

Tensor-train methods for sequential state and parameter learning in state-space models

Yiran Zhao

Tiangang Cui

School of Mathematics and Statistics

University of Sydney

New South Wales 2006, Australia

YIRAN.ZHAO@SYDNEY.EDU.AU

TIANGANG.CUI@SYDNEY.EDU.AU

Editor: Ryan Adams

Abstract

We consider sequential state and parameter learning in state-space models with intractable state transition and observation processes. By exploiting low-rank tensor train (TT) decompositions, we propose new sequential learning methods for joint parameter and state estimation under the Bayesian framework. Our key innovation is the introduction of scalable function approximation tools such as TT for recursively learning the sequentially updated posterior distributions. The function approximation perspective of our methods offers tractable error analysis and potentially alleviates the particle degeneracy faced by many particle-based methods. In addition to the new insights into the algorithmic design, our methods complement conventional particle-based methods. Our TT-based approximations naturally define conditional Knothe–Rosenblatt (KR) rearrangements that lead to parameter estimation, filtering, smoothing and path estimation accompanying our sequential learning algorithms, which open the door to removing potential approximation bias. We also explore several preconditioning techniques based on either linear or nonlinear KR rearrangements to enhance the approximation power of TT for practical problems. We demonstrate the efficacy and efficiency of our proposed methods on several state-space models, in which our methods achieve state-of-the-art estimation accuracy and computational performance.

Keywords: tensor train, Knothe–Rosenblatt rearrangement, state-space models, sequential Monte Carlo, uncertainty quantification, transport maps

1. Introduction

State-space models have been widely used in mathematical and statistical modelling to analyze time-varying complex phenomena (Cappé et al., 2006; Kantas et al., 2009). Examples include time series analysis in finance, temporal pattern recognition in bioinformatics, forecast in meteorology, and more. A typical state-space model is also referred to as a hidden Markov model, which consists of two stochastic processes $\{\mathbf{X}_t\}_{t \geq 0}$ and $\{\mathbf{Y}_t\}_{t \geq 1}$. The state transition process $\{\mathbf{X}_t\}_{t \geq 0}$ is an \mathcal{X} -valued latent Markov process specified by a conditional transition density

$$(\mathbf{X}_t | \mathbf{X}_{0:t-1} = \mathbf{x}_{0:t-1}) \equiv (\mathbf{X}_t | \mathbf{X}_{t-1} = \mathbf{x}_{t-1}) \sim f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}), \quad (1)$$

and an initial density $p(\mathbf{x}_0|\boldsymbol{\theta})$. Here \mathbf{X}_t represents the hidden state of the underlying system at an integer-valued time index $t \geq 0$, and the notation $\mathbf{x}_{i:j}$ denotes vectors $(\mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_j)$ from a sequence $\{\mathbf{x}_t\}_{t \geq 0}$. The \mathcal{Y} -valued observation process $\{\mathbf{Y}_t\}_{t \geq 1}$ is characterized by the likelihood function $g(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta})$ as

$$(\mathbf{Y}_t|\mathbf{X}_{0:t} = \mathbf{x}_{0:t}, \mathbf{Y}_{1:t-1} = \mathbf{y}_{1:t-1}) \equiv (\mathbf{Y}_t|\mathbf{X}_t = \mathbf{x}_t) \sim g(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}), \quad (2)$$

where \mathbf{Y}_t represents the observable at time index $t > 0$. In (1) and (2), $\boldsymbol{\theta} \in \Theta$ is a set of parameters that govern the state transition process and the observation process.

Example 1 (*Volatility of financial instruments*) *The stochastic volatility model consists of the return of an asset $\{\mathbf{Y}_t\}_{t \geq 1}$ and the logarithm of its squared volatility $\{\mathbf{X}_t\}_{t \geq 0}$. In the simplest setup, the logarithm of the squared volatility $\{\mathbf{X}_t\}_{t \geq 0}$ is a scalar-valued autoregressive process with order one, that is, an AR(1) process, while the scalar-valued return $\{\mathbf{Y}_t\}_{t \geq 1}$ is the observable determined by the volatility. We arrange the unknown parameters to appear in both the state transition process and the observation process. The system is described by*

$$\begin{cases} \mathbf{X}_t = \gamma \mathbf{X}_{t-1} + \sigma \varepsilon_t^{(x)} \\ \mathbf{Y}_t = \varepsilon_t^{(y)} \beta \exp(\frac{1}{2} \mathbf{X}_t) \end{cases},$$

where $\Theta = (\gamma, \sigma, \beta)$ are model parameters, $\varepsilon_t^{(x)}$ and $\varepsilon_t^{(y)}$ are independent and identically distributed (i.i.d.) standard Gaussian random variables, and the initial state is given as $\mathbf{X}_0 \sim \mathcal{N}(0, \frac{\sigma^2}{1-\gamma^2})$. The goal is to estimate the parameter Θ and the hidden states $\{\mathbf{X}_t\}_{t \geq 0}$ from the observed returns $\{\mathbf{Y}_t\}_{t \geq 1}$.

1.1 Sequential learning problems

Following the state-space model defined in (1) and (2), random variables $(\boldsymbol{\theta}, \mathbf{X}_{0:t}, \mathbf{Y}_{1:t})$ have the joint density

$$p(\boldsymbol{\theta}, \mathbf{x}_{0:t}, \mathbf{y}_{1:t}) = p(\boldsymbol{\theta}) p(\mathbf{x}_0|\boldsymbol{\theta}) \prod_{j=1}^t \left(f(\mathbf{x}_j|\mathbf{x}_{j-1}, \boldsymbol{\theta}) g(\mathbf{y}_j|\mathbf{x}_j, \boldsymbol{\theta}) \right), \quad (3)$$

where $p(\boldsymbol{\theta})$ and $p(\mathbf{x}_0|\boldsymbol{\theta})$ are prescribed prior densities. Conditioned on all available data $\mathbf{Y}_{1:t} = \mathbf{y}_{1:t}$ at time t , the trajectory of the states $\mathbf{X}_{0:t}$ and the parameter $\boldsymbol{\theta}$ jointly follow the posterior density

$$p(\boldsymbol{\theta}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) = \frac{p(\boldsymbol{\theta}, \mathbf{x}_{0:t}, \mathbf{y}_{1:t})}{p(\mathbf{y}_{1:t})}, \quad (4)$$

where $p(\mathbf{y}_{1:t})$ is an unknown constant commonly referred to as the evidence. We aim to design sequential algorithms that simultaneously solve the following inference problems at each time t :

- **Filtering** that infers the current state, which is the marginal conditional random variable

$$(\mathbf{X}_t|\mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim p(\mathbf{x}_t|\mathbf{y}_{1:t}) := \int p(\boldsymbol{\theta}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) d\mathbf{x}_{0:t-1} d\boldsymbol{\theta}. \quad (5)$$

- **Parameter estimation** that infers the unknown parameter, which is the marginal conditional random variable

$$(\Theta | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim p(\theta | \mathbf{y}_{1:t}) := \int p(\theta, \mathbf{x}_{0:t} | \mathbf{y}_{1:t}) d\mathbf{x}_{0:t}. \quad (6)$$

- **Path estimation** that infers the trajectory of the states, which are the marginal conditional random variables

$$(\mathbf{X}_{0:t} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim p(\mathbf{x}_{0:t} | \mathbf{y}_{1:t}) := \int p(\theta, \mathbf{x}_{0:t} | \mathbf{y}_{1:t}) d\theta. \quad (7)$$

- **Smoothing.** As a result of increasing state dimensions and the observation size, path estimation becomes increasingly challenging over time. Smoothing is a related learning problem focusing on the lower-dimensional marginal conditional random variable

$$(\mathbf{X}_k | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim p(\mathbf{x}_k | \mathbf{y}_{1:t}) := \int p(\theta, \mathbf{x}_{0:t} | \mathbf{y}_{1:t}) d\mathbf{x}_{0:k-1} d\mathbf{x}_{k+1:t} d\theta \quad (8)$$

for some previous time index $k < t$.

In the presence of unknown parameters and stochastic noises in the state transition density (1), all the above-mentioned learning problems are essentially characterizations of marginal random variables. Furthermore, the random variables of interest are intractable, in the sense that the characteristic functions and moments often do not admit analytical forms, and it is infeasible to directly simulate these random variables. Thus, we need to design numerical methods to solve these problems. In principle, algorithms for solving these problems need to be *sequential* by design—the computation of new solutions to the problems (5)–(8) at time t only relies on either previous solutions at $t - 1$ for forward algorithms, or solutions at $t + 1$ for backward algorithms. We refer readers to Evensen et al. (2022); Kantas et al. (2009); Reich and Cotter (2015) and Särkkä (2013) for further details and references of these sequential learning problems.

1.2 Related work and our contributions

Assuming the parameter θ is known, the classical filtering problem aims to estimate the distribution of the current state conditioned on the data available up to t , that is, $(\mathbf{X}_t | \Theta = \theta, \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$. For linear models with Gaussian noises, the filtering density yields a closed-form solution given by the classical Kalman filter (Kalman, 1960). The extended Kalman filter (Anderson and Moore, 2012; Einicke and White, 1999) and the ensemble Kalman filter (Evensen, 2003) generalize the Kalman filter to nonlinear problems. These generalizations use linearization and Monte Carlo sampling to create Gaussian approximations and sample-based approximations of posterior densities. As a result, these methods are flexible to implement and often robust with respect to dimensionality. However, the Gaussian assumption used by various Kalman filters introduces unavoidable approximation errors to the resulting filtering density, which makes them only suitable for tracking the states for nonlinear problems.

Beyond the Gaussian form, sequential Monte Carlo (SMC) methods recursively maintain a set of weighted particles to fully characterize the filtering density (5) for nonlinear

problems. In the simplest form, the bootstrap filter (Gordon et al., 1993) updates weighted particles using the state transition process and reweights them using the likelihood function with newly collected data. It is well-known that the weights of the bootstrap filter tend to degenerate over time, a phenomenon commonly referred to as the particle degeneracy (Doucet and Johansen, 2009; Snyder et al., 2008; Liu et al., 2001; Pitt and Shephard, 1999). To mitigate the particle degeneracy, the auxiliary particle filter (Pitt and Shephard, 1999, 2001) and various advanced resampling techniques such as residual resampling (Carpenter et al., 1999), systematic resampling (Kitagawa, 1996), and resample-move (Gilks and Berzuini, 2001) have been developed. See Del Moral et al. (2012); Doucet and Johansen (2009); Reich and Cotter (2015) and Maskell and Gordon (2002) for comprehensive reviews of these techniques.

An emerging trend is the development of transport map methods that overcome the particle degeneracy by transforming the forecast particles into equally weighted particles following the filtering density. The ensemble transform filter (Reich, 2013) builds such transformations via a discrete optimal transport problem. The coupling technique of Spantini et al. (2022) relaxes the Gaussian assumption of the ensemble Kalman filter by building conditional Knothe–Rosenblatt (KR) rearrangements (Knothe et al., 1957; Rosenblatt, 1952). The implicit sampling method (Chorin and Tu, 2009; Morzfeld et al., 2012) couples a reference density, for example, a standard Gaussian, to a particular approximation of the filtering density via implicit maps. Methods based on neural networks are also developed. For example, Gottwald and Reich (2021) build random feature maps on delayed coordinates to accelerate filtering, and Hoang et al. (2021) approximate conditional mean filters using neural networks.

One of our innovations is treating the filtering problem and other aforementioned sequential learning problems as recursive function approximation problems. This perspective opens the door to applying scalable function approximation tools in sequential learning. In particular, we employ tensor-train (TT) decompositions (Hackbusch, 2012; Oseledets and Tyrtshnikov, 2010; Oseledets, 2011) to design non-parametric solution ansatzes for recursively approximating probability densities in sequential learning. As a result, our new algorithms significantly generalize the linear Gaussian assumptions used by various Kalman filters and bypass sample-based empirical approximations used in various particle filters. This not only makes our algorithms computationally efficient but also capable of fully quantifying uncertainties in sequential learning problems. The computed TT decompositions in our algorithms naturally lead to a sequence of KR rearrangements, enabling us to design an accompanying particle filter to correct for biases due to function approximation errors. Moreover, we use these KR rearrangements to develop a backward sampling algorithm, in line with many particle smoothing methods—for example, Bresler (1986); Briers et al. (2010); Fearnhead et al. (2010); Godsill et al. (2004)—to solve the path estimation and smoothing problems.

Another significant challenge in sequential learning problems is to estimate unknown parameters, which intrinsically requires marginalizing over the state path to evaluate the posterior parameter density (6). The particle Markov chain Monte Carlo (MCMC) method (Andrieu et al., 2010) and the method developed in Särkkä et al. (2015) use inner-loop particle filters or Gaussian filters to estimate the computationally intractable marginalized posterior parameter density $p(\boldsymbol{\theta}|\mathbf{y}_{1:t})$, and then build outer-loop Markov chain transition

kernels to sample posterior parameters. These methods are designed for batched data and thus may have difficulties updating the solution in an online manner with new observations. To estimate posterior parameters online, the SMC² method (Chopin et al., 2013) and its variants—for example, the nested particle filter (Crisan and Miguez, 2018)—run an outer-loop particle filter to sequentially update the parameter distribution. Since these methods are based on particle filters, they often need a restart step to handle the particle degeneracy. As a result, they need to sacrifice either the computational complexity (SMC² has a computational complexity quadratic in time) or the rate of convergence (the nested particle filter has a convergence rate of $N^{-1/4}$, where N is the sample size) in online estimation.

Our new algorithms are directly applicable to the online parameter estimation problem. In fact, our new algorithms solve all four aforementioned sequential learning problems in an integrated manner by recursively approximating the joint posterior density of parameters and states over time. The separability of TT decompositions naturally leads to marginal posterior densities for solving the parameter estimation and filtering problems. This also allows us to build accompanying KR rearrangements to correct approximation biases and solve the smoothing and path estimation problems. More importantly, we show that the accumulation of these density approximation errors has a linear rate under mild assumptions. Within the same framework, we also present preconditioning techniques to improve the approximation power of TT decompositions. On a range of numerical examples, we show that our method outperforms the SMC² method and can produce meaningful results in real-world applications.

1.3 Outline

We propose a basic version of the TT-based algorithm that recursively approximates the time-varying joint posterior densities in Section 2. The separable form of TT provides approximations to the marginal densities required in learning problems (5)–(8). The basic TT-based algorithm is useful to outline the design principles and key steps of our proposed methods. However, the rank truncation used by TT may not preserve non-negativity of density functions. This leads to difficulties in removing the approximation bias using techniques such as importance sampling. We overcome this limitation in Section 3 by integrating the square-root approximation technique of Cui and Dolgov (2022) into the basic algorithm of Section 2. This preserves the recursive learning pattern of the basic algorithm for problems (5)–(8) and leads to non-negative approximations by construction. More importantly, the separable form of the non-negative TT naturally leads to a sequence of KR rearrangements that couple reference random variables and the marginal random variables of interest. We use the resulting (conditional) KR rearrangements to design particle filters and particle smoothers accompanying TT approximations to remove estimation biases.

