger)

covariance matrices with smallest deviations in almost all tests. Moreover, DIRT is computationally more efficient in terms of the CPU time, because the evaluation of DIRT can take advantage of vector instructions.

In this example, HINT gives the worst results since the joint density estimation from samples is a much higher dimensional problem compared to the posterior density approximation. In particular, the dimension of y in the predator-prey model is $2n_T = 26$, so the total dimensionality of the problem increases to 34. This required us to construct HINT networks with 8 blocks containing 200×200 weights each. This totalled to 2691308 trainable parameters in the entire HINT. We trained the networks using $5 \cdot 10^6$ 10^4 training samples for 50 epochs, consisting of 500 batches of 10^5 samples each. The other (e.g. ADAM) parameters are left unchanged from [35]. The training took 5.8h on a NVidia GeForce GTX 1650 Max-Q GPU card. To avoid disproportionate scaling of axes in Fig. 6 compared to other methods, we put just an indicative marker for HINT. The actual error (60) with 10^5 test samples taken directly from HINT was 16.7. Using the test samples as proposals in the MCMC rejection against the exact posterior gives a rejection rate of 97% and IACT of 127, and the covariance matrix computed from the rejected samples gives the Förstner–Moonen distance of 0.1. This indicates that the data-driven joint density estimation should be more applicable for lower-dimensional data, whereas if one is only interested in the posterior, the function approximation methods seem to be a better choice.

6.2 Lorenz-96

This is a widely used benchmark model in atmospheric circulations. We consider a Lorenz-96 model that is specified by the system of ODEs

$$\frac{dP_i(t)}{dt} = (P_{i+1} - P_{i-2})P_{i-1} - P_i + 8, \quad \text{for } i = 1, \dots, d, \quad (61)$$

with periodic boundary conditions and an unknown initial condition $P_i(0) = x_i$ for $i = 1, \dots, d$. The state dimension is set to $d = 40$. Observing noisy states with even




Fig. 7 Lorenz-96 model, true initial state (solid red) and posterior expectation, mean (blue dashed) ± 2 standard deviations (shaded area) (Color figure online)

indices at the final time $T = 0.1$, we aim to infer the initial state x in this example. This way, we have observed data $y \in \mathbb{R}^{\frac{d}{2}}$ and can define a forward model $G : \mathcal{X} \mapsto \mathbb{R}^{\frac{d}{2}}$ in the form of $G(x) = [P_{2k}(T)]_{k=1}^{\frac{d}{2}}$ to represent simulated observables for a given initial condition x .

Assuming i.i.d. normal noise in the observed data and assigning a truncated normal prior density to the initial condition, we have the unnormalized posterior density

$$\pi(x) = \exp\left(-\frac{1}{2\sigma^2} \|G(x) - y\|_2^2\right) \prod_{k=1}^d \left(\mathbb{I}_{[-10,10]}(x_k) \exp\left(-\frac{1}{2}(x_k - 1)^2\right)\right).$$

We use a synthetic data set $y = G(x_{\text{true}}) + \eta$, where x_{true} is drawn from $\sim \mathcal{N}(1, 10^{-4}I_d)$, and η is a realisation of the i.i.d. zero mean normal noise with the standard deviation $\sigma = 10^{-1}$.

For the TT-cross approximations, we use the truncated normal reference measure on $[-3, 3]^d$, piecewise linear basis functions with $n = 15$ interior collocation points, $\text{MaxIt} = 1$ TT-cross iteration, and TT ranks $R_{\text{max}} = 15$. DIRT is built with the tempered density

$$\pi_k(x) = \exp\left(-\frac{\beta_k}{2\sigma^2} \|G(x) - y\|_2^2\right) \cdot \prod_{k=1}^d \left(\mathbb{I}_{[-10,10]}(x_k) \exp\left(-\frac{\beta_k^{0.25}}{2}(x_k - 1)^2\right)\right),$$

with $\beta_0 = 10^{-4}$ and $\beta_{k+1} = \sqrt{10} \cdot \beta_k$. This way, we need $L = 8$ layers to reach the posterior density. A weaker tempering of the prior is used to reduce its impact on the intermediate levels. This allows most of the intermediate DIRT levels to be used to bridge the more complicated likelihood. This setup requires a total of 1.2×10^6 density evaluations in TT-cross at all layers and provides an average ESS of $N/1.55$ in IRT-IS and an average IACT of 2.6 in IRT-MCMC.

Using the posterior density, we can quantify the uncertainty of the inferred initial state and make predictions of the terminal state. The predicted initial state is shown in Fig. 7. Note that the chaotic regime of Lorenz-96 makes it difficult to predict the

unobserved odd coordinates. Nevertheless, DIRT demonstrates high numerical and sampling efficiency in approximating this complicated posterior.

6.3 Elliptic PDE

In the third example, we apply both SIRT and DIRT to the classical inverse problem governed by the stochastic diffusion equation

$$-\nabla \cdot (\kappa_d(s; x)\nabla u(s)) = 0 \quad \text{on } s \in D := (0, 1)^2, \tag{62}$$

with Dirichlet boundary conditions $u|_{s_1=0} = 1$ and $u|_{s_1=1} = 0$ on the left and right boundaries, and homogeneous Neumann conditions on other boundaries. The goal is to infer the unknown diffusion coefficient $\kappa_d(s; x)$ from incomplete observations of the potential function $u(s)$. Here, we adopt the same setup used in [17,56].

6.3.1 Posterior Density

The unknown diffusion coefficient $\kappa_d(s; x)$ is parametrized by a d -dimensional random variable x . We take each of the parameters $x_k, k = 1, \dots, d$, to be uniformly distributed on $[-\sqrt{3}, \sqrt{3}]$. Then, for any $x \in [-\sqrt{3}, \sqrt{3}]^d$ and $s = (s_1, s_2) \in D$, the logarithm of the diffusion coefficient at s is defined by the following expansion

$$\ln \kappa_d(s; x) = \sum_{k=1}^d x_k \sqrt{\eta_k} \cos(2\pi\rho_1(k)s_1) \cos(2\pi\rho_2(k)s_2), \tag{63}$$

where

$$\eta_k = \frac{k^{-(\nu+1)}}{\sum_{k=1}^d k^{-(\nu+1)}}, \quad \rho_1(k) = k - \frac{\tau(k)^2 + \tau(k)}{2}, \quad \text{and} \quad \rho_2(k) = \tau(k) - \rho_1(k),$$

with $\tau(k) = \lfloor \frac{1}{2}(\sqrt{1+k/2} - 1) \rfloor$. To discretise the PDE in (62), we tessellate the spatial domain D with a uniform Cartesian grid with mesh size h . Then, we replace the infinite dimensional solution $u \in V \equiv H^1(D)$ by the continuous, piecewise bilinear finite element (FE) approximation $u_h \in V_h$ associated with the discretisation grid. To find u_h , we solve the resulting Galerkin system using a sparse direct solver. A fixed discretisation with $d = 11, h = 2^{-6}$, and $\nu = 2$ is used in this example.

The observed data $y \in \mathbb{R}^m$ consist of m local averages of the potential function $u(s)$ over subdomains $D_i \subset D, i = 1, \dots, m$. To simulate the observable model outputs, we define the forward model $G^h : \mathcal{X} \mapsto \mathbb{R}^m$ with

$$G_i^h(x) = \frac{1}{|D_i|} \int_{D_i} u_h(s; x) ds, \quad i = 1, \dots, m.$$

The subdomains D_i are squares with side length $2/(\sqrt{m} + 1)$ centred at the interior vertices of a uniform Cartesian grid on $D = [0, 1]^2$ with grid size $1/(\sqrt{m} + 1)$,

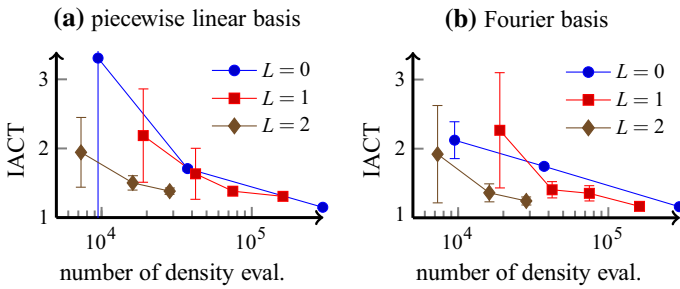


Fig. 8 Elliptic PDE with $\sigma^2 = 10^{-2}$ and $m = 3^2$. IACT vs. number of density evaluations in DIRT for different number of layers: $L = 0$ ($\beta = 1$), $L = 1$ ($\beta = \{0.1, 1\}$) and $L = 2$ ($\beta = \{0.1, \sqrt{0.1}, 1\}$). Note R_{\max} varying from 8 to 32 for $L = 0$, but only from 4 to 8 for $L = 2$

which form an overlapping partition of D . Synthetic data for these m local averages are produced from the “true” parameter $x_{\text{true}} = (1.5, \dots, 1.5)$ by adding i.i.d. zero mean normally distributed noise with the standard deviation σ . This way, we have the unnormalized posterior density

$$\pi(x) = \exp\left(-\frac{1}{2\sigma^2}\|G^h(x) - y\|_2^2\right) \prod_{k=1}^d \left(\mathbb{I}_{[-\sqrt{3}, \sqrt{3}]}(x_k)\right).$$

6.3.2 Numerical Results

In this example, we compare the impact of different tempering schemes, different numbers of measurements, and different measurement noise levels on DIRT. We also compare different basis functions used in the DIRT construction. In all experiments, we feed $N = 2^{16}$ independent samples generated by DIRT to both IRT-MCMC and IRT-IS.