The smoothing procedure accompanying our TT-based learning algorithm shares many similarities with the joint KR rearrangements developed in Spantini et al. (2018). In particular, our joint KR rearrangement for the simultaneous path and parameter estimation (see Section 3.4) follows the same sparsity pattern as that of Spantini et al. (2018), as a result of the hidden Markov structure of the learning problem. Our key innovation lies in algorithmic design. The work of Spantini et al. (2018) adopts a *variational* approach that builds an approximate map by minimizing the statistical divergence of the pushforward of some ref-

erence density under the candidate map from the target density. The training procedure of the variational approach is often quite involved—the objective function presents many local minima and each optimization iteration requires repeated evaluations of the target density at transformed reference variables under the candidate map. In comparison, our algorithmic design uses multilinear function approximations to directly approximate marginal and conditional densities of interest. In addition to the benefit of bypassing nonlinear optimization, we can exploit the TT-based approximations to compute the *exact* KR rearrangement of the *approximate* density in both lower-triangular and upper-triangular forms. This enables us to solve the filtering problem (5) and the smoothing problem (8) using different conditional KR rearrangements constructed from the same approximate density.

In Section 4, we discuss the error accumulation rate of our proposed algorithms. In Section 5, we present a set of preconditioning techniques for our sequential learning problems to further enhance the approximation power of TT. These preconditioning techniques are also given in the form of KR rearrangements. In Section 6, we present a set of numerical experiments to demonstrate the efficacy and efficiency of our proposed methods and compare them to particle-based methods. Our code is accessible online via <https://github.com/DeepTransport/tensor-ssm-paper-demo>.

1.4 Notation

The dimensionalities of the state space \mathcal{X} , the data space \mathcal{Y} , and the parameter space Θ are m , n , and d , respectively. We assume that the spaces \mathcal{X} , \mathcal{Y} , and Θ can be expressed as Cartesian products. We denote the j -th element of a vector \mathbf{x} by x_j . The index $t \in \mathbb{N}$ is specifically reserved for denoting the time throughout the paper. For a state random vector \mathbf{X}_t and its realization \mathbf{x}_t at time t , their j -th elements are denoted by $X_{t,j}$ and $x_{t,j}$, respectively. The same convention applies to the data vector \mathbf{Y}_t and its realization \mathbf{y}_t .

For a vector $\mathbf{x}_t \in \mathbb{R}^m$ and an index j , it is convenient to group a subset of elements as follows. The vector $\mathbf{x}_{t,<j} = (x_{t,1}, \dots, x_{t,j-1})$ collects the first $j - 1$ elements, and the vector $\mathbf{x}_{t,>j} = (x_{t,j+1}, \dots, x_{t,m})$ collects the last $m - j$ elements. Similarly, we have $\mathbf{x}_{t,\leq j} = (x_{t,1}, \dots, x_{t,j})$, $\mathbf{x}_{t,\geq j} = (x_{t,j}, \dots, x_{t,m})$, and $\mathbf{x}_{t,\leq m} \equiv \mathbf{x}_{t,\geq 0} \equiv \mathbf{x}_t$.

We consider probability measures absolutely continuous with respect to the Lebesgue measure. We denote normalized posterior probability densities in the sequential learning problems by p and its unnormalized version by π . Approximations to these densities are denoted by \hat{p} and $\hat{\pi}$, respectively. The Hellinger distance between random variables with densities p and \hat{p} is defined by

$$D_{\text{H}}(p, \hat{p}) = \left(\frac{1}{2} \int (\sqrt{p(\mathbf{x})} - \sqrt{\hat{p}(\mathbf{x})})^2 d\mathbf{x} \right)^{\frac{1}{2}}.$$

We denote the m -dimensional uniform random variable on a unit hypercube by $\Xi \sim \text{uniform}(\boldsymbol{\xi}; [0, 1]^m)$, in which we drop $[0, 1]^m$ in the definition when no confusion arises.

Consider a diffeomorphism $\mathcal{S}: \mathcal{X} \rightarrow \mathcal{U}$, where $\mathcal{X}, \mathcal{U} \subseteq \mathbb{R}^m$. The density of the transformed variable $\mathcal{S}(\mathbf{X})$ where $\mathbf{X} \sim p$, is the *pushforward* of p under \mathcal{S} , which takes the form

$$\mathcal{S}_{\#} p(\mathbf{u}) = p(\mathcal{S}^{-1}(\mathbf{u})) |\nabla \mathcal{S}^{-1}(\mathbf{u})|.$$

Similarly, the density of the transformed variable $\mathcal{S}^{-1}(\mathbf{U})$, where $\mathbf{U} \sim \eta$, is the *pullback* of the density of \mathbf{U} under \mathcal{S} , which takes the form

$$\mathcal{S}^\# \eta(\mathbf{x}) = \eta(\mathcal{S}(\mathbf{x})) |\nabla \mathcal{S}(\mathbf{x})|.$$

Here $|\cdot|$ denotes the absolute value of the determinant of a matrix.

2. TT-based recursive posterior approximation

We first discuss the recursive formula outlining the design principle of sequential learning problems. We then introduce the TT decomposition and present a basic implementation of the TT-based sequential learning algorithm.

2.1 Recursive state and parameter learning

Following the definition of the joint density of $(\Theta, \mathbf{X}_{0:t}, \mathbf{Y}_{1:t})$ in (3), the density of the posterior random variables $(\Theta, \mathbf{X}_{0:t} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ also has a recursive form

$$p(\theta, \mathbf{x}_{0:t} | \mathbf{y}_{1:t}) = \frac{p(\theta, \mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1}) f(\mathbf{x}_t | \mathbf{x}_{t-1}, \theta) g(\mathbf{y}_t | \mathbf{x}_t, \theta)}{p(\mathbf{y}_t | \mathbf{y}_{1:t-1})}, \quad (9)$$

where $p(\mathbf{y}_t | \mathbf{y}_{1:t-1})$ is a computationally intractable conditional evidence. Marginalizing both sides over $\mathbf{X}_{0:t-2}$, the parameter Θ and the adjacent states $(\mathbf{X}_t, \mathbf{X}_{t-1})$ jointly follow the posterior density

$$p(\theta, \mathbf{x}_t, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) = \frac{p(\theta, \mathbf{x}_{t-1} | \mathbf{y}_{1:t-1}) f(\mathbf{x}_t | \mathbf{x}_{t-1}, \theta) g(\mathbf{y}_t | \mathbf{x}_t, \theta)}{p(\mathbf{y}_t | \mathbf{y}_{1:t-1})}. \quad (10)$$

The above formula outlines basic steps needed by a sequential estimation algorithm to solve the filtering problem (5) and the parameter estimation problem (6):

1. At time $t-1$, the posterior random variables $(\Theta, \mathbf{X}_{t-1} | \mathbf{Y}_{1:t-1} = \mathbf{y}_{1:t-1})$ has the density $p(\theta, \mathbf{x}_{t-1} | \mathbf{y}_{1:t-1})$.
2. At time t , using the state transition density $f(\mathbf{x}_t | \mathbf{x}_{t-1}, \theta)$ and the likelihood function $g(\mathbf{y}_t | \mathbf{x}_t, \theta)$, we can compute the joint posterior density $p(\theta, \mathbf{x}_t, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$, up to the unknown conditional evidence $p(\mathbf{y}_t | \mathbf{y}_{1:t-1})$.
3. We obtain the density of the new posterior random variables $(\Theta, \mathbf{X}_t | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ by solving a marginalization problem

$$p(\theta, \mathbf{x}_t | \mathbf{y}_{1:t}) = \int p(\theta, \mathbf{x}_t, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) d\mathbf{x}_{t-1}. \quad (11)$$

In these steps, the marginalization in (11) plays a key role in sequentially updating the joint posterior random variables $(\Theta, \mathbf{X}_t | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$. For example, the bootstrap filter (Gordon et al., 1993) solves the marginalization using a weighted update of conditional samples—it first draws conditional random variables $\mathbf{X}_t^{(i)} | \mathbf{X}_{t-1}^{(i)}$ that follow the state transition density $f(\mathbf{x}_t | \mathbf{X}_{t-1}^{(i)})$ conditioned on each of the previous particles, and then updates the

weights using the likelihood function $g(\mathbf{y}_t | \mathbf{X}_t^{(i)})$. The path estimation (7) and the smoothing (8) can be solved by additional backward propagation steps given the solutions to the filtering and parameter estimation problems. In the rest of this section, we will introduce the TT decomposition that solves the marginalization problem (11) by function approximation, and then present a basic TT-based algorithm to outline the procedure for solving the filtering and parameter estimation problems. Algorithms for solving the smoothing problem will be discussed in later sections.

2.2 TT decomposition

The central computational tool used in this work is the functional TT decomposition (Bigoni et al., 2016; Gorodetsky et al., 2019). Consider we have a general multivariate function $h : \mathcal{X} \rightarrow \mathbb{R}$ where $\mathcal{X} \in \mathbb{R}^m$ is a Cartesian product. Then, one can approximately decompose $h(\mathbf{x})$ in the following form

$$h(\mathbf{x}) \approx \hat{h}(\mathbf{x}) = \sum_{\alpha_0=1}^{r_0} \sum_{\alpha_1=1}^{r_1} \cdots \sum_{\alpha_m=1}^{r_m} \mathbf{H}_1^{(\alpha_0, \alpha_1)}(x_1) \cdots \mathbf{H}_k^{(\alpha_{k-1}, \alpha_k)}(x_k) \cdots \mathbf{H}_m^{(\alpha_{m-1}, \alpha_m)}(x_m),$$

where $r_0 = r_m = 1$ and the summation ranges r_0, r_1, \dots, r_m are called TT ranks. Each scalar-valued univariate function $\mathbf{H}_k^{(\alpha_{k-1}, \alpha_k)}(x_k)$ is represented as a linear combination of a set of ℓ_k basis functions $\{\phi_k^{(1)}(x_k), \dots, \phi_k^{(\ell_k)}(x_k)\}$, which yields

$$\mathbf{H}_k^{(\alpha_{k-1}, \alpha_k)}(x_k) = \sum_{j=1}^{\ell_k} \phi_k^{(j)}(x_k) \mathbf{A}_k[\alpha_{k-1}, j, \alpha_k],$$

where $\mathbf{A}_k \in \mathbb{R}^{r_{k-1} \times \ell_k \times r_k}$ is an order-3 coefficient tensor. Examples of the basis functions include piecewise polynomials, orthogonal functions, radial basis functions, etc. Grouping all scalar-valued univariate functions $\mathbf{H}_k^{(\alpha_{k-1}, \alpha_k)}(x_k)$ for each coordinate x_k yields a matrix-valued function $\mathbf{H}_k(x_k) : \mathcal{X}_k \rightarrow \mathbb{R}^{r_{k-1} \times r_k}$, which is referred to as the k -th tensor core. This way, the decomposed function can also be expressed as a sequence of multiplications of matrix-valued univariate functions, which is given by

$$\hat{h}(\mathbf{x}) = \mathbf{H}_1(x_1) \cdots \mathbf{H}_k(x_k) \cdots \mathbf{H}_m(x_m).$$

The TT decomposition can be efficiently computed via alternating linear schemes together with cross interpolation (Bigoni et al., 2016; Gorodetsky et al., 2019; Oseledets and Tyrtyshnikov, 2010). We employ the functional extension of the alternating minimal energy method with a residual-based rank adaptation of Dolgov and Savostyanov (2014). It requires only $\mathcal{O}(m\ell r^2)$ evaluations of the function f and $\mathcal{O}(m\ell r^3)$ floating point operations, where $\ell = \max_k \ell_k$ and $r = \max_k r_k$. In general, the maximal rank r depends on the dimension m and can be large when the function h concentrates in some part of its domain. Some theoretical results exist that provide rank bounds. For example, the work of Rohrbach et al. (2022) establishes specific bounds for certain multivariate Gaussian densities that depend poly-logarithmically on m , while the work of Griebel and Harbrecht (2023) proves dimension-independent bounds for general functions in weighted spaces with dominating mixed smoothness.

The integral of the factorized function \hat{h} can be simplified to the integral of certain univariate tensor cores. For example, the integral over x_k can be calculated as

$$\begin{aligned} \int \hat{h}(x_1, x_2, \dots, x_m) dx_k &= \int \mathbf{H}_1(x_1) \cdots \mathbf{H}_k(x_k) \cdots \mathbf{H}_m(x_m) dx_k \\ &= \mathbf{H}_1(x_1) \cdots \overline{\mathbf{H}}_k \cdots \mathbf{H}_m(x_m), \end{aligned}$$

where the overlined matrix $\overline{\mathbf{H}}_k = \int \mathbf{H}_k(x_k) dx_k \in \mathbb{R}^{r_{k-1} \times r_k}$ is obtained by integrating the k -th tensor core elementwisely. This way, the cost of the integration problem scales linearly in the number of variables and quadratically in tensor ranks. This opens the door to solving the marginalization step in sequential learning problems.

2.3 Basic algorithm

We integrate the TT decomposition and the recursive formula in Section 2.1 to design a basic TT-based algorithm for solving the sequential learning problems. Although this basic algorithm does not require a particular variable ordering, we order the variables as $(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ to be compatible with algorithms that will be introduced in Section 3.2.

At time $t-1$, we suppose the density of the posterior variables $(\boldsymbol{\Theta}, \mathbf{X}_{t-1} | \mathbf{Y}_{1:t-1} = \mathbf{y}_{1:t-1})$ is approximated by a TT decomposition, that is,

$$p(\boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t-1}) \propto \hat{\pi}(\boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t-1}).$$

Here $h_1(\mathbf{x}) \propto h_2(\mathbf{x})$ denotes that $h_1(\mathbf{x})$ is approximately proportional to $h_2(\mathbf{x})$, that is, they are close to each other after normalization:

$$\frac{h_1(\mathbf{x})}{\int h_1(\mathbf{x}) d\mathbf{x}} \approx \frac{h_2(\mathbf{x})}{\int h_2(\mathbf{x}) d\mathbf{x}}.$$

Then, for the new observed data \mathbf{y}_t , we can recursively approximate the new density of the joint posterior random variables $(\boldsymbol{\Theta}, \mathbf{X}_t | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$, the filtering density of $(\mathbf{X}_t | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ and the posterior parameter density $(\boldsymbol{\Theta} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ as follows.

Algorithm 1: TT-based sequential estimation.

- (a) **Non-separable approximation.** Following (10), the density of the joint posterior random variables $(\mathbf{X}_t, \boldsymbol{\Theta}, \mathbf{X}_{t-1} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ yields a non-separable, unnormalized approximation q_t in the form of

$$q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) := \hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}), \quad (12)$$

which can be evaluated pointwisely.

- (b) **Separable approximation.** We re-approximate the non-separable unnormalized density q_t by a TT decomposition $\hat{\pi}$, that is,

$$\begin{aligned} q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) &\approx \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) \\ &= \mathbf{F}_1(x_{t,1}) \cdots \mathbf{F}_m(x_{t,m}) \mathbf{G}_1(\theta_1) \cdots \mathbf{G}_d(\theta_d) \mathbf{H}_1(x_{t-1,1}) \cdots \mathbf{H}_m(x_{t-1,m}), \end{aligned}$$

where \mathbf{F} , \mathbf{G} and \mathbf{H} denote tensor cores for the state at time t , the parameters, and the state at time $t-1$, respectively.

(c) **Integration.** By integrating the TT-approximation $\hat{\pi}$, we are able to approximate the density of the posterior random variables $(\Theta, \mathbf{X}_t | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ by

$$\begin{aligned} \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta} | \mathbf{y}_{1:t}) &= \int \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) d\mathbf{x}_{t-1} \\ &= \bar{F}_1(x_{t,1}) \cdots \bar{F}_m(x_{t,m}) \bar{G}_1(\theta_1) \cdots \bar{G}_d(\theta_d) (\bar{H}_1 \cdots \bar{H}_m), \end{aligned}$$

and the normalizing constant

$$c_t = \int \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) d\mathbf{x}_t d\boldsymbol{\theta} d\mathbf{x}_{t-1} = \bar{F}_1 \cdots \bar{F}_m \bar{G}_1 \cdots \bar{G}_d \bar{H}_1 \cdots \bar{H}_m.$$

The densities of the posterior parameters $(\Theta | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ and the filtering state $(\mathbf{X}_t | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ can be approximated in a similar way.

In the next step $t + 1$, we can apply the same procedure using the newly computed approximation $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta} | \mathbf{y}_{1:t})$.

The above basic TT-based algorithm reveals some key principles of the algorithms introduced in this paper. Given the existing approximation $\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})$, the joint posterior random variables $(\mathbf{X}_t, \Theta, \mathbf{X}_{t-1} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ yield a non-separable, unnormalized approximate density q_t in step (a) of Alg. 1. Since q_t cannot be directly marginalized, we re-approximate q_t using a new TT in step (b) of Alg. 1, which enables further marginalizations in step (c) of Alg. 1.

In Section 3, we design new algorithms to remove estimation biases due to approximation errors. In Alg. 1, the rank truncation used by TT may not preserve non-negativity of density functions. An analogue is that the truncated singular value decomposition of a non-negative matrix may contain negative entries. Non-negativity is essential for designing transport maps for debiasing. We will overcome this barrier by using an alternative form of the TT decomposition in Section 3. Then in Section 4, we analyze the accumulation of approximation errors over time. Since the complexity and the approximation power of TT-based algorithms crucially rely on the tensor ranks, we develop preconditioning methods in Section 5 to improve the efficiency of our proposed algorithms.

3. Squared-TT algorithms, debiasing and smoothing

We first review the squared-TT method for building KR rearrangements, which is originally introduced in Cui and Dolgov (2022). We then integrate the resulting transport maps into the recursive procedure defined in Section 2.3 to sequentially solve the filtering problem (5) and the parameter estimation problem (6) with sample-based debiasing. Furthermore, we also design an algorithm to solve the path estimation problem (7) and the smoothing problem (8) with additional backward propagation steps.

3.1 Squared TT and KR rearrangement

Consider the normalized target probability density $p(\mathbf{x}) = \frac{1}{z} \pi(\mathbf{x})$ with $\mathbf{x} \in \mathbb{R}^m$, in which z is an unknown constant, and we can only evaluate the unnormalized density $\pi(\mathbf{x})$. To

preserve non-negativity in function approximation, one can decompose the square root of $\pi(\mathbf{x})$, that is,

$$\sqrt{\pi(\mathbf{x})} \approx \phi(\mathbf{x}) = \mathbf{H}_1(x_1) \cdots \mathbf{H}_m(x_m).$$

Then, the approximate density function $\phi(\mathbf{x})^2$ is non-negative for all \mathbf{x} by construction. Given some reference tensor-product probability density $\lambda(\mathbf{x}) := \prod_{i=1}^m \lambda_i(x_i)$ such that $\sup_{\mathbf{x}} \pi(\mathbf{x})/\lambda(\mathbf{x}) < \infty$ and a sufficiently small constant $\tau > 0$, we can further construct a defensive version of the approximate density function

$$\hat{p}(\mathbf{x}) = \frac{1}{\hat{z}} \hat{\pi}(\mathbf{x}), \quad \hat{\pi}(\mathbf{x}) = \phi(\mathbf{x})^2 + \tau \lambda(\mathbf{x}), \quad \hat{z} = \int \hat{\pi}(\mathbf{x}) d\mathbf{x}. \quad (13)$$

The approximate density \hat{p} with the defensive term $\lambda(\mathbf{x})$ satisfies

$$\sup_{\mathbf{x}} \frac{p(\mathbf{x})}{\hat{p}(\mathbf{x})} = \sup_{\mathbf{x}} \frac{\hat{z} p(\mathbf{x})}{\phi(\mathbf{x})^2 + \tau \lambda(\mathbf{x})} < \frac{\hat{z}}{\tau z} \sup_{\mathbf{x}} \frac{\pi(\mathbf{x})}{\lambda(\mathbf{x})} < \infty.$$

We introduce the defensive term $\tau \lambda(\mathbf{x})$ to ensure that the target density $p(\mathbf{x})$ is absolutely continuous with respect to the approximate density $\hat{p}(\mathbf{x})$. This way, when the approximation $\hat{p}(\mathbf{x})$ is used as the proposal density in importance sampling, for example, in the case of correcting the bias of our recursive posterior approximation scheme, the resulting estimators can fulfil the requirement of the central limit theorem. Under mild assumptions, the following lemma establishes the L^2 error of $\sqrt{\hat{\pi}}$ and the Hellinger distance between the normalized density p and its approximation \hat{p} .

Lemma 1 *Suppose the TT approximation ϕ satisfies $\|\phi - \sqrt{\pi}\|_{L^2} \leq \epsilon$ and the constant τ satisfies $\tau \leq \|\phi - \sqrt{\pi}\|_{L^2}^2$. Then, the L^2 error of $\sqrt{\hat{\pi}}$ defined in (13) satisfies $\|\sqrt{\hat{\pi}} - \sqrt{\pi}\|_{L^2} \leq \sqrt{2}\epsilon$. The Hellinger distance between p and its normalized approximation \hat{p} defined in (13) satisfies*

$$D_{\text{H}}(\hat{p}, p) \leq \frac{\sqrt{2}}{\sqrt{z}} \|\sqrt{\hat{\pi}} - \sqrt{\pi}\|_{L^2} \leq \frac{2\epsilon}{\sqrt{z}}.$$

Proof The bounds on the L^2 error and the Hellinger distance are shown in Proposition 4 and Theorem 1 of Cui and Dolgov (2022), respectively. \blacksquare

As outlined in Section 2.3, marginalization of the approximation is a key operation in deriving the recursive algorithm. For the squared TT approximation defined in (13), the following (Cui and Dolgov, 2022, Proposition 2) gives its marginal density.

Proposition 2 *We consider the normalized approximation \hat{p} defined in (13). For the TT decomposition $\phi(\mathbf{x}) = \mathbf{H}_1(x_1) \cdots \mathbf{H}_m(x_m)$ and a given index $1 \leq k < m$, we define the left accumulated tensor core as $\mathbf{H}_{\leq k}(\mathbf{x}_{\leq k}) = \mathbf{H}_1(x_1) \cdots \mathbf{H}_k(x_k) \in \mathbb{R}^{1 \times r_k}$ and the right accumulated tensor core as $\mathbf{H}_{> k}(\mathbf{x}_{> k}) = \mathbf{H}_{k+1}(x_{k+1}) \cdots \mathbf{H}_m(x_m) \in \mathbb{R}^{r_k \times 1}$. Then the marginal density of \hat{p} takes the form*

$$\hat{p}(\mathbf{x}_{\leq k}) = \int \hat{p}(\mathbf{x}) d\mathbf{x}_{> k} = \frac{1}{\hat{z}} \left(\sum_{\gamma_k=1}^{r_k} \left(\sum_{\alpha_k=1}^{r_k} \mathbf{H}_{\leq k}^{(\alpha_k)}(\mathbf{x}_{\leq k}) \mathbf{L}_{> k}^{(\alpha_k, \gamma_k)} \right)^2 + \tau \lambda(\mathbf{x}_{\leq k}) \right), \quad (14)$$

where $\lambda(\mathbf{x}_{\leq k}) = \prod_{i=1}^k \lambda_i(x_i)$ and $\mathbf{L}_{>k} \in \mathbb{R}^{r_k \times r_k}$ is the (lower-triangular) Cholesky decomposition of the accumulated mass matrix

$$\mathbf{M}_{>k}^{(\alpha_k, \beta_k)} = \int \mathbf{H}_{>k}^{(\alpha_k)}(\mathbf{x}_{>k}) \mathbf{H}_{>k}^{(\beta_k)}(\mathbf{x}_{>k}) d\mathbf{x}_{>k}, \quad \alpha_k = 1, \dots, r_k, \quad \beta_k = 1, \dots, r_k.$$

As shown in Cui and Dolgov (2022), the Cholesky decomposition in Proposition 2 can be recursively computed from $k = m - 1$ to $k = 1$. The normalizing constant \hat{z} can also be computed in a similar way as the last iteration of the recursion. Computing all the marginal densities $\hat{p}(x_1), \dots, \hat{p}(\mathbf{x}_{\leq m-1})$ requires $\mathcal{O}(mlr^3)$ floating point operations (flops).

The normalized approximate density \hat{p} defines a new random variable $\hat{\mathbf{X}}$. Then using the marginal densities constructed in Proposition 2, the densities of random variables $\hat{X}_1, \hat{X}_2 | \hat{\mathbf{X}}_{<1}, \dots, \hat{X}_k | \hat{\mathbf{X}}_{<k}, \dots, \hat{X}_m | \hat{\mathbf{X}}_{<m}$ are given by $\hat{p}(x_1)$ and

$$\hat{p}(x_k | \mathbf{x}_{<k}) = \frac{\hat{p}(\mathbf{x}_{\leq k})}{\hat{p}(\mathbf{x}_{<k})} = \frac{\phi(\mathbf{x}_{<k})^2 + \tau \lambda(\mathbf{x}_{\leq k})}{\phi(\mathbf{x}_{<k})^2 + \tau \lambda(\mathbf{x}_{<k})}, \quad k = 2, 3, \dots, m,$$

respectively. The corresponding distribution functions

$$F_1(x_1) = \int_{-\infty}^{x_1} \hat{p}(x'_1) dx'_1 \quad \text{and} \quad F_k(x_k | \mathbf{x}_{<k}) = \int_{-\infty}^{x_k} \hat{p}(x'_k | \mathbf{x}_{<k}) dx'_k$$

define an order-preserving map $\mathcal{F} : \mathbb{R}^m \mapsto [0, 1]^m$ in the form of

$$\mathcal{F}(\mathbf{x}) = [F_1(x_1), \dots, F_k(x_k | \mathbf{x}_{<k}), \dots, F_m(x_m | \mathbf{x}_{<m})]^\top,$$

which is referred to as the KR rearrangement (Rosenblatt, 1952; Knothe et al., 1957). The map \mathcal{F} transforms a random variable $\hat{\mathbf{X}} \sim \hat{p}(\mathbf{x})$ into a uniform random variable $\Xi \sim \text{uniform}(\boldsymbol{\xi})$, that is, $\Xi = \mathcal{F}(\hat{\mathbf{X}})$. Since the k -th component of \mathcal{F} depends on only the previous $k - 1$ coordinates, both \mathcal{F} and its inverse \mathcal{F}^{-1} have triangular structures, and thus can be evaluated dimension-by-dimension.

The construction of the marginal densities in (14), and hence the conditional densities $\{\hat{p}(x_k | \mathbf{x}_{<k})\}_{k=1}^m$, requires $\mathcal{O}(mlr^3)$ flops. These conditional densities can be constructed before evaluating the transformations \mathcal{F} and \mathcal{F}^{-1} . Evaluating the KR rearrangement \mathcal{F} consists of the dimension-by-dimension evaluation of the conditional densities $\hat{p}(x_k | \mathbf{x}_{<k})$ and the computation of the corresponding distribution functions $\{F_k\}_{k=1}^m$. For each sample, the former is essentially a sequence of vector-matrix products with a total cost of $\mathcal{O}(mlr^2)$ flops, while the latter can be computed algebraically exactly with a total cost of $\mathcal{O}(ml(\log \ell + r) + ml)$ flops, where the $\mathcal{O}(ml(\log \ell + r))$ term is for applying pseudo-spectral methods to construct the distribution function and the $\mathcal{O}(ml)$ term is for evaluating the resulting distribution function. This way, evaluating \mathcal{F} for N samples costs $\mathcal{O}(mlr^3 + Nmlr^2 + Nml(\log \ell + r) + Nml)$. For evaluating the inverse transformation \mathcal{F}^{-1} , we need to numerically solve a sequence of root finding problems for inverting each distribution function F_k . Since F_k is strictly increasing and bounded by construction, applying either the regula falsi method or Newton's method for a fixed number of iterations (denoted by c) can accurately invert F_k up to machine precision, where c is often less than 10. This way, evaluating the inverse transform \mathcal{F}^{-1} for N samples costs $\mathcal{O}(mlr^3 + Nmlr^2 + Nml(\log \ell + r) + Nmlc)$ flops. Therefore, for cases where the sample size N is large and the TT rank r is sufficiently high, the evaluations of \mathcal{F} and \mathcal{F}^{-1} cost $\mathcal{O}(Nmlr^2)$ flops.

Remark 3 *The marginal densities $\hat{p}(\mathbf{x}_{\geq k})$ where $k = 2, \dots, m$, can be computed in a similar way from the first coordinate to the last coordinate. This way, the resulting KR rearrangement is upper-triangular, that is,*

$$\mathcal{F}^u(\mathbf{x}) = [F_1(x_1|\mathbf{x}_{>1}), \dots, F_k(x_k|\mathbf{x}_{>k}), \dots, F_m(x_m)]^\top,$$

as the variable dependency follows a reverse order.

3.2 Non-negativity-preserving algorithm and conditional maps

Following the steps of the basic algorithm (Alg. 1), we first define a new sequential learning algorithm using the non-negativity-preserving approximation presented in (13), and then construct the associated conditional KR rearrangements. The conditional KR rearrangements will be used to sequentially generate weighted samples to correct for the approximation bias.

Algorithm 2: Sequential estimation using squared-TT approximations.