In Fig. 8, we compare DIRT with three different tempering sequences $\beta = [\beta_0, \dots, \beta_L]$, varying the grid size n and the TT ranks R_{\max} . Note that with $L = 0$, we have the single-layer SIRT. The reported number of density function evaluations is a sum of the numbers of evaluations in TT-cross at all layers. We use the truncated normal reference measure on $[-4, 4]^d$ with both piecewise linear and Fourier bases for the multilayer DIRT.

With the noise variance $\sigma^2 = 10^{-2}$ and a rather small data size $m = 3^2$, the posterior density is relatively simple to characterise and hence, can be tackled directly using the single-layer SIRT (see the case $L = 0$ in Fig. 8 and [17]). However, the multilayer DIRT uses much smaller number of collocation points and TT ranks for producing an approximate posterior density with the same accuracy. Here, the three-layer DIRT needs only 10% of the density evaluations required for the single-layer counterpart.

Next, we test the multilayer DIRT on more difficult posterior densities, with larger numbers of measurements and smaller observation noise. We set the number of collocation points to be $n = 16$, maximum TT-cross iteration to be $\text{MaxIt} = 1$, and maximum TT rank to be $R_{\max} = 12$. In Fig. 9, we fix $\sigma^2 = 10^{-2}$ and vary the number




Fig. 9 Elliptic PDE with varying numbers of measurements m . Total number of density evaluations in all layers (a), reciprocal sample size (b), and IACT (c). Tempering is carried out with $\beta_0 = 4^{-\lceil \log_4 m \rceil}$, $\beta_{k+1} = 4 \cdot \beta_k$. TT-cross parameters: $n = 16$, $R_{\max} = R_0 = 12$, and $\text{MaxIT} = 1$

of measurements. Since halving the measurement grid size $2/(\sqrt{m} + 1)$ corresponds to multiplying m by approximately a factor of 4, we use a different tempering strategy, starting with $\beta_0 = 4^{-\lceil \log_4 m \rceil}$, and setting $\beta_{k+1} = 4 \cdot \beta_k$ for next layers. This way, the number of layers grows proportionally to $\log m$, and the number of density evaluations in TT-cross for fixed TT ranks is also proportional to $\log m$, which can be confirmed by Fig. 9a. Here, we can see that the Fourier basis is significantly more accurate than the piecewise-linear basis for the same grid size. With the linear basis, both IACT and $N/\text{ESS}(N)$ grow logarithmically in the number of measurements. With the Fourier basis, the IACT stays almost constant below 1.5 and the $N/\text{ESS}(N)$ stays almost constant below 1.1, increasing slightly only for the most difficult case with $m = 31^2$ (Fig. 9b, c). With increasing number of measurements, the likelihood becomes more concentrated. This makes it more challenging to characterise the posterior using prior-based approaches such as QMC [56] or single-layer TT approximation. For example, even with a much larger number of collocation points $n = 65$ and 5 iterations of TT-cross (giving a maximal TT rank of 41), we still cannot produce reasonable results for $m = 15^2$ with the single-layer SIRT.

We carry out an additional test with decreasing noise variance σ^2 . In Fig. 10, we fix $m = 15^2$ and vary σ^2 from 10^{-1} to 10^{-5} . In this experiment, fixing TT ranks becomes insufficient for representing posterior densities with low observation noise. In particular, the piecewise linear basis does not have sufficient accuracy for the case of the smallest noise variance. In contrast, the Fourier basis can still retain low IACT and $N/\text{ESS}(N)$ for low noise variance cases, where IACT and $N/\text{ESS}(N)$ grow proportionally to $\log \sigma$. Together with the log-scaling of the number of evaluations, the effective complexity of the entire IRT-MCMC and IRT-IS schemes becomes *poly-logarithmic* in the variance. Although the Fourier basis is computationally more expensive to evaluate than the piecewise-linear basis, with a factor of 2.5 in the worst case scenario in this experiment, this additional computational effort is well compensated by a much higher accuracy. This makes DIRT a viable approach for a range of concentrated distributions.

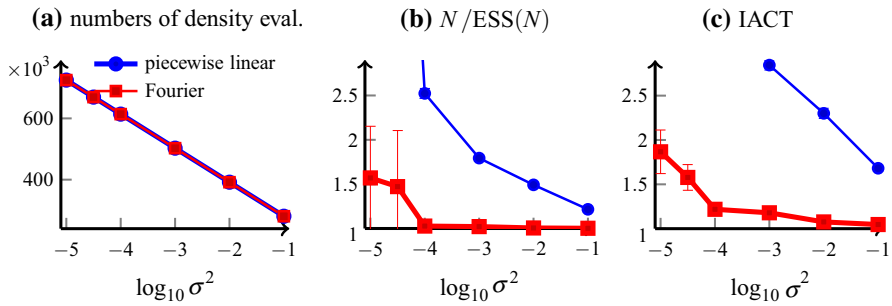


Fig. 10 Elliptic PDE with varying noise variances σ^2 . Total numbers of density evaluations in all layers (a), reciprocal sample size (b), and IACT (c). Tempering is carried out with $\beta_0 = 0.1\sigma^2$, $\beta_{k+1} = \sqrt{10} \cdot \beta_k$. TT parameters: $n = 16$, TT rank 20, one TT-cross iteration

6.4 Parabolic PDE

In the fourth example, we consider an inverse problem of identifying the diffusion coefficient of a two-dimensional parabolic PDE from point observations of its solution. In the problem domain $D = [0, 3] \times [0, 1]$, with boundary ∂D , we model the time-varying potential function $p(s, t)$ for given diffusion coefficient field $\kappa_d(s)$ and forcing function $f(s, t)$ using the heat equation

$$\frac{\partial p(s, t)}{\partial t} = \nabla \cdot (\kappa_d(s; x) \nabla p(s, t)) + f(s, t), \quad s \in D, \quad t \in [0, T], \quad (64)$$

where $T = 10$. Parabolic PDEs of this type are widely used in modeling groundwater flow, optical diffusion tomography, the diffusion of thermal energy, and numerous other common scenarios for inverse problems. Let $\partial D_n = \{s \in \partial D \mid s_2 = 0\} \cup \{s \in \partial D \mid s_2 = 1\}$ denote the top and bottom boundaries, and $\partial D_d = \{s \in \partial \Omega \mid s_1 = 0\} \cup \{s \in \partial \Omega \mid s_1 = 3\}$ denote the left and right boundaries. For $t \geq 0$, we impose the mixed boundary condition:

$$p(s, t) = 0, \quad \forall s \in \partial D_d, \quad \text{and} \quad (\kappa_d(s; \theta) \nabla p(s, t)) \cdot \mathbf{n}(s) = 0, \quad \forall x \in \partial D_n,$$

where $\mathbf{n}(s)$ is the outward normal vector on the boundary. We also impose a zero initial condition, i.e. $p(s, 0) = 0, \forall s \in D$, and let the potential field be driven by a time-invariant forcing function

$$f(s, t) = c \left(\exp \left(-\frac{1}{2r^2} \|s - a\|^2 \right) - \exp \left(-\frac{1}{2r^2} \|s - b\|^2 \right) \right), \quad \forall t \geq 0,$$

with $r = 0.05$, which is the superposition of two normal-shaped sink/source terms centered at $a = (0.5, 0.5)$ and $b = (2.5, 0.5)$, scaled by a constant $c = 5\pi \times 10^{-5}$.




Fig. 11 Setup of the parabolic example. Logarithm of the “true” diffusion coefficient and $d = 27$ centers of the process convolution prior (blue crosses) (a). The potential function $p_h(s, t; x_{\text{true}})$ at $t = 0.1$, computed with $h = 1/80$ (b). The potential function $p_h(s, t; x_{\text{true}})$ at $t = 10$ (c). Black dots in (b) and (c) are locations of measurements

6.4.1 Posterior Density

The logarithm of the diffusion coefficient, $\ln \kappa_d(s; x)$, is endowed with the process convolution prior [29],

$$\ln \kappa_d(s; x) = \ln \bar{\kappa} + \sum_{k=1}^d x_k \exp\left(-\frac{1}{2} \|s - s^{(k)}\|^2\right), \tag{65}$$

where $d = 27$, $\ln \bar{\kappa} = -5$, each coefficient x_k follows a standard normal prior $\mathcal{N}(0, 1)$ (which can be truncated to $[-5, 5]$ with sufficient accuracy), and $s^{(k)}$, $k = 1, \dots, d$ are centers of the kernel functions (shown as blue crosses in Fig. 11a). Similarly to the previous example, the potential function $p(s, t)$ in (64) is approximated by $p_h(s, t)$ using the finite element method with piecewise bilinear basis functions and implicit Euler time integration.