- (a) **Non-separable approximation.** At time $t-1$, suppose the posterior random variables $(\Theta, \mathbf{X}_{t-1}|\mathbf{Y}_{1:t-1})$ has a TT-based (unnormalized) approximate density $\hat{\pi}(\boldsymbol{\theta}, \mathbf{x}_{t-1}|\mathbf{y}_{1:t-1})$. The density of the new joint posterior random variables $(\mathbf{X}_t, \Theta, \mathbf{X}_{t-1}|\mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ yields a non-separable, unnormalized approximation

$$q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) := \hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta}|\mathbf{y}_{1:t-1})f(\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\theta})g(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}). \quad (15)$$

- (b) **Separable approximation.** We re-approximate the square root of q_t by a TT decomposition ϕ_t , that is,

$$\begin{aligned} \sqrt{q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})} &\approx \phi_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) \\ &= \mathbf{G}_1(x_{t,1}) \cdots \mathbf{G}_m(x_{t,m}) \mathbf{F}_1(\theta_1) \cdots \mathbf{F}_d(\theta_d) \mathbf{H}_1(x_{t-1,1}) \cdots \mathbf{H}_m(x_{t-1,m}). \end{aligned}$$

and choose a constant τ_t such that $\tau_t \leq \|\phi_t - \sqrt{q_t}\|_{L^2}^2$. Following the construction in (13), this gives a non-negative approximation

$$\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}|\mathbf{y}_{1:t}) = \phi_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})^2 + \tau_t \lambda(\mathbf{x}_t) \lambda(\boldsymbol{\theta}) \lambda(\mathbf{x}_{t-1}). \quad (16)$$

- (c) **Integration.** Applying Proposition 2 from the right variable $x_{t-1,m}$ to the left variable $x_{t-1,1}$, we obtain $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}|\mathbf{y}_{1:t}) = \int \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}|\mathbf{y}_{1:t}) d\mathbf{x}_{t-1}$.

In the next step $t+1$, we can apply the same procedure using the newly computed marginal approximation $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}|\mathbf{y}_{1:t})$.

Using Proposition 2, we can integrate the approximate density $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}|\mathbf{y}_{1:t})$ from the right variable $x_{t-1,m}$ to the left variable $x_{t-1,1}$ to define the lower-triangular KR rearrangement

$\mathcal{F}_t^l : \mathcal{X} \times \Theta \times \mathcal{X} \rightarrow [0, 1]^{2m+d}$ in the form of

$$\mathcal{F}_t^l(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \begin{bmatrix} \mathcal{F}_{t,t}^l(\mathbf{x}_t) \\ \mathcal{F}_{t,\theta}^l(\boldsymbol{\theta} | \mathbf{x}_t) \\ \mathcal{F}_{t,t-1}^l(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\theta}) \end{bmatrix} = \begin{bmatrix} F_{t,t,1}(x_{t,1}) \\ \vdots \\ F_{t,t,m}(x_{t,m} | \mathbf{x}_{t,<m}) \\ \vdots \\ F_{t,\theta,k}(\theta_k | \mathbf{x}_t, \boldsymbol{\theta}_{<k}) \\ \vdots \\ F_{t,t-1,1}(x_{t-1,1} | \mathbf{x}_t, \boldsymbol{\theta}) \\ \vdots \\ F_{t,t-1,m}(x_{t-1,m} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<m}) \end{bmatrix}, \quad (17)$$

in which each scalar-valued function $F_{t,\cdot,\cdot}$ is a (conditional) distribution function. For example, we have

$$F_{t,t-1,k}(x_{t-1,k} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<k}) = \int_{-\infty}^{x_{t-1,k}} \frac{\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<k}, x'_{t-1,k} | \mathbf{y}_{1:t})}{\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<k} | \mathbf{y}_{1:t})} dx'_{t-1,k}. \quad (18)$$

The inverse map $(\mathcal{F}_t^l)^{-1}$ transforms uniform random variables to the approximate posterior random variables $(\widehat{\mathbf{X}}_t, \widehat{\boldsymbol{\Theta}}, \widehat{\mathbf{X}}_{t-1} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim \hat{p}$, where $\hat{p}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) \propto \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$.

Note that the construction of conditional distribution function, for example, (18), does not require the normalizing constant of the approximate density $\hat{\pi}$, because the marginal densities share the same normalizing constant. Thus, by only integrating the state \mathbf{x}_{t-1} dimension-by-dimension from the right variable $x_{t-1,m}$ to the left variable $x_{t-1,1}$, we obtain the lower conditional KR rearrangement

$$\mathcal{F}_{t,t-1}^l(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\theta}) = \begin{bmatrix} F_{t,t-1,1}(x_{t-1,1} | \mathbf{x}_t, \boldsymbol{\theta}) \\ \vdots \\ F_{t,t-1,m}(x_{t-1,m} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<m}) \end{bmatrix}. \quad (19)$$

The following proposition shows that the lower conditional map $\mathcal{F}_{t,t-1}^l$ defines a *backward* sampler, which generates samples backward in time from the conditional density $\hat{p}(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{y}_{1:t})$. This is particularly useful for defining smoothing algorithms.

Proposition 4 *Conditioned on $\widehat{\mathbf{X}}_t = \mathbf{x}_t$ and $\widehat{\boldsymbol{\Theta}} = \boldsymbol{\theta}$, inverting the lower conditional KR rearrangement $\mathcal{F}_{t,t-1}^l$ defined in (19), we transform a uniform random variable to the conditional random variable*

$$(\widehat{\mathbf{X}}_{t-1} | \widehat{\mathbf{X}}_t = \mathbf{x}_t, \widehat{\boldsymbol{\Theta}} = \boldsymbol{\theta}, \mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim \hat{p}(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{y}_{1:t}).$$

Proof Conditioned on $\widehat{\mathbf{X}}_t = \mathbf{x}_t$ and $\widehat{\boldsymbol{\Theta}} = \boldsymbol{\theta}$, the Jacobian of $\mathcal{F}_{t,t-1}^l(\cdot | \mathbf{x}_t, \boldsymbol{\theta})$, denoted by $\mathbf{J} \in \mathbb{R}^{m \times m}$, is a lower triangular matrix with diagonal entries

$$\mathbf{J}_{kk} = \frac{\partial F_{t,t-1,k}(x_{t-1,k} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<k})}{\partial x_{t-1,k}} = \hat{p}(x_{t-1,k} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<k}, \mathbf{y}_{1:t}), \quad k = 1, \dots, m,$$

by (18). Thus, the pullback of the uniform density under $\mathcal{F}_{t,t-1}^l(\cdot | \mathbf{x}_t, \boldsymbol{\theta})$ is

$$(\mathcal{F}_{t,t-1}^l)^\# \text{uniform}(\mathbf{x}_t) = |\mathbf{J}| = \prod_{k=1}^m \hat{p}(x_{t-1,k} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1,<k}, \mathbf{y}_{1:t}) = \hat{p}(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\theta}, \mathbf{y}_{1:t}).$$

This concludes the result. ■

Similarly, one can integrate the state \mathbf{x}_t in the approximation $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$ from the left variable $x_{t,1}$ to the right variable $x_{t,m}$ to define an upper conditional KR rearrangement

$$\mathcal{F}_{t,t}^u(\mathbf{x}_t | \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \begin{bmatrix} F_{t,t,1}(x_{t,1} | \mathbf{x}_{t,>1}, \boldsymbol{\theta}, \mathbf{x}_{t-1}) \\ \vdots \\ F_{t,t,m}(x_{t,m} | \boldsymbol{\theta}, \mathbf{x}_{t-1}) \end{bmatrix}, \quad (20)$$

Inverting the conditional map $\mathcal{F}_{t,t}^u$, we are able to transform a uniform random variable to the conditional random variable

$$(\widehat{\boldsymbol{\Theta}}_t | \widehat{\boldsymbol{\Theta}} = \boldsymbol{\theta}, \widehat{\mathbf{X}}_{t-1} = \mathbf{x}_{t-1}, \mathbf{Y}_{1:t} = \mathbf{y}_{1:t}) \sim \hat{p}(\mathbf{x}_t | \boldsymbol{\theta}, \mathbf{x}_{t-1}, \mathbf{y}_{1:t}).$$

This way, the upper conditional map $\mathcal{F}_{t,t}^u$ defines a *forward* sampler that generates samples forward in time. This will be used in particle filtering.

3.3 Particle filter

The upper conditional map $\mathcal{F}_{t,t}^u$ in (20) naturally define a particle filter accompanying the recursive approximations in Alg. 2. Suppose weighted samples $\{\widehat{\boldsymbol{\Theta}}^{(i)}, \widehat{\mathbf{X}}_{0:t-1}^{(i)}, W_{t-1}^{(i)}\}_{i=1}^N$ follow the joint posterior density $p(\boldsymbol{\theta}, \mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1})$ at time $t-1$. Conditioned on each pair of $\{\widehat{\boldsymbol{\Theta}}^{(i)}, \widehat{\mathbf{X}}_{t-1}^{(i)}\}_{i=1}^N$, we can invert $\mathcal{F}_{t,t}^u$ to obtain a new sample

$$\widehat{\mathbf{X}}_t^{(i)} = (\mathcal{F}_{t,t}^u)^{-1}(\boldsymbol{\Xi}^{(i)} | \widehat{\boldsymbol{\Theta}}^{(i)}, \widehat{\mathbf{X}}_{t-1}^{(i)}), \quad \boldsymbol{\Xi}^{(i)} \sim \text{uniform}(\boldsymbol{\xi}; [0, 1]^m). \quad (21)$$

The sample $\widehat{\mathbf{X}}_t^{(i)}$ follows the conditional density $\hat{p}(\mathbf{x}_t | \boldsymbol{\theta}, \mathbf{x}_{t-1}, \mathbf{y}_{1:t})$. Expanding the sample state path $\widehat{\mathbf{X}}_{0:t-1}^{(i)}$ by $\widehat{\mathbf{X}}_t^{(i)}$, the updated weighted samples $\{\widehat{\boldsymbol{\Theta}}^{(i)}, \widehat{\mathbf{X}}_{0:t}^{(i)}, W_{t-1}^{(i)}\}_{i=1}^N$ jointly follow the normalized density

$$\hat{p}(\mathbf{x}_t | \boldsymbol{\theta}, \mathbf{x}_{t-1}, \mathbf{y}_{1:t}) p(\boldsymbol{\theta}, \mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1}). \quad (22)$$

After applying the sampling procedure in (21), we can update the weights of the updated samples $\{\widehat{\boldsymbol{\Theta}}^{(i)}, \widehat{\mathbf{X}}_{0:t}^{(i)}, W_{t-1}^{(i)}\}_{i=1}^N$ according to the ratio of the new joint posterior density $p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})$ in (9) to the sampling density in (22), which take the form

$$r_f(\boldsymbol{\theta}, \mathbf{x}_{0:t}) = \frac{p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})}{\hat{p}(\mathbf{x}_t | \boldsymbol{\theta}, \mathbf{x}_{t-1}, \mathbf{y}_{1:t}) p(\boldsymbol{\theta}, \mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1})}.$$

Following the recursive form of $p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})$ in (9), we have

$$r_f(\boldsymbol{\theta}, \mathbf{x}_{0:t}) \propto \omega_f(\boldsymbol{\theta}, \mathbf{x}_{t-1:t}) := \frac{f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta})}{\hat{p}(\mathbf{x}_t | \boldsymbol{\theta}, \mathbf{x}_{t-1}, \mathbf{y}_{1:t})}. \quad (23)$$

Thus, by updating the weights using $W_t^{(i)} = W_{t-1}^{(i)} \omega_f(\widehat{\boldsymbol{\Theta}}^{(i)}, \widehat{\mathbf{X}}_{t-1:t}^{(i)})$ followed by the renormalization

$$W_t^{(i)} = \frac{W_{t-1}^{(i)}}{\sum_{i=1}^N W_{t-1}^{(i)}} \quad \text{for } i = 1, \dots, N,$$

we have the reweighed samples $\{\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{0:t}^{(i)}, W_t^{(i)}\}_{i=1}^N$ that follow the joint posterior density $p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})$. Alg. 3 provides details of the particle filter accompanying the sequential learning procedure in Alg. 2.

Algorithm 3: One iteration of the particle filter accompanying Alg. 2.

- (a) At time $t-1$, suppose we have weighted samples $\{\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{0:t-1}^{(i)}, W_{t-1}^{(i)}\}_{i=1}^N$ following the joint density $p(\boldsymbol{\theta}, \mathbf{x}_{0:t-1} | \mathbf{y}_{1:t-1})$.
- (b) At time t , given the approximate density $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$ computed in Alg. 2, apply Proposition 2 to integrate $\hat{\pi}$ from the left variable $x_{t,1}$ to the right variable $x_{t,m}$ to define a conditional KR rearrangement $\mathcal{F}_{t,t}^u$ in the form of (20).
- (c) For each of $\{\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{0:t-1}^{(i)}\}_{i=1}^N$, invert $\mathcal{F}_{t,t}^u$ as in (21) to generate a new sample $\widehat{\mathbf{X}}_t^{(i)}$, and then update the weights according to $W_t^{(i)} = W_{t-1}^{(i)} \omega_f(\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{t-1:t}^{(i)})$, where the function $\omega_f(\cdot, \cdot)$ is defined in (23).
- (d) Renormalize the weights $W_t^{(i)} \leftarrow W_t^{(i)} / \sum_{i=1}^N W_t^{(i)}$.

In Alg. 3, constructing the conditional densities in step (b), which defines the conditional KR rearrangement $\mathcal{F}_{t,t}^u$, requires $\mathcal{O}(mlr^3)$ flops. This step does not involve any sampling, and thus its cost is independent of the sample size N . Then, in step (c), generating N samples from the resulting conditional KR rearrangement requires an additional $\mathcal{O}(Nm\ell(r^2 + \log \ell + r + c))$ flops, where c is the maximum number of iterations used in the root finding algorithms, as discussed in Sec. 3.1.

In summary, Alg. 3 can be viewed as a standard particle filter with TT approximations to be the proposal density. Commonly used resampling techniques in SMC methods can also be used in Alg. 3 as a rejuvenation to re-balance the weights.

3.4 Path estimation and smoothing

Similar to other particle filter algorithms, weights computed by Alg. 3 may degenerate over time. This has an intuitive explanation in the joint state and parameter estimation context. The locations of the parameter samples are fixed at the initial time, so that the weights (and hence the effective sample size) necessarily degenerate because the posterior parameter density may concentrate with more data observed over time. We overcome this issue by drawing random samples from the last approximation $\hat{p}(\mathbf{x}_T, \boldsymbol{\theta}, \mathbf{x}_{T-1} | \mathbf{y}_{1:T})$ —which is conditioned on all available data by construction—and removing the approximation bias by designing a path estimation algorithm accompanying Alg. 2 to replace the particle filtering. The path estimation algorithm also naturally leads to particle smoothing.

Our path estimation algorithm uses the same sequence of non-negativity-preserving approximations constructed in Alg. 2. It starts at the final time T and recursively samples backward in time using the lower conditional map $\mathcal{F}_{t,t-1}^l$ defined in (19). At time T , initial samples $\{\widehat{\mathbf{X}}_{T-1}^{(i)}, \widehat{\mathbf{X}}_T^{(i)}, \widehat{\Theta}^{(i)}\}_{i=1}^N$ are generated from $\hat{p}(\mathbf{x}_T, \boldsymbol{\theta}, \mathbf{x}_{T-1} | \mathbf{y}_{1:T})$ using a full KR

rearrangement \mathcal{F}_T^l defined in (17). Then, at each of steps $t = T-1, \dots, 1$, for each of $\{\hat{\Theta}^{(i)}, \hat{\mathbf{X}}_{t:T}^{(i)}\}_{i=1}^N$, we invert the lower conditional map $\mathcal{F}_{t,t-1}^l$ to obtain

$$\hat{\mathbf{X}}_{t-1}^{(i)} = (\mathcal{F}_{t,t-1}^l)^{-1}(\Xi_{t-1}^{(i)} | \hat{\Theta}^{(i)}, \hat{\mathbf{X}}_t^{(i)}), \quad \Xi_{t-1}^{(i)} \sim \text{uniform}(\boldsymbol{\xi}; [0, 1]^m). \quad (24)$$

Following from Proposition 4, after completing backward recursion, each pair of the parameter sample and state path sample in $\{\hat{\Theta}^{(i)}, \hat{\mathbf{X}}_{0:T}^{(i)}\}_{i=1}^N$ follows the joint density

$$\hat{p}(\boldsymbol{\theta}, \mathbf{x}_{0:T} | \mathbf{y}_{1:T}) = \hat{p}(\mathbf{x}_T, \boldsymbol{\theta}, \mathbf{x}_{T-1} | \mathbf{y}_{1:T}) \prod_{t=1}^{T-1} \hat{p}(\mathbf{x}_{t-1} | \boldsymbol{\theta}, \mathbf{x}_t, \mathbf{y}_{1:t}). \quad (25)$$

Conditioned on $\boldsymbol{\theta}$ and \mathbf{x}_t , the state \mathbf{x}_{t-1} is independent of the data $\mathbf{y}_{t:T}$ observed at future times, and hence we have $p(\mathbf{x}_{t-1} | \boldsymbol{\theta}, \mathbf{x}_t, \mathbf{y}_{1:t}) = p(\mathbf{x}_{t-1} | \boldsymbol{\theta}, \mathbf{x}_t, \mathbf{y}_{1:T})$. Using this identity, the original joint posterior density $p(\boldsymbol{\theta}, \mathbf{x}_{0:T} | \mathbf{y}_{1:T})$ defined in (4) has the same factorized form as the density in (25). Therefore, the above backward sampling procedure can have a reasonable efficiency if the approximations built in Alg. 2 are sufficiently accurate.

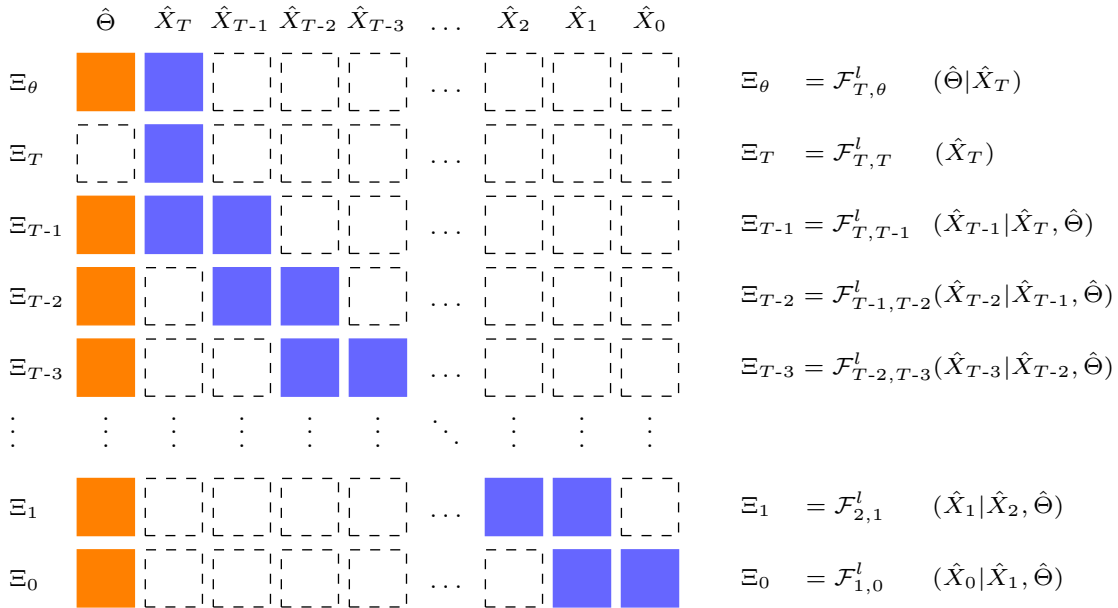


Figure 1: Conditional dependency structures of the joint map \mathcal{F} used in particle smoothing. The blue entries indicate the conditional dependency on the states, and the orange entries indicate the conditional dependency on the parameters. The 2×2 sub-matrix in the top left corner represents the structure of the first two blocks of the map \mathcal{F}_T^l .

In fact, the backward recursion can be considered as the inverse of a joint KR rearrangement $\mathcal{F} : \Theta \times \mathcal{X}^{T+1} \rightarrow [0, 1]^{d+(m+1) \times T}$ that transforms approximate posterior random variables $(\hat{\Theta}, \hat{\mathbf{X}}_{0:T} | \mathbf{Y}_{1:T} = \mathbf{y}_{1:T})$ that follow the density (25), to the reference uniform random variables. As illustrated in Figure 1, the joint KR rearrangement \mathcal{F} has a block sparse structure. The non-zero blocks in the first column reflect the conditional dependency of each

of the state variables on the parameter, while the non-zero blocks on the first sub-diagonal reflect the Markov dependency of state variables. From this perspective, our TT-based construction of the joint KR rearrangement \mathcal{F} can be viewed as a constructive proof and a numerical implementation of the decomposition theorem (Spantini et al., 2018, Theorem 12) for the joint parameter and state estimation in general state-space models.

For a pair of sample parameter and sample state path $(\widehat{\Theta}^{(i)} = \boldsymbol{\theta}, \widehat{\mathbf{X}}_{0:T}^{(i)} = \mathbf{x}_{0:T})$ generated by the backward recursion, it yields a weighted representation of the posterior sample with the importance weight

$$r_b(\boldsymbol{\theta}, \mathbf{x}_{0:T}) = \frac{p(\boldsymbol{\theta}, \mathbf{x}_{0:T} | \mathbf{y}_{1:T})}{\hat{p}(\boldsymbol{\theta}, \mathbf{x}_{0:T} | \mathbf{y}_{1:T})}.$$

Following the recursive form of $p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})$, the ratio $r_b(\boldsymbol{\theta}, \mathbf{x}_{0:T})$ has an unnormalized computable form

$$\begin{aligned} r_b(\boldsymbol{\theta}, \mathbf{x}_{0:T}) &\propto \omega_b(\boldsymbol{\theta}, \mathbf{x}_{0:T}) \\ &:= p(\boldsymbol{\theta})p(\mathbf{x}_0 | \boldsymbol{\theta}) \left(\prod_{t=1}^{T-1} \frac{f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta})g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta})}{\hat{p}(\mathbf{x}_{t-1} | \boldsymbol{\theta}, \mathbf{x}_t, \mathbf{y}_{1:t})} \right) \frac{f(\mathbf{x}_T | \mathbf{x}_{T-1}, \boldsymbol{\theta})g(\mathbf{y}_T | \mathbf{x}_T, \boldsymbol{\theta})}{\hat{p}(\mathbf{x}_T, \boldsymbol{\theta}, \mathbf{x}_{T-1} | \mathbf{y}_{1:T})}. \end{aligned} \quad (26)$$

We summarize the particle smoothing procedure in Alg. 4. The set of weighted samples $\{\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{0:T}^{(i)}, W^{(i)}\}_{i=1}^N$ is an unbiased representation of the joint posterior random variables $(\Theta, \mathbf{X}_{0:T} | \mathbf{y}_{1:T})$. It mitigates the particle degeneracy because the samples are drawn conditioned all observed data $\mathbf{y}_{1:T}$.

Algorithm 4: Path estimation algorithm accompanying Alg. 2.

- (a) Generate samples $\{\widehat{\mathbf{X}}_T^{(i)}, \widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{T-1}^{(i)}\}_{i=1}^N$ from $\hat{p}(\mathbf{x}_T, \boldsymbol{\theta}, \mathbf{x}_{T-1} | \mathbf{y}_{1:T})$ using \mathcal{F}_T^l .
- (b) For $t = T - 1, \dots, 1$, do the following
 - For each of $\{\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{t:T}^{(i)}\}_{i=1}^N$, invert the lower conditional map $\mathcal{F}_{t,t-1}^l$ as in (24) to generate a new sample $\widehat{\mathbf{X}}_{t-1}^{(i)}$.
- (c) For each of $\{\widehat{\Theta}^{(i)}, \widehat{\mathbf{X}}_{0:T}^{(i)}\}_{i=1}^N$, compute the weight $W^{(i)}$ using (26).
- (d) Normalize the weights as $W^{(i)} \leftarrow W^{(i)} / \sum_{i=1}^N W^{(i)}$.

In our setting, applying one step of particle filtering and applying one step of particle smoothing have the same computational complexity. For particle filtering, integrals used for constructing the upper conditional map $\mathcal{F}_{t,t}^u$ in (20) take the opposite direction to the integrals used for building $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta} | \mathbf{y}_{1:t})$ in Alg. 2. This incurs an additional $\mathcal{O}(mlr^3)$ flops. In comparison, constructing the lower conditional map $\mathcal{F}_{t,t-1}^l$ in (19) for particle smoothing shares the same integrals as in Alg. 2, and thus does not incur additional computational cost.

Marginalizing the square TT approximation (16) from either end has an $\mathcal{O}(\ell r^3)$ computational complexity per coordinate. This is much lower than marginalizing variables in the middle, which has an $\mathcal{O}(\ell r^6)$ computational complexity per coordinate. Thus, the variable

ordering $(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ is computationally advantageous for both smoothing and filtering. One can also consider the variable ordering $(\boldsymbol{\theta}, \mathbf{x}_t, \mathbf{x}_{t-1})$ if only the smoothing algorithm is used for debiasing, as Alg. 2 and the lower conditional map in Alg. 4 only need to integrate the state \mathbf{x}_{t-1} .

4. Error analysis

At the core of Alg. 1 and 2 is the recursive approximation of the joint densities of the posterior random variables $(\mathbf{X}_t, \boldsymbol{\Theta}, \mathbf{X}_{t-1} | \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$. We begin this section with an analysis of the accumulation of approximation errors in general state-space models. Based on this result, we provide an upper bound on the TT-approximation errors of Alg. 1 and show the stability of Alg. 2. Although most of the analysis presented centers around the TT decomposition, our results on the error accumulation rate are readily generalizable to other function approximation schemes applied to state-space models. For instance, one can consider the deep polynomial method of Cui et al. (2023) and the sum-of-square approximation of Zanger et al. (2024). The same error accumulation analysis will apply.

4.1 Error decomposition and propagation

Starting with the initial unnormalized joint posterior density

$$\pi(\mathbf{x}_1, \boldsymbol{\theta}, \mathbf{x}_0 | \mathbf{y}_1) = p(\boldsymbol{\theta})p(\mathbf{x}_0 | \boldsymbol{\theta})f(\mathbf{x}_1 | \mathbf{x}_0, \boldsymbol{\theta})g(\mathbf{y}_1 | \mathbf{x}_1, \boldsymbol{\theta}),$$

we define the unnormalized joint posterior density recursively as

$$\begin{aligned} \pi(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) &= \int \pi(\mathbf{x}_{t-1}, \boldsymbol{\theta}, \mathbf{x}_{t-2} | \mathbf{y}_{1:t-1}) d\mathbf{x}_{t-2} f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}) \\ &= \pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}), \end{aligned}$$

which follows a similar derivation to that of (10). We use shorthands $p_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) := p(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$ and $\hat{p}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) := \hat{p}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$ to respectively denote the joint posterior density at time t and its TT-based approximation. Similar notation is adopted for representing their unnormalized counterparts. We omit the input variables when no confusion arises. The density π_t has the normalizing constant $z_t := \int \pi(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) d\mathbf{x}_t d\boldsymbol{\theta} d\mathbf{x}_{t-1}$, which is indeed the evidence $p(\mathbf{y}_{1:t})$ introduced in (4).

Recall that in each iteration of Alg. 1 and 2, the unnormalized joint posterior density π_t cannot be directly evaluated since it is the density of marginal random variables. Instead, we introduce the non-separable, unnormalized approximation q_t in step (a) of both algorithms, and then approximate q_t using a TT-based approximation $\hat{\pi}_t$ in step (b). As shown in (12) and (15), the function q_t takes the form

$$q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}),$$

where $\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})$ is the marginalization of the previous TT-based approximation. The construction of the function q_t enables pointwise function evaluations that are necessary for building the new TT decomposition used by $\hat{\pi}_t$.

This reveals that each iteration of Alg. 1 and 2 can be considered as a two-step approximation procedure: the unnormalized joint density π_t is approximated by the function

q_t , from which we construct the TT-based approximation $\hat{\pi}_t$. Considering the triangle inequality for some distance metric $D(\cdot, \cdot)$, the error $D(\hat{\pi}_t, \pi_t)$ has the decomposition

$$D(\hat{\pi}_t, \pi_t) \leq D(\hat{\pi}_t, q_t) + D(q_t, \pi_t), \quad (27)$$

where $D(\hat{\pi}_t, q_t)$ is the *approximation error* in the current iteration and $D(q_t, \pi_t)$ is the *propagation error* from the previous iteration. In particular, we consider the L^1 distance of unnormalized densities and the L^2 distance of the square roots of unnormalized densities as the distance metrics for bounding these errors. The L^1 distance is useful for approximations that cannot preserve non-negativity (for example, Alg. 1) and is analogous to the total variation distance of probability measures. The L^2 distance on square-root densities is directly related to the Hellinger distance (see Lemma 1) and we use it to analyze Alg. 2. We aim to understand how the approximation errors accumulate over time in the recursive algorithms. Based on the following assumption, we first analyze the propagation error in Propositions 5 and 6.

Assumption 1 *For the state transition process and the observation process, we assume either of the following bounds holds*

$$C_t^{(g)} = \sup_{\mathbf{x}_t \in \mathcal{X}, \boldsymbol{\theta} \in \Theta} g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}) < \infty, \quad (28)$$

$$C_t^{(f)} = \sup_{\mathbf{x}_{t-1} \in \mathcal{X}, \boldsymbol{\theta} \in \Theta} \int f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}) d\mathbf{x}_t < \infty. \quad (29)$$

Generally speaking, any Lipschitz continuous f and g satisfy the above assumptions, which are usually the case in many applications. For example, an observation model with the commonly used Gaussian noise has a bounded $C_t^{(g)}$ for all t .