The observed data $y \in \mathbb{R}^{m \times n_T}$ consist of the time-varying potential function $p(s, t)$ measured at $m = 13$ locations (shown as black dots in Fig. 11b, c) at $n_T = 10$ discrete time points equally spaced between $t = 1$ and $t = 10$. To simulate the observable model outputs, we define the forward model $G^h : \mathcal{X} \mapsto \mathbb{R}^{m \times n_T}$ with

$$G^h_{i,j}(x) = p_h(s_i, t_j; x), \quad i = 1, \dots, m, \quad j = 1, \dots, n_T.$$

Using a “true” parameter x_{true} drawn from the prior distribution and a forward model with $h = 1/80$, synthetic data $y \in \mathbb{R}^{m \times n_T}$ are produced by adding i.i.d. normal noise with zero mean and the standard deviation $\sigma = 1.65 \times 10^{-2}$ to $G^h(x_{\text{true}})$. The corresponding $\ln \kappa_d(s; x_{\text{true}})$ and the simulated potential function at several time snapshots are shown in Fig. 11. The standard deviation $\sigma = 1.65 \times 10^{-2}$ corresponds to a signal-to-noise ratio of 10. This way, we have the unnormalized posterior density

$$\pi(x) = \exp\left(-\frac{1}{2\sigma^2} \|G^h(x) - y\|_F^2\right) \prod_{k=1}^d \left(\mathbb{I}_{[-5,5]}(x_k) \exp\left(-\frac{1}{2} x_k^2\right)\right).$$

6.4.2 Numerical Results

To construct DIRT, we employ a geometric grading in β , refining towards 1,

$$\log_{10} \beta_k \in \{-5, -4, -3, -2.5, -2, -1.5, -1, -0.75, -0.5, -0.25, 0\}.$$

The posterior is very concentrated in this example, so we employ separate tempering of prior and likelihood in the bridging densities,

$$\pi_k(x) = \exp\left(-\frac{\beta_k}{2\sigma^2} \|G^h(x) - y\|_F^2\right) \prod_{k=1}^d \left(\mathbb{I}_{[-5,5]}(x_k) \exp\left(-\frac{\beta_k^{0.01}}{2} x_k^2\right)\right).$$

in which a weakly tempered prior is used. We use a truncated normal reference measure on the domain $(-4, 4]^d$ with the Fourier basis to build DIRT. In TT-cross, a maximum iteration $\text{MaxIt} = 1$ without enrichment ($\text{Rho} = 0$) is used. The number of collocation points in each dimension is set to be $n = 16$, and the TT ranks are chosen to be $\text{R0} = \text{R}_{\max} = \text{R}_k$, where

$$\text{R}_k \in \{15, 15, 15, 15, 15, 15, 13, 9, 9, 8, 7\}$$

at the k th layer of DIRT.

The PDE in (64) is computationally expensive to solve. Here, our goal is to explore the posterior density defined by a forward model, G^{h_f} , with refined grid size $h_f = 1/80$. A coarse forward model, G^{h_c} with $h_c = 1/20$, and an intermediate forward model, G^{h_m} with $h_m = 1/40$, are used in defining the bridge densities to speed-up the DIRT construction. This multilevel construction shares similarities with the multi-fidelity preconditioning strategy of [51], except that DIRT is based on TT rather than optimisation and our multilevel models are blended into the bridging densities. In numerical experiments, we consider the CPU time of solving the coarse model evaluation as one work unit. The CPU times for evaluating the intermediate model and the fine model are about 12.5 work units and 160 work units, respectively.

In the first experiment, we employ the coarse forward model, G^{h_c} , to compare the sampling performance of DIRT with that of DRAM. The results are reported in Fig. 12a, where the number of independent samples is calculated as the length of the Markov chain divided by the estimated IACT. The estimated IACTs for DRAM and DIRT are about 132 and 3.04, respectively, and the importance sampling with DIRT produces $\text{ESS} = N/1.5$. For DRAM, we exclude the burn-in samples in the number of work units, whereas the number of work units for the DIRT includes the construction cost of DIRT (993392 density evaluations). In this experiment, despite the high construction cost, DIRT can generate a Markov chain with almost independent samples, which is significantly more efficient than DRAM. Furthermore, the construction cost of DIRT will be less significant if one needs to generate more posterior samples, as shown in Fig. 12a.

In the second experiment, we demonstrate the construction of DIRT using not only the bridge densities with different temperatures, but also the forward models with

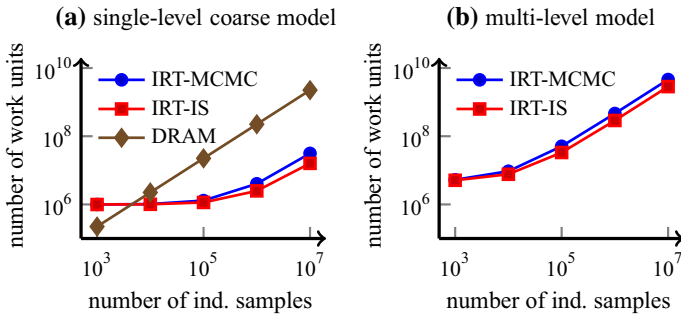


Fig. 12 Number of independent samples computed by IRT-MCMC, IRT-IS, and DRAM versus the total computational cost. **a** Comparison using a single level coarse model. **b** Comparison using the multilevel model in the DIRT construction. In both plots, the construction costs of IRT-MCMC and IRT-IS are included, whereas the burn-in cost of DRAM is not included. The work unit is the computational cost of one coarse model evaluation

different grid resolutions. For initial temperatures such that $\beta_k < 10^{-0.5}$, we use the coarse forward model G^{h_c} . For $\beta_k = 10^{-0.5}$ and $\beta_k = 10^{-0.25}$, we use the intermediate forward model G^{h_m} . For $\beta_k = 1$, we use the fine forward model G^{h_f} , so that the fine model is used to define the target posterior density. We need 915024 58544, and 19824 evaluations of the coarse, intermediate, and fine models, respectively, to construct DIRT. Once the DIRT is constructed, Algorithm 1 generates a Markov chain with IACT 2.87 that samples the posterior defined by the fine model. Again, the importance sampling is more efficient with $ESS = N/1.78$. The number of independent samples versus the number of work units is reported in Fig. 12b. In this experiment, it is computationally infeasible to apply DRAM directly (or any MCMC in general) to sample the posterior defined by the fine model. In contrast, the evaluation of DIRT and the corresponding posterior densities can be embarrassingly parallelised, which can further accelerate the posterior inference using high-performance computers. The IRT-IS algorithm can bypass the construction of Markov chains, which makes it suitable to be integrated into multilevel Monte Carlo or multilevel quasi-Monte Carlo estimators to improve the convergence rate of the computation of posterior expectations. We leave this as a future research question.

7 Conclusion

We have enabled functional tensor decompositions of complicated and concentrated continuous probability density functions that suffer from impractically large tensor ranks when approximated directly. Instead, we build an adaptive sequential change of coordinates that drives the target density towards a product function. This change of variables is realised by the composition of order-preserving SIRTs computed from functional TT decompositions of ratios of bridging densities. Each of the ratio functions recovers one scale of correlations of the target density, and hence, it can be approximated with fixed TT ranks. Together with the triangular structure of the Rosenblatt transport, this makes the total complexity linear in the number of variables.

This deep composition of the inverse Rosenblatt transports shares similarities with deep neural networks with nonlinear activation functions. However, DIRT has several advantages.

- Each DIRT layer, defined by the bridging densities, can be associated with the scale of noise or observation function. Any prior knowledge of model hierarchies can improve the selection of bridging densities. In contrast, the influence of a particular fully connected layer in a neural network is difficult to predict or understand.
- DIRT layers can be computed independently. As soon as the layer is approximated up to the desired accuracy, it can be saved and never recomputed again. This enables a simple interactive construction, where the tuning parameters can be set layer per layer. Neural networks require optimisation of all layers simultaneously.
- The construction of each DIRT layer is powered by efficient TT-cross algorithms, which can converge much faster than the stochastic gradient descent used by neural networks in many cases. The dense linear algebra operations used by TT decompositions can take full advantage of modern CPU and GPU vectorisations, whereas an embarrassing parallelism with respect to target density evaluations is well scalable to modern high-performance computers.

This work opens many potential applications and further enhancements of DIRT. For example, the transport maps defined by DIRT can be naturally extended to approximate the optimal biasing density in importance sampling, which can be valuable for solving rare event simulations. In Sect. 6.4, we offered some preliminary investigation on constructing DIRT using multilevel models. The multilevel idea can be further integrated with DIRT to improve the convergence rate of the importance sampling estimator. For problems involving extremely high-dimensional or infinite-dimensional random variables, DIRT can be combined with the likelihood-informed subspace (LIS) [13,14,59,66] to characterise the highly non-Gaussian effective random variable dimensions identified by LIS. In addition, for sequential Bayesian inference, we can apply DIRT to iteratively characterise the filtered posterior measures changing over time, where the evolution of the random states and time-dependent observations naturally defines a sequence of bridging measures.

Acknowledgements The author would like to thank Y. Marzouk and R. Scheichl for many insightful discussions. T. Cui acknowledges support from the Australian Research Council, under grant number CE140100049. S. Dolgov acknowledges support from the International Visitor Program of Sydney Mathematical Research Institute, and from the EPSRC New Investigator Award EP/T031255/1.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendices

Appendix A: Construction of Functional TT by Cross Interpolation

Here, we recall an alternating iteration algorithm for constructing the TT decomposition of a multivariate function $h : \mathcal{X} \mapsto \mathbb{R}$. We seek an TT of the form

$$h(x) \approx \tilde{h}(x) = \sum_{\alpha_0=1}^{r_0} \sum_{\alpha_1=1}^{r_1} \cdots \sum_{\alpha_d=1}^{r_d} H_1^{(\alpha_0, \alpha_1)}(x_1) \cdots H_k^{(\alpha_{k-1}, \alpha_k)}(x_k) \cdots H_d^{(\alpha_{d-1}, \alpha_d)}(x_d), \tag{66}$$

with $r_0 = r_d = 1$. Each univariate function $H_k^{(\alpha_{k-1}, \alpha_k)}(x_k) : \mathcal{X}_k \mapsto \mathbb{R}$ is represented as a linear combination of a set of n_k basis functions $\{\phi_k^{(i)}(x_k)\}_{i=1}^{n_k}$, that is,

$$H_k^{(\alpha_{k-1}, \alpha_k)}(x_k) = \sum_{i=1}^{n_k} \phi_k^{(i)}(x_k) \mathbf{A}_k[\alpha_{k-1}, i, \alpha_k], \tag{67}$$

where $\mathbf{A}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ is (the coefficient tensor of) the k th TT core. The number of degrees of freedom in the TT decomposition, that is, in the tensors $\{\mathbf{A}_k\}_{k=1}^d$, is linear in d provided the TT ranks r_0, \dots, r_d are bounded. For the numerical efficiency, it is essential that the TT cores can be constructed using a similar number of evaluations of $h(x)$. This can be achieved using cross interpolation methods. The following definition is used to construct cross interpolations.

Definition 1 For each variable x_k , we consider a set of *interpolation basis functions* that can be represented by a vector-valued function

$$\Phi_k(x_k) = [\phi_k^{(1)}(x_k), \dots, \phi_k^{(n_k)}(x_k)] \in \mathbb{R}^{1 \times n_k},$$

and a set of collocation points $X_k = \{x_k^{(i)}\}_{i=1}^{n_k}$ such that the $n_k \times n_k$ dimensional Vandermonde matrix

$$\Phi_k(X_k)[i, j] \equiv \phi_k^{(j)}(x_k^{(i)}), \tag{68}$$

is an identity matrix. A typical construction is the (piecewise) Lagrange basis functions defined by a point set X_k . We can also construct the interpolation basis from other basis functions of a separable Hilbert space, denoted by

$$\Psi_k(x_k) = [\psi_k^{(1)}(x_k), \dots, \psi_k^{(n_k)}(x_k)] \in \mathbb{R}^{1 \times n_k},$$

and a point set X_k with a nonsingular Vandermonde matrix by setting

$$\Phi_k(x_k) = \Psi_k(x_k) \Psi_k(X_k)^{-1}.$$

Specifically, if $\Psi_k(x_k)$ is a set of λ_k -orthogonal functions and X_k are the roots of the function $\psi_k^{(n_k+1)}(x_k)$, we recover the pseudo-spectral methods and have

$$\Psi_k(X_k)^{-1} = \Psi_k(X_k)^\top \text{diag}(\omega_k),$$

where the vector $\omega_k \in \mathbb{R}^{n_k}$ contains quadrature weights associated with X_k and $\text{diag}(\cdot)$ brings a vector into a diagonal matrix.

Furthermore, we define the mass matrix $M_k = \int \Phi_k(x_k)^\top \Phi_k(x_k) \lambda_k(x_k) dx_k$. We let $L_k \in \mathbb{R}^{n_k \times n_k}$ be the Cholesky factor of the mass matrix, i.e. $L_k L_k^\top = M_k$. For an interpolation basis constructed from λ_k -orthogonal functions and the roots of $\psi_k^{(n_k+1)}(x_k)$, we have $M_k = \text{diag}(\omega_k)$.

Two-Dimensional Case

Consider first the TT decomposition of a bivariate function

$$h(x_1, x_2) \approx \tilde{h}(x_1, x_2) = \sum_{\alpha=1}^r H_1^{(\alpha)}(x_1) H_2^{(\alpha)}(x_2), \tag{69}$$

where the rank- r cores

$$H_1(x_1) \equiv \Phi_1(x_1) A_1 \in \mathbb{R}^{1 \times r} \quad \text{and} \quad H_2(x_2) \equiv A_2 \Phi_2(x_2)^\top \in \mathbb{R}^{r \times 1}$$

are specified by basis functions $\Phi_1(x_1) \in \mathbb{R}^{1 \times n_1}$ and $\Phi_2(x_2) \in \mathbb{R}^{1 \times n_2}$, and the corresponding coefficient matrices $A_1 \in \mathbb{R}^{n_1 \times r}$ and $A_2 \in \mathbb{R}^{r \times n_2}$, respectively. We aim to recover A_1 and A_2 such that the L^2 norm of the error

$$\|h(x_1, x_2) - H_1(x_1) H_2(x_2)\|_{L^2_\lambda(\mathcal{X})},$$

is minimised. Note that with interpolation bases, the matrices A_1 and A_2 are also pointwise evaluations of the functions $H_1(x_1)$ and $H_2(x_2)$ at collocation points X_1 and X_2 , respectively. This way, $h(x_1, x_2)$ yields a discrete approximation

$$h(x_1, x_2) \approx \Phi_1(x_1) h(X_1, X_2) \Phi_2(x_2)^\top, \tag{70}$$

where $h(X_1, X_2) \equiv [h(x_1^{(i)}, x_2^{(j)})] \in \mathbb{R}^{n_1 \times n_2}$ for $x_1^{(i)} \in X_1$ and $x_2^{(j)} \in X_2$ is the matrix of nodal values of $h(x_1, x_2)$ similarly to (68). This way, the L^2 error of the continuous factorisation yields a discrete approximation

$$\begin{aligned} \|h(x_1, x_2) - H_1(x_1) H_2(x_2)\|_{L^2_\lambda(\mathcal{X})} &\approx \|\Phi_1(x_1) (h(X_1, X_2)^\top - A_1 A_2) \Phi_2(x_2)^\top\|_{L^2_\lambda(\mathcal{X})} \\ &= \|L_1^\top (h(X_1, X_2) - A_1 A_2) L_2\|_F. \end{aligned} \tag{71}$$

Thus, we can recover the matrices A_1 and A_2 by solving some low-rank matrix decomposition of $h(X_1, X_2)$. However, assembling the matrix $h(X_1, X_2)$ requires evaluating

the function $h(x_1, x_2)$ at the Cartesian union of the collocation points $X_1 \times X_2$, which can be computationally prohibitive in the generalisation to $d > 2$.

Instead, we can use some *interpolation point sets* $\bar{X}_1 \subset X_1$ and $\bar{X}_2 \subset X_2$ of cardinality $\#\bar{X}_1 = \#\bar{X}_2 = r$ such that the matrix $h(\bar{X}_1, \bar{X}_2) \in \mathbb{R}^{r \times r}$ is nonsingular, and rank- r *interpolation cores*

$$G_1(x_1) \equiv \Phi_1(x_1)B_1 \in \mathbb{R}^{1 \times r} \quad \text{and} \quad G_2(x_2) \equiv B_2\Phi_2(x_2)^\top \in \mathbb{R}^{r \times 1},$$

with $B_1 \in \mathbb{R}^{n_1 \times r}$ and $B_2 \in \mathbb{R}^{r \times n_2}$, to approximate $h(x_1, x_2)$ by interpolation. The interpolation cores satisfy the property that $G_1(\bar{X}_1)$ and $G_2(\bar{X}_2)$ are identity matrices. This yields interpolated approximations to $h(x_1, x_2)$, for example,

$$h(x_1, x_2) \approx G_1(x_1)h(\bar{X}_1, x_2) \quad \text{and} \quad h(x_1, x_2) \approx G_1(x_1)h(\bar{X}_1, \bar{X}_2)G_2(x_2).$$

This way, the goal becomes identifying the optimal point sets (\bar{X}_1, \bar{X}_2) and the cores (G_1, G_2) that minimise the interpolated rank- r factorisation error

$$\|h(x_1, x_2) - G_1(x_1)h(\bar{X}_1, \bar{X}_2)G_2(x_2)\|_{L^2_\lambda(\mathcal{X})}. \tag{72}$$

In practice, an *alternating direction* strategy can be employed to solve the above non-linear minimisation problem via a sequence of subproblems at a lower computational cost compared to that of the full matrix decomposition induced by (71). For example, we start from some initial guess of B_2 and \bar{X}_2 to solve for B_1 and \bar{X}_1 via the minimisation problem

$$B_1, \bar{X}_1 = \arg \min_{B'_1, \bar{X}'_1} \|h(x_1, \bar{X}_2)B_2\Phi_2(x_2)^\top - \Phi_1(x_1)B'_1h(\bar{X}'_1, \bar{X}_2)B_2\Phi_2(x_2)^\top\|_{L^2_\lambda(\mathcal{X})}, \tag{73}$$

then we use the updated B_1 and \bar{X}_1 to renew B_2 and \bar{X}_2 via

$$B_2, \bar{X}_2 = \arg \min_{B'_2, \bar{X}'_2} \|\Phi_1(x_1)B_1h(\bar{X}_1, x_2) - \Phi_1(x_1)B'_1h(\bar{X}_1, \bar{X}'_2)B'_2\Phi_2(x_2)^\top\|_{L^2_\lambda(\mathcal{X})},$$

and repeat until convergence. Given the collocation points \bar{X}_1 and \bar{X}_2 , the coefficient matrices B_1 and B_2 satisfy a simple quadratic optimisation and can be computed from

$$B_1h(\bar{X}_1, \bar{X}_2) = h(X_1, \bar{X}_2) \quad \text{and} \quad h(\bar{X}_1, \bar{X}_2)B_2 = h(\bar{X}_1, X_2), \tag{74}$$

respectively. Solving (74) only requires $(n_1 + n_2 - r)r$ evaluations of $h(x_1, x_2)$.

In (74), one needs to find the interpolation point sets \bar{X}_1 and \bar{X}_2 so that the resulting interpolation operator is an optimal approximation to the projection operator that spans the same linear subspace. However, finding the optimal interpolation point sets is an NP-hard problem. In practice, accurate quasi-optimal solutions can be obtained by greedy algorithms such as the (discrete) empirical interpolation [1,8] or the maximum volume (*MaxVol*) [22–24] methods. Here, we outline the procedure of the *MaxVol*

algorithm [22] for solving (73), which can be equivalently expressed as the problem of searching for an index set $\mathcal{I} \subset \{1, \dots, n\}$ of cardinality r such that the norm of $\mathbf{B} = \mathbf{H}\mathbf{H}^{-1} \in \mathbb{R}^{n \times r}$ is minimized. Here, $\overline{\mathbf{H}} = \mathbf{H}[\mathcal{I}, :] \in \mathbb{R}^{r \times r}$ is the submatrix of a given matrix \mathbf{H} in the MATLAB notation. For example, one can have $\mathbf{H} = h(\mathbf{X}_1, \overline{\mathbf{X}}_2)$, and then, the interpolation point set $\overline{\mathbf{X}}_1$ is given by \mathcal{I} and the coefficient matrix is set by $\mathbf{B}_1 = \mathbf{B}$.