Proposition 5 *Suppose either bound of Assumption 1 holds. The propagation error $\|q_t - \pi_t\|_{L^1}$ at time t is bounded by the previous total error*

$$\|q_t - \pi_t\|_{L^1} \leq C_t^{(h)} \|\hat{\pi}_{t-1} - \pi_{t-1}\|_{L^1}, \quad h \in \{f, g\}.$$

Proof By the construction of q_t , we express the propagation error as

$$\|q_t - \pi_t\|_{L^1} = \int \left| \hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) - \pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) \right| r(\mathbf{x}_{t-1}, \boldsymbol{\theta}) d\mathbf{x}_{t-1} d\boldsymbol{\theta},$$

where $r(\mathbf{x}_{t-1}, \boldsymbol{\theta}) = \int f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}) d\mathbf{x}_t$. Either condition (28) or (29) leads to $\sup_{\mathbf{x}_{t-1} \in \mathcal{X}, \boldsymbol{\theta} \in \Theta} r(\mathbf{x}_{t-1}, \boldsymbol{\theta}) \leq C_t^{(h)}$ for $h \in \{f, g\}$, and thus we have

$$\|q_t - \pi_t\|_{L^1} \leq C_t^{(h)} \|\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) - \pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})\|_{L^1}.$$

Applying Lemma 15 in Appendix A, the L^1 distance of the marginal densities is bounded by that of the joint densities as

$$\|\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}) - \pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})\|_{L^1} \leq \|\hat{\pi}_{t-1} - \pi_{t-1}\|_{L^1},$$

which leads to the result. ■

Proposition 6 *Suppose either bound of Assumption 1 holds. The propagation error $\|\sqrt{q_t} - \sqrt{\pi_t}\|_{L^2}$ at time t is bounded by the previous total error*

$$\|\sqrt{q_t} - \sqrt{\pi_t}\|_{L^2} \leq \sqrt{C_t^{(h)}} \|\sqrt{\hat{\pi}_{t-1}} - \sqrt{\pi_{t-1}}\|_{L^2}, \quad h \in \{f, g\}.$$

Proof By the construction of q_t , we express the squared propagation error as

$$\|\sqrt{q_t} - \sqrt{\pi_t}\|_{L^2}^2 = \int \left(\sqrt{\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} - \sqrt{\pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} \right)^2 r(\mathbf{x}_{t-1}, \boldsymbol{\theta}) d\mathbf{x}_{t-1} d\boldsymbol{\theta},$$

where $r(\mathbf{x}_{t-1}, \boldsymbol{\theta}) = \int f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}) d\mathbf{x}_t$. Either condition (28) or (29) leads to $\sup_{\mathbf{x}_{t-1} \in \mathcal{X}, \boldsymbol{\theta} \in \Theta} r(\mathbf{x}_{t-1}, \boldsymbol{\theta}) \leq C_t^{(h)}$ for $h \in \{f, g\}$, and thus we have

$$\begin{aligned} \|\sqrt{q_t} - \sqrt{\pi_t}\|_{L^2} &\leq \left(C_t^{(h)} \int \left(\sqrt{\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} - \sqrt{\pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} \right)^2 d\mathbf{x}_{t-1} d\boldsymbol{\theta} \right)^{1/2} \\ &= \sqrt{C_t^{(h)}} \left\| \sqrt{\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} - \sqrt{\pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} \right\|_{L^2}. \end{aligned}$$

Applying Lemma 14 in Appendix A, the L^2 distance between the marginal densities is bounded by the L^2 distance between the joint densities as

$$\left\| \sqrt{\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} - \sqrt{\pi(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})} \right\|_{L^2} \leq \|\sqrt{\hat{\pi}_{t-1}} - \sqrt{\pi_{t-1}}\|_{L^2},$$

which leads to the result. \blacksquare

Based on the error propagation bounds, we can then analyze the total approximation errors of our sequential estimation algorithms. In the following, we show the total errors of Alg. 1 and 2 assuming the approximation error in each iteration is bounded.

Theorem 7 *Suppose either bound of Assumption 1 holds and the approximation $\hat{\pi}_t$ in Alg. 1 satisfies $\|\hat{\pi}_t - q_t\|_{L^1} \leq \epsilon_t$. For $t > 0$, the total error satisfies*

$$\|\hat{\pi}_t - \pi_t\|_{L^1} \leq C_t^{(h)} \|\pi_{t-1} - \hat{\pi}_{t-1}\|_{L^1} + \epsilon_t \leq \sum_{k=1}^t C^{t-k} \epsilon_k, \quad C = \sup_t C_t^{(h)},$$

for $h \in \{f, g\}$.

Proof The result follows from the triangle inequality (27) and induction. \blacksquare

Theorem 8 *Suppose either bound of Assumption 1 holds and the approximation $\hat{\pi}_t$ in Alg. 2 satisfies $\|\sqrt{\hat{\pi}_t} - \sqrt{q_t}\|_{L^2} \leq \epsilon_t$. For $t > 0$, the Hellinger distance between p_t and its approximation \hat{p}_t is bounded by*

$$D_H(\hat{p}_t, p_t) \leq \frac{\sqrt{2}}{\sqrt{p(\mathbf{y}_{1:t})}} \sum_{k=1}^t C^{(t-k)/2} \epsilon_k,$$

where $C = \sup_t C_t^{(h)}$ with $h \in \{f, g\}$ and $p(\mathbf{y}_{1:t})$ is the evidence.

Proof We recall that there is no approximation used in Alg. 2 at the initial step $t = 0$, that is, $\epsilon_0 = 0$, as we only have the prior for \mathbf{X}_0 and Θ . Applying the triangle inequality (27) and induction, we have

$$\|\sqrt{\hat{\pi}_t} - \sqrt{\pi_t}\|_{L^2} \leq \sqrt{C_t^{(h)}} \|\sqrt{\pi_{t-1}} - \sqrt{\hat{\pi}_{t-1}}\|_{L^2} + \epsilon_t \leq \sum_{k=1}^t C^{(t-k)/2} \epsilon_k.$$

Applying Lemma 1, the Hellinger distance between \hat{p}_t and p_t is bounded by the L^2 distance between $\sqrt{\hat{\pi}_t}$ and $\sqrt{\pi_t}$, that is,

$$D_H(\hat{p}_t, p_t) \leq \frac{\sqrt{2}}{\sqrt{z_t}} \|\sqrt{\hat{\pi}_t} - \sqrt{\pi_t}\|_{L^2},$$

where $z_t = p(\mathbf{y}_{1:t})$ by definition. This concludes the result. \blacksquare

By Lemmas 13 and 15 in Appendix A, either the Hellinger distance or the L^1 distance between the exact marginal and the approximate marginal is bounded from above by the corresponding distance between the joint densities. Thus, approximations of the filtering density $p(\mathbf{x}_t | \mathbf{y}_{1:t})$ and the posterior parameter density $p(\boldsymbol{\theta} | \mathbf{y}_{1:t})$ obtained by Alg. 1 and 2, which are marginalizations of the joint approximation $\hat{p}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$, follow the same error bounds derived here.

4.2 Error analysis for TT approximations

For Alg. 1, we first show in Proposition 9 that the approximation error $\|\hat{\pi}_t - q_t\|_{L^1}$ can be bounded based on certain Sobolev-type smoothness assumptions. Then, we can apply the result of Theorem 7 to bound the total error.

Proposition 9 *For any time $t > 0$, we express the product of the state transition density and the likelihood function as*

$$h_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta}) g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta}).$$

We assume that state space \mathcal{X} and the parameter space Θ are compact in the sense that they admit finite Lebesgue measures and there exist some $K \in \mathbb{N}$ and $s \geq 1$ such that $h_t \in W^{K+1, 2s}(\mathcal{X} \times \Theta \times \mathcal{X})$. At time $t = 1$, we further assume that the (unnormalized) prior density $\pi(\mathbf{x}_0, \boldsymbol{\theta}) := \pi(\mathbf{x}_0 | \boldsymbol{\theta}) \pi(\boldsymbol{\theta})$ belongs to $W^{K+1, 2r}(\mathcal{X} \times \Theta)$ for some $r \geq 1$ such that $\frac{1}{r} + \frac{1}{s} = 1$. Then, for some error controlling factor $\epsilon_t \in (0, 1)$ at time $t > 0$, there exists a TT decomposition $\hat{\pi}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ with ranks

$$r_1 = \lceil \epsilon_t^{-1/K} \rceil \quad \text{and} \quad r_k = \lceil \epsilon_t^{-k/K} \rceil \quad \text{for} \quad k = 2, 3, \dots, 2m + d - 1,$$

such that the approximation error in step (b) of Alg. 1 is bounded by

$$\|\hat{\pi}_t - q_t\|_{L^1} \leq C_{\text{TT}} \sqrt{\Omega(\mathcal{X} \times \Theta \times \mathcal{X})(2m + d - 1)} \epsilon_t,$$

where $\Omega(\cdot)$ denotes the measure of the space and C_{TT} is a constant independent of $\hat{\pi}_t$, q_t , and the dimension.

Proof Our goal is to establish the Sobolev-type smoothness of the function q_t , so that results of Griebel and Harbrecht (2023) can be used to derive the TT ranks for some target error measured in the L^2 norm. The bound on the L^1 norm follows from that on the L^2 norm using the Cauchy–Schwartz inequality.

For brevity, we define the previous marginal density $\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})$ as

$$\hat{\varphi}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \begin{cases} \hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1}), & t > 1 \\ \pi(\mathbf{x}_0, \boldsymbol{\theta}), & t = 1 \end{cases},$$

which takes a constant value over \mathbf{x}_t . For $t > 1$, since the marginal TT approximation $\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})$ is a linear combination of multivariate polynomial basis functions, the function $\hat{\varphi}_t$ is an analytical function. Thus, we have $\hat{\varphi}_t \in W^{K+1, 2r}(\mathcal{X} \times \Theta \times \mathcal{X})$ for any $K \in \mathbb{N}$, $r \geq 1/2$, and $t > 1$.

For $t > 0$, the L^2 norm of q_t can be expressed as $\|q_t\|_{L^2}^2 = \|\hat{\varphi}_t^2 h_t^2\|_{L^1}$. Applying the Hölder inequality, we have

$$\|q_t\|_{L^2} \leq \|\hat{\varphi}_t^2\|_{L^r}^{1/2} \|h_t^2\|_{L^s}^{1/2} = \|\hat{\varphi}_t\|_{L^{2r}} \|h_t\|_{L^{2s}},$$

for $r, s \in [0, \infty]$ with $\frac{1}{r} + \frac{1}{s} = 1$. This way, we have $\|q_t\|_{L^2} < \infty$ for all $t > 0$ by our assumptions and the analyticity of $\hat{\varphi}_t$ for $t > 1$. For derivatives of q_t , with the multi-index notation and the general Leibniz rule (see Constantine and Savits (1996)), we have

$$\partial^\alpha q_t = \sum_{\beta \leq \alpha} \binom{\alpha}{\beta} (\partial^{\alpha-\beta} \hat{\varphi}_t) (\partial^\beta h_t),$$

where $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are multi-indices. Following a similar derivation as in the L_2 norm case, we can show that $\|\partial^\alpha q_t\|_{L^2} < \infty$ for all $|\boldsymbol{\alpha}| \leq K + 1$, where $|\boldsymbol{\alpha}| = \sum \alpha_i$. Therefore, we have $q_t \in W^{K+1, 2}(\mathcal{X} \times \Theta \times \mathcal{X})$ for all $t > 0$.

Given a target error controlling factor $\varepsilon_t \in (0, 1)$, Theorem 4 of Griebel and Harbrecht (2023) states that there exists a TT decomposition $\hat{\pi}_t$ with ranks $r_1 = \lceil \varepsilon_t^{-1/K} \rceil$ and $r_k = \lceil \varepsilon_t^{-2k/K} \rceil$ for $k = 2, 3, \dots, 2m+d-1$ such that $\|\hat{\pi}_t - q_t\|_{L^2} \leq C_{\text{TT}} \sqrt{2m+d-1} \varepsilon_t$. Then, by the Cauchy–Schwartz inequality, we have

$$\|\hat{\pi}_t - q_t\|_{L^1} \leq C_{\text{TT}} \sqrt{\Omega(\mathcal{X} \times \Theta \times \mathcal{X})} \|\hat{\pi}_t - q_t\|_{L^2},$$

and thus the result follows. ■

Proposition 9 only considers bounded spaces. This can be easily extended to unbounded spaces by considering a weight measure and a slightly modified approximation scheme. See Cui and Dolgov (2022); Cui et al. (2023) and references therein for details. The result of Griebel and Harbrecht (2023) establishes sufficient conditions of the *a priori* error bound of TT decompositions using smoothness. In practice, there are many examples of low-rank functions that do not satisfy the smoothness assumption, for example, products of discontinuous univariate functions. To the authors’ knowledge, the analysis of TT ranks for general functions is still an active research area and the results of Griebel and Harbrecht (2023) may not provide a precise practical guideline for building TT decompositions. Nonetheless,

Proposition 9 provides some confidence in applying TT decompositions in the sequential state and parameter learning problems.

It is worth mentioning that the error controlling factor ε_t in Proposition 9 determines the TT ranks and hence the computational complexity for building a TT decomposition. Proposition 9 suggests that with a prescribed ε_t , the error of a TT decomposition may increase with the dimensionality of the problem. For a class of high-dimensional problems where the Sobolev norm is equipped with suitably decaying weights (Dick et al., 2013; Griebel and Harbrecht, 2023; Sloan and Woźniakowski, 1998), one can obtain a dimension-free error bound. Although such dimension-free bounds are not discussed here, extending Proposition 9 to such weighted Sobolev norms is trivial.

Corollary 10 *Suppose either bound of Assumption 1 holds and the smoothness of the densities f and g fulfil the requirement of Proposition 9. In each step of Alg. 1, suppose further the ranks of a TT approximation is chosen according to Proposition 9 so that it satisfies*

$$\|\hat{\pi}_t - q_t\|_{L^1} \leq C_{\text{TT}} \sqrt{\Omega(\mathcal{X} \times \Theta \times \mathcal{X})(2m + d - 1)} \varepsilon_t$$

for some error controlling factor $\varepsilon_t \in (0, 1)$. Then, for $t > 0$, the total error satisfies

$$\|\hat{\pi}_t - \pi_t\|_{L^1} \leq C_{\text{TT}} \sqrt{\Omega(\mathcal{X} \times \Theta \times \mathcal{X})(2m + d - 1)} \sum_{k=1}^t C^{t-k} \varepsilon_k,$$

where $C = \sup_t C_t^{(h)}$ with $h \in \{f, g\}$.