Algorithm 3 MaxVol

- 1: Choose an initial set \mathcal{I} and a stopping threshold $\delta > 0$.
 - 2: **while** $\max_{i,j} |\mathbf{B}[i, j]| > 1 + \delta$ **do**
 - 3: Let $i_*, j_* = \arg \max_{i,j} |\mathbf{B}[i, j]|$.
 - 4: Replace the index $\mathcal{I}[j_*]$ in the set by i_* .
 - 5: Recompute $\overline{\mathbf{H}} = \mathbf{H}[\mathcal{I}, :]$ and $\mathbf{B} = \mathbf{H}\mathbf{H}^{-1}$.
 - 6: **end while**
-

Given an initial index set, which can be chosen as the r dominant pivots from Gaussian elimination, *MaxVol* proceeds as Algorithm 3. Note that $\mathbf{B}[\mathcal{I}, :] \in \mathbb{R}^{r \times r}$ is an identity matrix by construction. *MaxVol* ensures that no other row is more “important” by searching for a *dominant* submatrix $\overline{\mathbf{H}}$ such that $|\mathbf{B}[i, j]| \leq 1 + \delta$, which is a proxy to the *maximum volume* submatrix $\overline{\mathbf{H}}_*$ such that $|\det(\overline{\mathbf{H}}_*)| = \max_{\mathcal{I}} |\det(\mathbf{H}[\mathcal{I}, :])|$. The update of \mathbf{B} can be computed efficiently via the Sherman–Morrison–Woodbury formula [22] with a total cost of $\mathcal{O}(nr^2)$ per iteration.

For the numerical stability, it is beneficial to compute the thin generalised QR factorization $\mathbf{H}\mathbf{R} = h(\mathbf{X}_1, \overline{\mathbf{X}}_2)$, where the matrix \mathbf{H} has M_1 -orthonormal columns. This way, the set of functions $\Phi_1(x_1)\mathbf{H}$ forms a λ_1 -orthogonal basis. The factorisation $\mathbf{H}\mathbf{R} = h(\mathbf{X}_1, \overline{\mathbf{X}}_2)$ can be obtained by the thin QR factorization $\mathbf{Q}\mathbf{R} = \mathbf{L}_1^T h(\mathbf{X}_1, \overline{\mathbf{X}}_2)$ and $\mathbf{H} = \mathbf{L}_1^{-T}\mathbf{Q}$. Then, one can apply *MaxVol* to \mathbf{H} , which is the evaluation of $\Phi_1(x_1)\mathbf{H}$ at \mathbf{X}_1 , to select the index set \mathcal{I} , and thus, the interpolation points $\overline{\mathbf{X}}_1 \subset \mathbf{X}_1$. We have $\mathbf{H}[\mathcal{I}, :] = h(\overline{\mathbf{X}}_1, \overline{\mathbf{X}}_2)\mathbf{R}^{-1}$, which yields $\mathbf{B} = \mathbf{H}\mathbf{H}[\mathcal{I}, :]\mathbf{R}^{-1} = h(\mathbf{X}_1, \overline{\mathbf{X}}_2)h(\overline{\mathbf{X}}_1, \overline{\mathbf{X}}_2)^{-1}$. Thus, we can set the coefficient matrix as $\mathbf{B}_1 = \mathbf{B}$ and define the interpolation core $\mathbf{G}_1(x_1) = \Phi_1(x_1)\mathbf{B}_1$ such that $\mathbf{G}_1(\mathbf{X}_1)$ is an identity matrix.

We can obtain $(\mathbf{B}_1, \overline{\mathbf{X}}_1)$ and $(\mathbf{B}_2, \overline{\mathbf{X}}_2)$ by applying *MaxVol* within alternating iterations. Then, we can set $\mathbf{A}_1 = \mathbf{B}_1$ and $\mathbf{A}_2 = h(\overline{\mathbf{X}}_1, \overline{\mathbf{X}}_2)\mathbf{B}_2$ to recover the factorisation in the form of (69).

Multi-dimensional Case

The TT-cross algorithm [46] recursively extends (74) to $d > 2$. In the first step, we assume that a reduced point set $\overline{\mathbf{X}}_{>1} = \{(x_2^{(\alpha_1)}, \dots, x_d^{(\alpha_1)})\}$ of r_1 points is given. We can for example draw it from some tractable reference measure. We compute an analogue of the first equation in (74)

$$\mathbf{A}_1 h(\overline{\mathbf{X}}_1, \overline{\mathbf{X}}_{>1}) = h(\mathbf{X}_1, \overline{\mathbf{X}}_{>1}) \in \mathbb{R}^{n_1 \times r_1},$$

where $h(\mathbb{X}_1, \bar{\mathbb{X}}_{>1}) = [h(x_1^{(\alpha_1)}, x_2^{(\alpha_1)}, \dots, x_d^{(\alpha_1)})]$ is a matrix filled with the function $h(x)$ evaluated at the “reduced” set of points $\mathbb{X}_1 \times \bar{\mathbb{X}}_{>1}$. Now, we apply *MaxVol* to compute reduced subsets $\mathbb{I}_1 \subset \{1, \dots, n_1\}$ and $\bar{\mathbb{X}}_{<2} = \mathbb{X}_1(\mathbb{I}_1) \subset \mathbb{X}_1$. Similarly to the matrix \mathbf{B}_1 in the two-dimensional case, we let the actual TT core be the “stabilized” matrix $\mathbf{A}_1 = \mathbf{H}_1 \mathbf{H}_1[\mathbb{I}_1, :]^{-1}$, where $\mathbf{H}_1 \mathbf{R}_1 = h(\mathbb{X}_1, \bar{\mathbb{X}}_{>1})$ is the generalised QR decomposition.

In the k th step, we assume reduced point sets $\bar{\mathbb{X}}_{<k} = \{x_1^{(\alpha_{k-1})}, \dots, x_{k-1}^{(\alpha_{k-1})}\}$ and $\bar{\mathbb{X}}_{>k} = \{x_{k+1}^{(\alpha_k)}, \dots, x_d^{(\alpha_k)}\}$ are given. We can compute a third-order tensor

$$\mathbf{H}_k := \left[h(x_1^{(\alpha_{k-1})}, \dots, x_{k-1}^{(\alpha_{k-1})}, x_k^{(i_k)}, x_{k+1}^{(\alpha_k)}, \dots, x_d^{(\alpha_k)}) \right] \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}, \tag{75}$$

which consists of evaluations of $h(x)$ at the Cartesian union of the sets $\bar{\mathbb{X}}_{<k} \times \mathbb{X}_k \times \bar{\mathbb{X}}_{>k}$. We let $\bar{\mathbb{X}}_{<1} = \bar{\mathbb{X}}_{>d} = \emptyset$ to enable the notation for all k . We can *unfold* \mathbf{H}_k into matrices of the form

$$\mathbf{H}_k^{(L)} \in \mathbb{R}^{(r_{k-1}n_k) \times r_k}, \text{ and } \mathbf{H}_k^{(R)} \in \mathbb{R}^{r_{k-1} \times (n_k r_k)}, \tag{76}$$

such that

$$\mathbf{H}_k[\alpha_{k-1}, i_k, \alpha_k] = \mathbf{H}_k^{(L)}[\alpha_{k-1} + (i_k - 1)r_{k-1}, \alpha_k] = \mathbf{H}_k^{(R)}[\alpha_{k-1}, i_k + (\alpha_k - 1)n_k].$$

The union of the indices α_{k-1} and i_k corresponds to the union of the point sets $\mathbb{X}_{\leq k} := \bar{\mathbb{X}}_{<k} \times \mathbb{X}_k$. Therefore, we can apply *MaxVol* to $\mathbf{H}_k^{(L)}$ (or a generalised QR factor thereof) to obtain a discrete set $\mathbb{I}_k \subset \{1, \dots, r_{k-1}n_k\}$, and take a subsample of $\mathbb{X}_{\leq k}$ for the next recursion step, $\bar{\mathbb{X}}_{<k+1} = \mathbb{X}_{\leq k}(\mathbb{I}_k)$. Similarly for the k th TT core, we define

$$\mathbf{B}_k^{(L)} = \mathbf{H}_k^{(L)} \mathbf{H}_k^{(L)}[\mathbb{I}_k, :]^{-1}, \tag{77}$$

$$\mathbf{A}_k[\alpha_{k-1}, i_k, \alpha_k] = \mathbf{B}_k^{(L)}[\alpha_{k-1} + (i_k - 1)r_{k-1}, \alpha_k]. \tag{78}$$

If the function $h(x)$ admits an exact TT decomposition, and the initial point sets were chosen such that all $\mathbf{H}_k^{(L)}$ are full-rank, the recursion defined above reconstructs the decomposition exactly. However, in practice the initial point sets can be a poor interpolation sets. In this case, we can *refine* them by carrying out several iterations. Having computed \mathbf{A}_d , we reverse the recursion and iterate backwards, computing discrete sets $\mathbb{J}_k \subset \{1, \dots, n_k r_k\}$ via *MaxVol* applied to $(\mathbf{H}_k^{(R)})^\top$, and setting $\bar{\mathbb{X}}_{>k-1} = \mathbb{X}_{\geq k}(\mathbb{J}_k)$, where $\mathbb{X}_{\geq k} = \mathbb{X}_k \times \bar{\mathbb{X}}_{>k}$.

The second key ingredient is the adaptation of TT ranks. The TT ranks can be easily reduced. For example, it is sufficient to replace the generalised QR factorization of $\mathbf{H}_k^{(L)}$ or $(\mathbf{H}_k^{(R)})^\top$ by a generalised SVD, where the singular values below the desired threshold are truncated. To increase the TT ranks, we can apply *oversampling*. Using the forward iteration (with k increasing) as an example, we can compute the tensor $\mathbf{H}_k \in \mathbb{R}^{r_{k-1} \times n_k \times (r_k + \rho_k)}$ on the enriched point set $\bar{\mathbb{X}}_{<k} \times \mathbb{X}_k \times (\bar{\mathbb{X}}_{>k} \cup \tilde{\mathbb{X}}_{>k})$, where $\tilde{\mathbb{X}}_{>k} = \{(x_{k+1}^{(\alpha_k)}, \dots, x_d^{(\alpha_k)})\}_{\alpha_k=1}^{\rho_k}$ are auxiliary points. These auxiliary points can be sampled at random [45], or more accurately, from a surrogate of the error [18]. In

the latter case, we carry out a second TT-cross to approximate the error $h(x) - \tilde{h}(x)$ by a TT decomposition with TT ranks $\rho_1, \dots, \rho_{d-1}$ and take the *MaxVol* points of the error as $\tilde{X}_{>k}$. This *enrichment* of the solution with error or residual information has proven to accelerate the convergence drastically even for small expansion ranks ρ_k when applied to solving linear systems [18]. The pseudocode of the TT-cross is provided in Algorithm 4.

Algorithm 4 TT-cross

- 1: Choose initial sets $\bar{X}_{<k}, k = 2, \dots, d$, stopping threshold $\delta > 0$, enrichment TT ranks ρ_k .
 - 2: **while** first iteration or $\|h(x) - \tilde{h}_{\text{prev}}(x)\| > \delta \|h(x)\|$ **do**
 - 3: **for** $k = d, d - 1, \dots, 2$ **do** ▷ backward iteration
 - 4: Sample \mathbf{H}_k as shown in (75), optionally expanding $\bar{X}_{<k}$ to $\bar{X}_{<k} \cup \tilde{X}_{<k}$.
 - 5: Compute \mathbb{J}_k from *MaxVol* on $(\mathbf{H}_k^{(R)})^\top$ or its SVD factor, let $\bar{X}_{>k-1} = X_{\geq k}(\mathbb{J}_k)$.
 - 6: **end for**
 - 7: **for** $k = 1, 2, \dots, d - 1$ **do** ▷ forward iteration
 - 8: Sample \mathbf{H}_k as shown in (75), optionally expanding $\bar{X}_{>k}$ to $\bar{X}_{>k} \cup \tilde{X}_{>k}$.
 - 9: Compute \mathbb{I}_k from *MaxVol* on $\mathbf{H}_k^{(L)}$ or its SVD factor, let $\bar{X}_{<k+1} = X_{\leq k}(\mathbb{I}_k)$.
 - 10: Reconstruct TT cores as shown in (77)–(78).
 - 11: **end for**
 - 12: Sample the last TT core $\mathbf{A}_d = \mathbf{H}_d$ as shown in (75).
 - 13: **end while**
-

The tensor in (75) suggests that the TT-cross requires $\sum_{k=1}^d r_{k-1} n_k r_k$ evaluations of $h(x)$ per iteration, which is proportional to the number of unknowns in the TT cores. To enhance the robustness (at the expense of a larger number of function evaluations), one may oversample X_k beyond n_k basis functions and use the *rectangular MaxVol* algorithm [42] to oversample $\mathbb{I}_k, \mathbb{J}_{k+1}$ beyond r_k indices. In this case, the matrix inverse in (77) is replaced by a pseudoinverse. For our DIRT framework, the standard *MaxVol* equipped with the error enrichment is sufficiently robust to factorise the ratio functions, so we proceed with Algorithm 4.

Appendix B: Proof of Proposition 2

Recall the marginal function

$$\hat{\pi}_{\leq k}(x_{\leq k}) = \gamma \prod_{i=k+1}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_k=1}^{r_k} \left(\mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-1}(x_{k-1}) \mathbf{P}_k^{(\cdot, \ell_k)}(x_k) \right)^2, \quad (79)$$

where $\mathbf{P}_k(x_k) : \mathcal{X}_k \mapsto \mathbb{R}^{r_{k-1} \times r_k}$ is given by (26) with a coefficient tensor $\mathbf{B}_k \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$. The next marginal function $\hat{\pi}_{<k}(x_{<k})$ is defined by

$$\hat{\pi}_{<k}(x_{<k}) = \int_{\mathcal{X}_k} \left(\gamma \prod_{i=k+1}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_k=1}^{r_k} \left(\mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-1}(x_{k-1}) \mathbf{P}_k^{(\cdot, \ell_k)}(x_k) \right)^2 \right) \lambda_k(x_k) dx_k$$

$$= \gamma \prod_{i=k}^d \lambda_i(\mathcal{X}_i) + \sum_{\ell_k=1}^{r_k} \int_{\mathcal{X}_k} \left(\mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-1}(x_{k-1}) \mathbf{P}_k^{(:, \ell_k)}(x_k) \right)^2 \lambda_k(x_k) dx_k. \tag{80}$$

The second term of (80) can be expressed as

$$\begin{aligned} & \sum_{\ell_k=1}^{r_k} \int_{\mathcal{X}_k} \left(\mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-1}(x_{k-1}) \mathbf{P}_k^{(:, \ell_k)}(x_k) \right)^2 \lambda_k(x_k) dx_k \\ &= \sum_{\ell_k=1}^{r_k} \int_{\mathcal{X}_k} \sum_{\alpha_{<k}=1}^{r_{k-1}} \mathbf{G}_{<k}^{(\alpha_{<k})}(x_{<k}) \mathbf{P}_k^{(\alpha_{<k}, \ell_k)}(x_k) \sum_{\beta_{k-1}=1}^{r_{k-1}} \mathbf{G}_{<k}^{(\beta_{k-1})}(x_{<k}) \mathbf{P}_k^{(\beta_{k-1}, \ell_k)}(x_k) \lambda_k(x_k) dx_k \\ &= \sum_{\alpha_{k-1}=1}^{r_{k-1}} \sum_{\beta_{k-1}=1}^{r_{k-1}} \mathbf{G}_{<k}^{(\alpha_{k-1})}(x_{<k}) \mathbf{G}_{<k}^{(\beta_{k-1})}(x_{<k}) \overline{\mathbf{M}}_k[\alpha_{k-1}, \beta_{k-1}], \end{aligned} \tag{81}$$

where the symmetric matrix $\overline{\mathbf{M}}_k \in \mathbb{R}^{r_{k-1} \times r_{k-1}}$ is given by

$$\overline{\mathbf{M}}_k[\alpha_{k-1}, \beta_{k-1}] = \sum_{\ell_k=1}^{r_k} \int_{\mathcal{X}_k} \mathbf{P}_k^{(\alpha_{k-1}, \ell_k)}(x_k) \mathbf{P}_k^{(\beta_{k-1}, \ell_k)}(x_k) \lambda_k(x_k) dx_k, \tag{82}$$

and $\mathbf{G}_{<k}$ is defined in (25). Plugging the expression (26) of $\mathbf{P}_k(x_k)$ into (82), we obtain

$$\begin{aligned} & \overline{\mathbf{M}}_k[\alpha_{k-1}, \beta_{k-1}] \\ &= \sum_{\ell_k=1}^{r_k} \int_{\mathcal{X}_k} \left(\sum_{i=1}^{n_k} \phi_k^{(i)}(x_k) \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] \right) \left(\sum_{j=1}^{n_k} \phi_k^{(j)}(x_k) \mathbf{B}_k[\beta_{k-1}, j, \ell_k] \right) \lambda_k(x_k) dx_k \\ &= \sum_{\ell_k=1}^{r_k} \sum_{i=1}^{n_k} \sum_{j=1}^{n_k} \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] \mathbf{B}_k[\beta_{k-1}, j, \ell_k] \int_{\mathcal{X}_k} \phi_k^{(i)}(x_k) \phi_k^{(j)}(x_k) \lambda_k(x_k) dx_k \\ &= \sum_{\ell_k=1}^{r_k} \sum_{i=1}^{n_k} \sum_{j=1}^{n_k} \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] \mathbf{B}_k[\beta_{k-1}, j, \ell_k] \mathbf{M}_k[i, j], \end{aligned} \tag{83}$$

where $\mathbf{M}_k \in \mathbb{R}^{n_k \times n_k}$ is the symmetric positive definite mass matrix defined in (22). Given the Cholesky decomposition $\mathbf{L}_k \mathbf{L}_k^\top = \mathbf{M}_k$, we have

$$\mathbf{M}_k[i, j] = \sum_{\tau=1}^{n_k} \mathbf{L}_k[i, \tau] \mathbf{L}_k[j, \tau].$$

Substituting the above identity into (83), we have

$$\overline{\mathbf{M}}_k[\alpha_{k-1}, \beta_{k-1}]$$

$$\begin{aligned}
 &= \sum_{\ell_k=1}^{r_k} \sum_{i=1}^{n_k} \sum_{j=1}^{n_k} \sum_{\tau=1}^{n_k} \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] \mathbf{B}_k[\beta_{k-1}, j, \ell_k] \mathbf{L}_k[i, \tau] \mathbf{L}_k[j, \tau] \\
 &= \sum_{\ell_k=1}^{r_k} \sum_{\tau=1}^{n_k} \left(\sum_{i=1}^{n_k} \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] \mathbf{L}_k[i, \tau] \right) \left(\sum_{j=1}^{n_k} \mathbf{B}_k[\beta_{k-1}, j, \ell_k] \mathbf{L}_k[j, \tau] \right). \tag{84}
 \end{aligned}$$