Proof The result is a direct consequence of Theorem 7. ■

In Alg. 2, the previous marginal density $\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})$ can be expressed as a sum of squares of multivariate polynomial functions (each in the TT form) to ensure non-negativity together with a defensive term (cf. Proposition 2). Applying the multivariate Faa di Bruno formula (Constantine and Savits, 1996), one can show that the square root of the previous marginal density $\sqrt{\hat{\pi}(\mathbf{x}_{t-1}, \boldsymbol{\theta} | \mathbf{y}_{1:t-1})}$ yields a bounded Sobolev norm for a finite order of differentiability K . However, the bound on the Sobolev norm may increase with K . Although Theorem 4 of Griebel and Harbrecht (2023) may still apply in this case for a finite K , it requires further analysis to understand the impact of increasing Sobolev norm on the *a priori* error analysis of the TT decomposition. In our numerical examples (Section 6.1), we demonstrate that Alg. 1 and 2 achieve comparable approximation accuracy. In the rest, we focus on analyzing the impact of the additional defensive term used in Alg. 2. We bound the approximation error $\|\sqrt{\hat{\pi}_t} - \sqrt{q_t}\|_{L^2}$ in Proposition 11—we omit the proof as it is a direct consequence of Lemma 1. Then in Corollary 12, we apply Theorem 8 to bound the total Hellinger error of Alg. 2 to complete the analysis.

Proposition 11 *In each iteration of Alg. 2, suppose a TT decomposition ϕ_t is constructed according to Proposition 9 such that its L^2 error satisfies*

$$\|\phi_t - \sqrt{q_t}\|_{L^2} < C_{\text{TT}} \sqrt{2m + d - 1} \varepsilon_t$$

for some error controlling factor $\varepsilon_t \in (0, 1)$. Suppose further the constant τ_t in the defensive term satisfies $\tau_t \leq \|\phi_t - \sqrt{q_t}\|_{L^2}^2$. Then, the approximation error satisfies $\|\sqrt{\hat{\pi}_t} - \sqrt{q_t}\|_{L^2} \leq C_{\text{TT}}\sqrt{2(2m+d-1)}\varepsilon_t$.

Corollary 12 *Suppose Proposition 11 and either bound of Assumption 1 hold. For $t > 0$, the total Hellinger distance between p_t and its approximation \hat{p}_t in Alg. 2 satisfies*

$$D_{\text{H}}(\hat{p}_t, p_t) \leq \frac{2C_{\text{TT}}\sqrt{2m+d-1}}{\sqrt{p(\mathbf{y}_{1:t})}} \sum_{k=1}^t C^{(t-k)/2} \varepsilon_k,$$

where $C = \sup_t C_t^{(h)}$ with $h \in \{f, g\}$ and $p(\mathbf{y}_{1:t})$ is the evidence.

5. Preconditioning methods

The key step of Alg. 2 is to approximate the square root of the non-separable, unnormalized density q_t by a TT decomposition $\phi_{q,t}$. In many applications, the temporally increasing data size and complex nonlinear interactions among parameters and states may concentrate posterior densities to some submanifold, and hence lead to potentially high ranks in $\phi_{q,t}$. This can make the TT-based algorithms computationally demanding. Rather than directly approximating q_t , here we present a preconditioning framework to improve TT's approximation efficiency and discuss how to apply the resulting preconditioned approximations in our sequential estimation algorithms (Alg. 2–4).

5.1 General framework

Our preconditioning procedure is guided by a (possibly unnormalized) bridging density $\rho_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ that is easier to approximate than $q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$. We introduce general reference random variables $(\mathbf{U}_t, \mathbf{U}_\theta, \mathbf{U}_{t-1})$ with tensor-product normalized density $\eta(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) = \eta(\mathbf{u}_t)\eta(\mathbf{u}_\theta)\eta(\mathbf{u}_{t-1})$, where $(\mathbf{U}_t, \mathbf{U}_{t-1})$ take values in \mathbb{R}^m and \mathbf{U}_θ takes values \mathbb{R}^d . The preconditioning procedure has the following conceptual steps.

1. **Change of coordinates.** We first construct a *preconditioning KR rearrangement* \mathcal{T}_t such that

$$\begin{aligned} (\mathcal{T}_t)_\# \rho_t(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) &= \rho_t(\mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})) |\nabla \mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})| \\ &\propto \eta(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}). \end{aligned} \quad (30)$$

The map \mathcal{T}_t defines a change of coordinates from $(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ to $(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})$.

2. **Preconditioning.** Applying the identity in (30), the pushforward of q_t under \mathcal{T}_t takes form

$$\begin{aligned} (\mathcal{T}_t)_\# q_t(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) &= q_t(\mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})) |\nabla \mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})| \\ &\propto q_{\#,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}), \end{aligned}$$

where

$$q_{\#,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) := \frac{q_t(\mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}))}{\rho_t(\mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}))} \eta(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) \quad (31)$$

is a non-negative function that can be evaluated pointwise. The pushforward density $q_{\sharp,t}$ can be viewed as the reference density η perturbed by the ratio q_t/ρ_t in the transformed coordinates $(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})$. With a suitable bridging density, the ratio q_t/ρ_t is significantly less concentrated, and hence may be easier to approximate. See Fig. 2 for an illustration.

3. **TT-approximation.** By approximating $\sqrt{q_{\sharp,t}}$ using a TT $\phi_{\sharp,t}$, that is,

$$\sqrt{q_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})} \approx \phi_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}),$$

we can follow the squared-TT approximation outlined in Section 3.1 to define an unnormalized approximate density

$$\hat{\nu}_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) := \phi_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})^2 + \tau_{\sharp,t}\eta(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) \quad (32)$$

where $\tau_{\sharp,t} \leq \|\phi_{\sharp,t} - \sqrt{q_{\sharp,t}}\|_{L^2}^2$. Applying Proposition 2 and denoting the normalized density by $\hat{\mu}_{\sharp,t} \propto \hat{\nu}_{\sharp,t}$, we obtain a KR rearrangement \mathcal{S}_t such that

$$(\mathcal{S}_t)_{\sharp} \hat{\mu}_{\sharp,t}(\boldsymbol{\xi}_t, \boldsymbol{\xi}_\theta, \boldsymbol{\xi}_{t-1}) = \text{uniform}(\boldsymbol{\xi}_t, \boldsymbol{\xi}_\theta, \boldsymbol{\xi}_{t-1}).$$

4. **Composition.** The above steps define a composite transformation $\mathcal{S}_t \circ \mathcal{T}_t$ that approximately pushes forward the normalized version of the density $q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ to a uniform density. Equivalently, we have

$$q_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) \propto (\mathcal{S}_t \circ \mathcal{T}_t)_{\sharp} \text{uniform}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}).$$

Thus, the pullback of the uniform density under $\mathcal{S}_t \circ \mathcal{T}_t$ defines a normalized approximate posterior density, which takes the form

$$\begin{aligned} \hat{p}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) &= (\mathcal{S}_t \circ \mathcal{T}_t)_{\sharp} \text{uniform}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) \\ &= |\nabla(\mathcal{S}_t \circ \mathcal{T}_t)(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})| \\ &\stackrel{(30)}{=} \frac{\hat{\mu}_{\sharp,t}(\mathcal{T}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}))}{\eta(\mathcal{T}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}))} \rho_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) \\ &= \frac{1}{\hat{z}_t} \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}), \end{aligned}$$

where

$$\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) = \frac{\hat{\nu}_{\sharp,t}(\mathcal{T}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}))}{\eta(\mathcal{T}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}))} \rho_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) \quad (33)$$

is the unnormalized approximate posterior density and

$$\hat{z}_t = \int \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) d\mathbf{x}_t d\boldsymbol{\theta} d\mathbf{x}_{t-1} = \int \hat{\nu}_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) d\mathbf{u}_t d\mathbf{u}_\theta d\mathbf{u}_{t-1},$$

is the normalizing constant. The last equality follows from the change of coordinates defined by \mathcal{T}_t .

There are many ways to defining the bridging density ρ_t and building the preconditioning map \mathcal{T}_t . In the following, we provide some examples of bridging densities relying on the particular structure of q_t , which is defined in (15).




Figure 2: Left: approximation of q_t (gray contours) using a single layer of TT-based approximation (red contours), in which Fourier basis of order 30 and a TT rank 24 is used. Middle: preconditioned density using the tempering technique (cf. Section 5.3), in which the nonlinear preconditioning transform is defined by a TT with Fourier basis of order 30 and a rank 12. Right: approximation of q_t (gray contours) using the preconditioned approximation (blue contours), in which Fourier basis of order 30 and a TT rank 12 is used in the preconditioned approximation.

5.2 Gaussian bridging and linear preconditioning

We consider a Gaussian approximation to the density q_t as the bridging density, that is, $\rho_t(\cdot) := \mathcal{N}(\cdot; \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, where $\boldsymbol{\mu}_t \in \mathbb{R}^{2m+d}$ is the mean vector and $\boldsymbol{\Sigma}_t \in \mathbb{R}^{(2m+d) \times (2m+d)}$ is the covariance matrix. In each iteration, we can draw random variables from the previous marginal approximation $\hat{\pi}(\boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t-1})$ and use one step of particle filter—for example, a bootstrap filter—to estimate $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ using particles. With $\rho_t(\cdot) := \mathcal{N}(\cdot; \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, we compute the Cholesky factorization of $\boldsymbol{\Sigma}_t$ to obtain $\mathbf{L}_t \mathbf{L}_t^\top = \boldsymbol{\Sigma}_t$, where \mathbf{L}_t is lower-triangular. This defines a lower-triangular linear KR rearrangement $\mathcal{T}_t^l(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \mathbf{L}_t^{-1}([\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}]^\top - \boldsymbol{\mu}_t)$ that transforms $(\mathbf{X}_t, \boldsymbol{\Theta}, \mathbf{X}_{t-1}) \sim \rho_t$ into reference variables $(\mathbf{U}_t, \mathbf{U}_\theta, \mathbf{U}_{t-1})$ following a tensor-product zero mean standard Gaussian density η . The upper-triangular linear KR rearrangement can also be derived by permuting the variables.

5.3 Tempering and nonlinear preconditioning

We employ the tempering idea, for example, Beskos et al. (2016); Gelman and Meng (1998); Herbst and Schorfheide (2019) and Kantas et al. (2014), to build an unnormalized bridging density

$$\rho_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \hat{\pi}(\boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t-1})^{\beta_\pi} f(\mathbf{x}_t | \mathbf{x}_{t-1}, \boldsymbol{\theta})^{\beta_f} g(\mathbf{y}_t | \mathbf{x}_t, \boldsymbol{\theta})^{\beta_g}, \quad (34)$$

with some constants $\beta_\pi, \beta_f, \beta_g \in [0, 1]$. When $\beta_\pi = \beta_f = \beta_g = 1$, it recovers the target density q_t . By choosing appropriate constants $(\beta_\pi, \beta_f, \beta_g)$ —for example, using the adaptation strategy discussed in Beskos et al. (2016) and Kantas et al. (2014)—to construct ρ_t that is less concentrate than q_t but retains some features of q_t . We can apply the squared approximation procedure to decompose $\sqrt{\rho_t}$ into a TT $\phi_{\rho,t}$, approximate ρ_t by

$$\hat{\rho}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) := \phi_{\rho,t}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})^2 + \tau_{\rho,t} \lambda(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}), \quad (35)$$

with some $\tau_{\rho,t} > 0$, and derive the corresponding KR rearrangement \mathcal{R}_t such that

$$(\mathcal{R}_t)_\# \hat{\rho}_t(\boldsymbol{\xi}_t, \boldsymbol{\xi}_\theta, \boldsymbol{\xi}_{t-1}) \propto \text{uniform}(\boldsymbol{\xi}_t, \boldsymbol{\xi}_\theta, \boldsymbol{\xi}_{t-1}).$$

Given a diagonal map \mathcal{D} such that $\mathcal{D}_\# \eta(\boldsymbol{\xi}_t, \boldsymbol{\xi}_\theta, \boldsymbol{\xi}_{t-1}) = \text{uniform}(\boldsymbol{\xi}_t, \boldsymbol{\xi}_\theta, \boldsymbol{\xi}_{t-1})$, composing \mathcal{D}^{-1} and \mathcal{R}_t , we obtain $\mathcal{T}_t = \mathcal{D}^{-1} \circ \mathcal{R}_t$ such that

$$(\mathcal{T}_t)_\# \hat{\rho}_t(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) \propto \eta(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}).$$

In this construction, the preconditioned target density $q_{\#,t}$ —which is originally defined in (31)—becomes

$$q_{\#,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) := \frac{q_t(\mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}))}{\hat{\rho}_t(\mathcal{T}_t^{-1}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}))} \eta(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}),$$

where the exact bridging density ρ_t is replaced with $\hat{\rho}_t$, which is the pushforward of the reference density under the preconditioning map \mathcal{T}_t . As a consequence, the unnormalized approximate density in (33) becomes

$$\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) = \frac{\hat{\nu}_{\#,t}(\mathcal{T}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}))}{\eta(\mathcal{T}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}))} \hat{\rho}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}).$$

5.4 Marginalization and conditional maps

After building the preconditioned approximation, we need to integrate the approximate density $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t})$ over \mathbf{x}_{t-1} , so that the marginal density

$$\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta} | \mathbf{y}_{1:t}) = \int \hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1} | \mathbf{y}_{1:t}) d\mathbf{x}_{t-1}$$

can be used in the next iteration of Alg. 2. In addition, we also need to derive the lower and upper conditional KR rearrangements to be used in path estimation and particle filtering, respectively. We use the lower-triangular KR rearrangements \mathcal{S}_t^l and \mathcal{T}_t^l , which respectively take the form

$$\mathcal{S}_t^l(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1}) = \begin{bmatrix} \mathcal{S}_{t,t}^l(\mathbf{u}_t) \\ \mathcal{S}_{t,\theta}^l(\mathbf{u}_\theta | \mathbf{u}_t) \\ \mathcal{S}_{t,t-1}^l(\mathbf{u}_{t-1} | \mathbf{u}_t, \mathbf{u}_\theta) \end{bmatrix} \quad \text{and} \quad \mathcal{T}_t^l(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1}) = \begin{bmatrix} \mathcal{T}_{t,t}^l(\mathbf{x}_t) \\ \mathcal{T}_{t,\theta}^l(\boldsymbol{\theta} | \mathbf{x}_t) \\ \mathcal{T}_{t,t-1}^l(\mathbf{x}_{t-1} | \mathbf{x}_t, \boldsymbol{\theta}) \end{bmatrix},$$

as well as nonlinear preconditioning defined above, to outline how to carry the marginalizations and conditional sampling.

Because the composition $\mathcal{S}_t^l \circ \mathcal{T}_t^l$ is also lower triangular, we can apply a similar derivation to that of Proposition 4 to compute the marginal density $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta} | \mathbf{y}_{1:t})$. The necessary steps for computing $\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta} | \mathbf{y}_{1:t})$ is integrated with the preconditioned density approximation procedure, and thus we summarize them together in Alg. 5, which replace steps (b) and (c) of the sequential estimation algorithm (Alg. 2).

Algorithm 5: Preconditioned replacements for steps (b) and (c) of the sequential estimation algorithm (Alg. 2).