Denoting $\mathbf{C}_k[\alpha_{k-1}, \tau, \ell_k] = \sum_{i=1}^{n_k} \mathbf{B}_k[\alpha_{k-1}, i, \ell_k] \mathbf{L}_k[i, \tau]$ and unfolding \mathbf{C}_k along the first coordinate similarly to (76) to obtain a matrix $\mathbf{C}_k^{(R)} \in \mathbb{R}^{r_{k-1} \times (n_k r_k)}$, we have

$$\begin{aligned}
 &\overline{\mathbf{M}}_k[\alpha_{k-1}, \beta_{k-1}] \\
 &= \sum_{\ell_k=1}^{r_k} \sum_{\tau=1}^{n_k} \mathbf{C}_k[\alpha_{k-1}, \tau, \ell_k] \mathbf{C}_k[\beta_{k-1}, \tau, \ell_k] = \sum_{\iota=1}^{n_k r_k} \mathbf{C}_k^{(R)}[\alpha_{k-1}, \iota] \mathbf{C}_k^{(R)}[\beta_{k-1}, \iota].
 \end{aligned}$$

Equivalently, we have $\overline{\mathbf{M}}_k = \mathbf{C}_k^{(R)} (\mathbf{C}_k^{(R)})^\top$. This way, computing the thin QR decomposition

$$\mathbf{Q}_k \mathbf{R}_k = (\mathbf{C}_k^{(R)})^\top,$$

we obtain the Cholesky decomposition $\mathbf{R}_k^\top \mathbf{R}_k = \overline{\mathbf{M}}_k$ where $\mathbf{R}_k \in \mathbb{R}^{r_{k-1} \times r_{k-1}}$ is upper-triangular. Substituting the identity

$$\overline{\mathbf{M}}_k[\alpha_{k-1}, \beta_{k-1}] = \sum_{\ell_{k-1}=1}^{r_{k-1}} \mathbf{R}_k[\ell_{k-1}, \alpha_{k-1}] \mathbf{R}_k[\ell_{k-1}, \beta_{k-1}],$$

into (81), the second term of (80) is defined by

$$\begin{aligned}
 &\sum_{\ell_k=1}^{r_k} \int_{\mathcal{X}_k} \left(\mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-1}(x_{k-1}) \mathbf{P}_k^{(:, \ell_k)}(x_k) \right)^2 \lambda_k(x_k) dx_k \\
 &= \sum_{\ell_{k-1}=1}^{r_{k-1}} \sum_{\alpha_{k-1}=1}^{r_{k-1}} \sum_{\beta_{k-1}=1}^{r_{k-1}} \mathbf{G}_{<k}^{(\alpha_{k-1})}(x_{<k}) \mathbf{G}_{<k}^{(\beta_{k-1})}(x_{<k}) \mathbf{R}_k[\ell_{k-1}, \alpha_{k-1}] \mathbf{R}_k[\ell_{k-1}, \beta_{k-1}] \\
 &= \sum_{\ell_{k-1}=1}^{r_{k-1}} \left(\sum_{\alpha_{k-1}=1}^{r_{k-1}} \mathbf{G}_{<k}^{(\alpha_{k-1})}(x_{<k}) \mathbf{R}_k[\ell_{k-1}, \alpha_{k-1}] \right)^2 \\
 &= \sum_{\ell_{k-1}=1}^{r_{k-1}} \left(\mathbf{G}_1(x_1) \cdots \mathbf{G}_{k-2}(x_{k-2}) \sum_{\alpha_{k-1}=1}^{r_{k-1}} \mathbf{G}_{k-1}^{(:, \alpha_{k-1})}(x_{k-1}) \mathbf{R}_k[\ell_{k-1}, \alpha_{k-1}] \right)^2, \tag{85}
 \end{aligned}$$

where the last line follows from the identity in (25). Following the recursive definition of $P_k(x_k)$ in Proposition 2, we have

$$P_{k-1}^{(\cdot, \ell_{k-1})}(x_{k-1}) = \sum_{\alpha_{k-1}=1}^{r_{k-1}} G_{k-1}^{(\cdot, \alpha_{k-1})}(x_{k-1}) R_k[\ell_{k-1}, \alpha_{k-1}]. \tag{86}$$

Substituting the definition of the TT core

$$G_{k-1}^{(\alpha_{k-1}-1, \alpha_{k-1})}(x_{k-1}) = \sum_{i=1}^{n_{k-1}} \phi_{k-1}^{(i)}(x_{k-1}) \mathbf{A}_{k-1}[\alpha_{k-2}, i, \alpha_{k-1}]$$

into (86), we have

$$P_{k-1}^{(\alpha_{k-2}, \ell_{k-1})}(x_{k-1}) = \sum_{i=1}^{n_{k-1}} \phi_{k-1}^{(i)}(x_{k-1}) \left(\sum_{\alpha_{k-1}=1}^{r_{k-1}} \mathbf{A}_{k-1}[\alpha_{k-2}, i, \alpha_{k-1}] R_k[\ell_{k-1}, \alpha_{k-1}] \right).$$

Thus, defining the coefficient tensor

$$\mathbf{B}_{k-1}[\alpha_{k-2}, i, \ell_{k-1}] = \sum_{\alpha_{k-1}=1}^{r_{k-1}} \mathbf{A}_{k-1}[\alpha_{k-2}, i, \alpha_{k-1}] R_k[\ell_{k-1}, \alpha_{k-1}],$$

we obtain the result of Proposition 2:

$$P_{k-1}^{(\alpha_{k-2}, \ell_{k-1})}(x_{k-1}) = \sum_{i=1}^{n_{k-1}} \phi_{k-1}^{(i)}(x_{k-1}) \mathbf{B}_{k-1}[\alpha_{k-2}, i, \ell_{k-1}].$$

By setting index $k = 1$ and repeating the above procedure, we can obtain the normalising constant $\hat{z} = \int_{\mathcal{X}_1} \hat{\pi}_{\leq 1}(x_1) \lambda_1(x_1) dx_1 = \gamma \prod_{i=1}^d \lambda_i(\mathcal{X}_i) + R_1^2$, where $R_1 \in \mathbb{R}$ as the unfolded \mathbf{C}_1 along the first coordinate is a row vector $\mathbf{C}_1^{(R)} \in \mathbb{R}^{1 \times (n_1 r_1)}$.

References

1. Maxime Barrault, Yvon Maday, Ngoc Cuong Nguyen, and Anthony T Patera, *An ‘empirical interpolation’ethod: application to efficient reduced-basis discretization of partial differential equations*, *Comptes Rendus Mathematique* **339** (2004), no. 9, 667–672.
2. Daniele Bigoni, Allan P Engsig-Karup, and Youssef M Marzouk, *Spectral tensor-train decomposition*, *SIAM Journal on Scientific Computing* **38** (2016), no. 4, A2405–A2439.
3. John P Boyd, *Chebyshev and fourier spectral methods*, Courier Corporation, 2001.
4. Michael C Brennan, Daniele Bigoni, Olivier Zahm, Alessio Spantini, and Youssef M Marzouk, *Greedy inference with structure-exploiting lazy maps*, *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 8330–8342.
5. Hans-Joachim Bungartz and Michael Griebel, *Sparse grids*, *Acta Numerica* **13** (2004), no. 1, 147–269.
6. Guillaume Carlier, Alfred Galichon, and Filippo Santambrogio, *From Knothe’s transport to Brenier’s map and a continuation method for optimal transport*, *SIAM Journal on Mathematical Analysis* **41** (2010), no. 6, 2554–2576.

7. Anthony Caterini, Rob Cornish, Dino Sejdinovic, and Arnaud Doucet, *Variational inference with continuously-indexed normalizing flows*, In: 37th Conference on Uncertainty in Artificial Intelligence (UAI 2021). <https://auai.org/uai2021/pdf/uai2021.30.pdf> arXiv preprint [arXiv:2007.05426](https://arxiv.org/abs/2007.05426) (2020).
8. Saifon Chaturantabud and Danny C Sorensen, *Nonlinear model reduction via discrete empirical interpolation*, *SIAM Journal on Scientific Computing* **32** (2010), no. 5, 2737–2764.
9. Ricky T Q Chen, Jens Behrmann, David K Duvenaud, and Joern-Henrik Jacobsen, *Residual flows for invertible generative modeling*, *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 9916–9926.
10. Rob Cornish, Anthony Caterini, George Deligiannidis, and Arnaud Doucet, *Relaxing bijectivity constraints with continuously indexed normalising flows*, *Proceedings of the 37th International Conference on Machine Learning*, PMLR, vol. 119, 2020, pp. 2133–2143.
11. Tiangang Cui, Youssef M Marzouk, and Karen E Willcox, *Data-driven model reduction for the Bayesian solution of inverse problems*, *International Journal for Numerical Methods in Engineering* **102** (2015), 966–990.
12. Tiangang Cui, Youssef M Marzouk, and Karen E Willcox, *Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction*, *Journal of Computational Physics* **315** (2016), 363–387.
13. Tiangang Cui, James Martin, Youssef M Marzouk, Antti Solonen, and Alessio Spantini, *Likelihood-informed dimension reduction for nonlinear inverse problems*, *Inverse Problems* **30** (2014), no. 11, 114015.
14. Tiangang Cui and Xin T Tong, *A unified performance analysis of likelihood-informed subspace methods*, arXiv preprint [arXiv:2101.02417](https://arxiv.org/abs/2101.02417) (2021).
15. Gianluca Detommaso, Tiangang Cui, Youssef M Marzouk, Alessio Spantini, and Robert Scheichl, *A Stein variational Newton method*, *Advances in Neural Information Processing Systems*, vol. 31, 2018, pp. 9169–9179.
16. Josef Dick, Frances Y Kuo, and Ian H Sloan, *High-dimensional integration: The quasi-Monte Carlo way*, *Acta Numerica* **22** (2013), 133–288.
17. Sergey Dolgov, Karim Anaya-Izquierdo, Colin Fox, and Robert Scheichl, *Approximation and sampling of multivariate probability distributions in the tensor train decomposition*, *Statistics and Computing* **30** (2020), 603–625.
18. Sergey Dolgov and Dmitry V Savostyanov, *Alternating minimal energy methods for linear systems in higher dimensions*, *SIAM Journal on Scientific Computing* **36** (2014), no. 5, A2248–A2271.
19. Wolfgang Förstner and Boudewijn Moonen, *A metric for covariance matrices*, *Geodesy—the Challenge of the 3rd Millennium*, Springer, 2003, pp. 299–309.
20. Andrew Gelman and Xiao-Li Meng, *Simulating normalizing constants: From importance sampling to bridge sampling to path sampling*, *Statistical Science* (1998), 163–185.
21. Mark Girolami and Ben Calderhead, *Riemann manifold Langevin and Hamiltonian Monte Carlo methods*, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73** (2011), no. 2, 123–214.
22. Sergei A Goreinov, Ivan V Oseledets, Dimitry V Savostyanov, Eugene E Tyrtshnikov, and Nikolay L Zamarashkin, *How to find a good submatrix*, *Matrix Methods: Theory, Algorithms And Applications: Dedicated to the Memory of Gene Golub*, World Scientific, 2010, pp. 247–256.
23. Sergei A Goreinov, Eugene E Tyrtshnikov, and Nikolai L Zamarashkin, *A theory of pseudoskeleton approximations*, *Linear Algebra and its Applications* **261** (1997), no. 1–3, 1–21.
24. Sergei A Goreinov, Nikolai L Zamarashkin, and Eugene E Tyrtshnikov, *Pseudo-skeleton approximations by matrices of maximal volume*, *Mathematical Notes* **62** (1997), no. 4, 515–519.
25. Alex Gorodetsky, Sertac Karaman, and Youssef M Marzouk, *A continuous analogue of the tensor-train decomposition*, *Computer Methods in Applied Mechanics and Engineering* **347** (2019), 59–84.
26. Michael Griebel and Helmut Harbrecht, *Analysis of tensor approximation schemes for continuous functions*, arXiv preprint [arXiv:1903.04234](https://arxiv.org/abs/1903.04234) (2019).
27. Heikki Haario, Marko Laine, Antonietta Mira, and Eero Saksman, *DRAM: efficient adaptive MCMC*, *Statistics and Computing* **16** (2006), no. 4, 339–354.
28. Wolfgang Hackbusch, *Tensor spaces and numerical tensor calculus*, vol. 42, Springer Science & Business Media, 2012.
29. Dave Higdon, *Space and space-time modeling using process convolutions*, *Quantitative methods for current environmental issues*, Springer, 2002, pp. 37–56.

30. 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** (2012), no. 2, A683–A713.
31. Koji Hukushima and Koji Nemoto, *Exchange Monte Carlo method and application to spin glass simulations*, Journal of the Physical Society of Japan **65** (1996), no. 6, 1604–1608.
32. Nikolas Kantas, Alexandros Beskos, and Ajay Jasra, *Sequential Monte Carlo methods for high-dimensional inverse problems: A case study for the Navier-Stokes equations*, SIAM/ASA Journal on Uncertainty Quantification **2** (2014), no. 1, 464–489.
33. Herbert Knothe, *Contributions to the theory of convex bodies.*, The Michigan Mathematical Journal **4** (1957), no. 1, 39–52.
34. Tamara G Kolda and Brett W Bader, *Tensor decompositions and applications*, SIAM Review **51** (2009), no. 3, 455–500.
35. Jakob Kruse, Gianluca Detommaso, Robert Scheichl, and Ullrich Koethe, *HINT: Hierarchical invertible neural transport for density estimation and Bayesian inference*, The AAAI Conference on Artificial Intelligence (AAAI), 2021, p. to appear.
36. Jun S Liu, *Monte Carlo strategies in scientific computing*, Springer, New York, 2001.
37. Qiang Liu and Dilin Wang, *Stein variational gradient descent: A general purpose Bayesian inference algorithm*, Advances In Neural Information Processing Systems, vol. 29, 2016, pp. 2378–2386.
38. Michael W Mahoney and Petros Drineas, *CUR matrix decompositions for improved data analysis*, Proceedings of the National Academy of Sciences **106** (2009), no. 3, 697–702.
39. Youssef M Marzouk, Tarek Moselhy, Matthew Parno, and Alessio Spantini, *Sampling via measure transport: An introduction*, Handbook of Uncertainty Quantification (2016), 1–41.
40. Xiao-Li Meng and Wing Hung Wong, *Simulating ratios of normalizing constants via a simple identity: a theoretical exploration*, Statistica Sinica (1996), 831–860.
41. Kerrie L Mengersen, Richard L Tweedie, et al., *Rates of convergence of the hastings and metropolis algorithms*, The Annals of Statistics **24** (1996), no. 1, 101–121.
42. Aleksandr Y Mikhalev and Ivan V Oseledets, *Rectangular maximum-volume submatrices and their applications*, Linear Algebra and its Applications **538** (2018), 187–211.
43. Tarek Moselhy and Youssef M Marzouk, *Bayesian inference with optimal maps*, Journal of Computational Physics **231** (2012), no. 23, 7815–7850.
44. Radford M Neal, *Sampling from multimodal distributions using tempered transitions*, Statistics and computing **6** (1996), no. 4, 353–366.
45. Ivan V Oseledets, *DMRG approach to fast linear algebra in the TT-format*, Comput. Meth. Appl. Math. **11** (2011), no. 3, 382–393.
46. Ivan V Oseledets and Eugene E Tyrtysnikov, *TT-cross approximation for multidimensional arrays*, Linear Algebra and its Applications **432** (2010), no. 1, 70–88.
47. Ivan V Oseledets, *Tensor-train decomposition*, SIAM Journal on Scientific Computing **33** (2011), no. 5, 2295–2317.
48. Art B Owen, *Monte Carlo theory, methods and examples*, 2013.
49. George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan, *Normalizing flows for probabilistic modeling and inference*, JMLR arXiv preprint arXiv:1912.02762; <https://www.jmlr.org/papers/volume22/19-1028/19-1028.pdf> (2019).
50. Matthew D Parno and Youssef M Marzouk, *Transport map accelerated Markov chain Monte Carlo*, SIAM/ASA Journal on Uncertainty Quantification **6** (2018), no. 2, 645–682.
51. Benjamin Peherstorfer and Youssef M Marzouk, *A transport-based multifidelity preconditioner for Markov chain Monte Carlo*, Advances in Computational Mathematics **45** (2019), no. 5-6, 2321–2348.
52. Danilo Rezende and Shakir Mohamed, *Variational inference with normalizing flows*, International Conference on Machine Learning, vol. 37, 2015, pp. 1530–1538.
53. Christian Robert and George Casella, *Monte Carlo statistical methods*, Springer Science & Business Media, 2013.
54. 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.
55. Murray Rosenblatt, *Remarks on a multivariate transformation*, The Annals of Mathematical Statistics **23** (1952), no. 3, 470–472.
56. Robert Scheichl, Andrew M Stuart, and Aretha L Teckentrup, *Quasi-Monte Carlo and Multilevel Monte Carlo methods for computing posterior expectations in elliptic inverse problems*, SIAM/ASA Journal on Uncertainty Quantification **5** (2017), no. 1, 493–518.

57. Jie Shen, Tao Tang, and Li-Lian Wang, *Spectral methods: algorithms, analysis and applications*, vol. 41, Springer Science & Business Media, 2011.
58. Alessio Spantini, Daniele Bigoni, and Youssef M Marzouk, *Inference via low-dimensional couplings*, *The Journal of Machine Learning Research* **19** (2018), no. 1, 2639–2709.
59. Alessio Spantini, Antti Solonen, Tiangang Cui, James Martin, Luis Tenorio, and Youssef M Marzouk, *Optimal low-rank approximations of Bayesian linear inverse problems*, *SIAM Journal on Scientific Computing* **37** (2015), no. 6, A2451–A2487.
60. Robert H Swendsen and Jian-Sheng Wang, *Replica Monte Carlo simulation of spin-glasses*, *Physical Review Letters* **57** (1986), no. 21, 2607.
61. Esteban G Tabak and Cristina V Turner, *A family of nonparametric density estimation algorithms*, *Communications on Pure and Applied Mathematics* **66** (2013), no. 2, 145–164.
62. Lloyd N Trefethen, *Approximation theory and approximation practice*, vol. 164, SIAM, 2019.
63. Giulio Trigila and Esteban G Tabak, *Data-driven optimal transport*, *Communications on Pure and Applied Mathematics* **69** (2016), no. 4, 613–648.
64. Cédric Villani, *Optimal transport: old and new*, vol. 338, Springer Science & Business Media, 2008.
65. Steven R White, *Density-matrix algorithms for quantum renormalization groups*, *Physical Review B* **48** (1993), no. 14, 10345.
66. Olivier Zahm, Tiangang Cui, Kody Law, Alessio Spantini, and Youssef M Marzouk, *Certified dimension reduction in nonlinear Bayesian inverse problems*, arXiv preprint [arXiv:1807.03712](https://arxiv.org/abs/1807.03712) (2018).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.