(b.1) Approximate the bridging density $\rho_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ by the TT-based approximate density $\hat{\rho}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ defined in (35).

- (b.2) Integrate $\hat{\rho}_t(\mathbf{x}_t, \boldsymbol{\theta}, \mathbf{x}_{t-1})$ from the right variable $x_{t-1,m}$ to the left variable $x_{t,1}$ using Proposition 2 to define lower-triangular preconditioning map \mathcal{T}_t^l , the lower conditional map $\mathcal{T}_{t,t-1}^l(\mathbf{x}_{t-1}|\mathbf{x}_t, \boldsymbol{\theta})$, and the marginal density $\hat{\rho}_t(\mathbf{x}_t, \boldsymbol{\theta})$.
- (b.3) Approximate the pushforward density $q_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})$ by the TT-based approximation $\hat{\nu}_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})$ defined in (32).
- (c.1) Integrate the last block of $\hat{\nu}_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta, \mathbf{u}_{t-1})$ from the right variable $u_{t-1,m}$ to the left variable $u_{t-1,1}$ using Proposition 2 to define the marginal density $\hat{\nu}_{\sharp,t}(\mathbf{u}_t, \mathbf{u}_\theta)$ and the lower conditional map $\mathcal{S}_{t,t-1}^l(\mathbf{u}_{t-1}|\mathbf{u}_t, \mathbf{u}_\theta)$.
- (c.2) Using the lower-triangular transformation $(\mathbf{u}_t, \mathbf{u}_\theta) = (\mathcal{T}_{t,t}^l(\mathbf{x}_t), \mathcal{T}_{t,\theta}^l(\boldsymbol{\theta}|\mathbf{x}_t))$, we have the marginal density

$$\hat{\pi}(\mathbf{x}_t, \boldsymbol{\theta}|\mathbf{y}_{1:t}) = \frac{\hat{\nu}_{\sharp,t}(\mathcal{T}_{t,t}^l(\mathbf{x}_t), \mathcal{T}_{t,\theta}^l(\boldsymbol{\theta}|\mathbf{x}_t))}{\eta(\mathcal{T}_{t,t}^l(\mathbf{x}_t), \mathcal{T}_{t,\theta}^l(\boldsymbol{\theta}|\mathbf{x}_t))} \hat{\rho}_t(\mathbf{x}_t, \boldsymbol{\theta}),$$

which will be used in the next iteration of the sequential estimation.

As a byproduct, steps (b.2) and (c.1) of Alg. 5 also define the last block of the composite map $\mathcal{S}_t^l \circ \mathcal{T}_t^l$, which can be expressed as

$$\mathcal{F}_{t,t-1}^l(\mathbf{x}_{t-1}|\mathbf{x}_t, \boldsymbol{\theta}) := \mathcal{S}_{t,t-1}^l(\mathcal{T}_{t,t-1}^l(\mathbf{x}_{t-1}|\mathbf{x}_t, \boldsymbol{\theta})|\mathcal{T}_{t,t}^l(\mathbf{x}_t), \mathcal{T}_{t,\theta}^l(\boldsymbol{\theta}|\mathbf{x}_t)).$$

This precisely gives the lower conditional map for the path estimation (Alg. 4).

6. Numerical results

We provide several numerical examples to demonstrate the efficiency of our TT-based methods. These include a linear Kalman filter controlled by unknown parameters, in which the posterior density of the parameters has an analytical form; a stochastic volatility model (cf. Example 1) commonly used in the literature for benchmarking sequential algorithms; a high-dimensional compartmental susceptible-infectious-removed model following the Austrian state adjacency map; and a dynamical system modelling the interaction of a predator-prey system. We also compare our method to SMC² when it is applicable. For all the numerical examples reported here, we compute the 25%, 50%, 75% quantiles of the effective sample sizes using 40 repeated experiments from smoothing samples obtained in Alg. 4.

6.1 Linear Kalman filter with unknown parameters

Setup. Our first example considers a linear Kalman filter with unknown parameters, in which the state process and the observation process take the form:

$$\begin{cases} \mathbf{X}_t - \mu = b(\mathbf{X}_{t-1} - \mu) + a\varepsilon_t^{(x)} \\ \mathbf{Y}_t = \mathbf{C}\mathbf{X}_t + d\varepsilon_t^{(y)} \end{cases}, \quad (36)$$

where $\varepsilon_t^{(x)}$ and $\varepsilon_t^{(y)}$ follow independent Gaussian distributions $\mathcal{N}(0, \mathbf{I}_m)$ and $\mathcal{N}(0, \mathbf{I}_n)$, and $\mathbf{C} \in \mathbb{R}^{n \times m}$ is the observation matrix. Here \mathbf{I}_m represents an $m \times m$ identity matrix. The

prior density of the initial state \mathbf{X}_0 is given by $p(\mathbf{x}_0|\mu) := \mathcal{N}(\mathbf{x}_0; \mu\mathbf{1}, \mathbf{I}_m)$, where $\mathbf{1}$ is the m -dimensional vector filled with ones. The parameters controlling the model are collected as $\boldsymbol{\theta} = (\mu, a, b, d)$.

Conditioned on the parameters, the joint state filtering density $p(\mathbf{x}_{0:t}|\boldsymbol{\theta}, \mathbf{y}_{1:t})$ follows a multivariable Gaussian distribution, which can be derived from the classical Kalman filter. Integrating over the states, we obtain an explicit form of the posterior parameter density $p(\boldsymbol{\theta}|\mathbf{y}_{1:t}) = \int p(\mathbf{x}_{0:t}, \boldsymbol{\theta}|\mathbf{y}_{1:t})d\mathbf{x}_{0:t}$. Although the posterior parameter density cannot be directly sampled, we can evaluate the density pointwisely for any given parameter $\boldsymbol{\theta}$, and thus we can use the model in (36) to benchmark the performance of various estimation algorithms—for example, this enables us to estimate the Hellinger distance between the exact posterior parameter density $p(\boldsymbol{\theta}|\mathbf{y}_{1:t})$ and its approximation $\hat{p}(\boldsymbol{\theta}|\mathbf{y}_{1:t})$ obtained by Alg. 2. For the sake of completeness, we include the derivation of the posterior parameter density $p(\boldsymbol{\theta}|\mathbf{y}_{1:t})$ in Appendix B.

In our numerical experiment, we take m and n to be three and let the total number of time steps T be 50. Additionally, we set $\mu = 0$ and set the observation matrix \mathbf{C} to be a prescribed random value. We impose the condition $a^2 + b^2 = 1$ such that the state process \mathbf{X}_t is stationary. Thus, the parameters to be estimated are reduced to $\boldsymbol{\theta} = (a, d)$. We impose a uniform prior density $p(\boldsymbol{\theta}) := \text{uniform}(\boldsymbol{\theta}; [0.4, 1]^2)$ on the parameters, and generate the synthetic data \mathbf{Y}_t using $a = 0.8$ and $d = 0.5$ in the numerical experiments. When building TT approximations, we transform the parameters to an unbounded domain using the inverse distribution function of the standard Gaussian distribution to facilitate the linear preconditioning techniques in Sec. 5.2.

Comparison of Alg. 1 and 2. We start with the comparison between Alg. 1 and 2. For both algorithms, we use 5 alternating least square (ALS) iterations to construct TT decompositions. Note that using two ALS iterations is sufficient in most of the situations presented here. We use a rather large number of ALS iterations to eliminate possible error sources in our benchmarks. The TT decomposition used here employs a piecewise Lagrange basis function defined by four subintervals and polynomials with order eight. We consider maximum TT ranks $r \in \{10, 20, 30\}$ in this comparison.

Following the discussion of Section 4, we use the relative L^1 error of posterior parameter densities, $\|\hat{\pi}(\boldsymbol{\theta}|\mathbf{y}_{1:t}) - \pi(\boldsymbol{\theta}|\mathbf{y}_{1:t})\|_{L^1} / \|\pi(\boldsymbol{\theta}|\mathbf{y}_{1:t})\|_{L^1}$, to benchmark the accuracy of these TT-based algorithms. Note that the denominator $\|\pi(\boldsymbol{\theta}|\mathbf{y}_{1:t})\|_{L^1}$ is also the normalizing constant of the posterior density. Fig. 3 shows the change of the relative L^1 error over time. For both algorithms, we observe that the error accumulations behave similarly, and relative errors reduce with increasing TT ranks. For all TT ranks, Alg. 2 is more accurate than Alg. 1. In the other numerical experiments, we only demonstrate the performance of Alg. 2 and its accompanying sampling algorithms.

Demonstration of Alg. 2 and debiasing. Since the non-negativity-preserving Alg. 2 is the workhorse of this paper, here we thoroughly demonstrate its accuracy and the sampling performance of the accompanying path estimation algorithm (Alg. 4) with various algorithmic settings. In all experiments, we use five ALS iterations to construct TT decompositions and piecewise basis functions defined on four subintervals with Lagrange polynomials. We denote the number of degrees of freedom of the basis functions by ℓ .

In the first set of numerical experiments, we vary the number of degrees of freedom of the piecewise basis functions ℓ and the maximum TT rank r to investigate their impacts

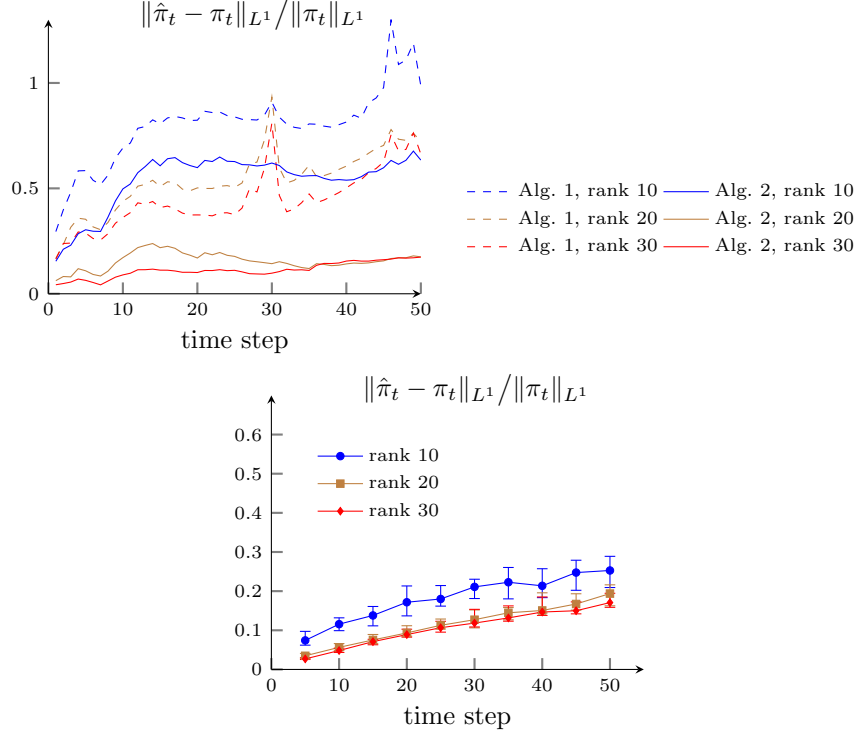


Figure 3: Linear Kalman filter with unknown parameters. The relative L^1 error of approximations of posterior parameter densities built by Alg. 1 and 2.

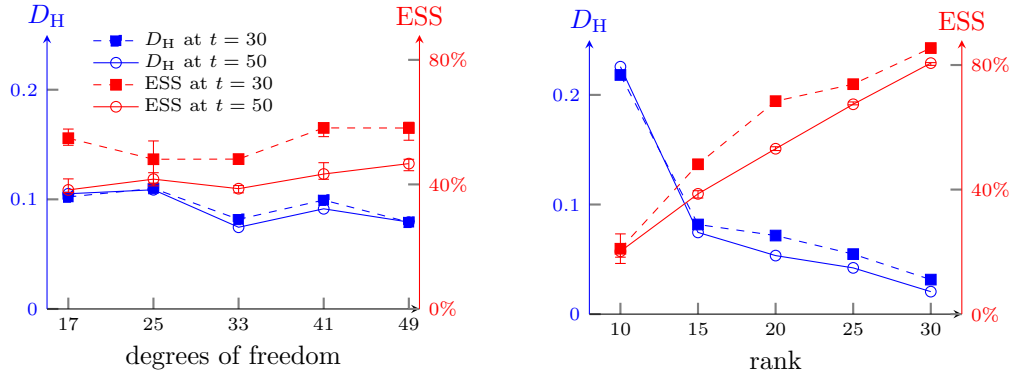


Figure 4: Linear Kalman filter with unknown parameters. The Hellinger distances (blue lines) between the theoretical posterior parameter density and its TT approximation, and the ESS (red lines) for the joint posterior of the states and the parameters at time $t \in \{30, 50\}$. Left: changing the number of degrees of freedom of the basis functions ℓ with a fixed maximum TT rank $r = 15$. Right: changing the maximum TT rank r with a fixed number of degrees of freedom of the basis functions $\ell = 33$.

on the accuracy of TT-based methods. We first fix the maximum TT rank r to be 15 and vary $\ell = \{17, 25, 33, 41, 49\}$ by fixing the number of subintervals to be four and increasing the order of the Lagrange polynomials. For each ℓ , we run Alg. 2 for t up to 50 and apply

the path estimation (Alg. 4) using $N = 1000$ sample paths at $t \in \{30, 50\}$. For each of the approximate posterior parameter densities, we report the Hellinger distance and the effective sample size (ESS)—for the joint density $p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})$ —at $t \in \{30, 50\}$ in the left plot of Fig. 4. We observe that increasing the number of degrees of freedom of the basis functions only marginally improves the accuracy. Then we fix the number of degrees of freedom ℓ to be 33 with order-eight Lagrange polynomials on four subintervals, and vary the maximum TT rank $r \in \{10, 15, 20, 25, 30\}$. The resulting Hellinger distances and ESSs are shown in the right plot of Fig. 4. In this case, we observe a significant error reduction with increasing TT ranks.




Figure 5: Linear Kalman filter with unknown parameters. Left: the Hellinger distance between the theoretical posterior parameter density $p(\boldsymbol{\theta} | \mathbf{y}_{1:t})$ and its TT approximations versus time. Middle: the ESS of the path estimation (Alg. 4) for sampling the joint posterior $p(\boldsymbol{\theta}, \mathbf{x}_{0:t} | \mathbf{y}_{1:t})$ versus time. Right: the ESS for sampling the posterior parameter density $p(\boldsymbol{\theta} | \mathbf{y}_{1:t})$ versus time.

In the second set of numerical experiments, we demonstrate the impact of preconditioning techniques. Here we consider two approximation ansatzes with $(\ell, r) = (33, 15)$ and $(\ell, r) = (33, 30)$ using linear preconditioning (cf. Section 5.2) and another ansatz $(\ell, r) = (33, 15)$ using nonlinear preconditioning (cf. Section 5.3). For nonlinear preconditioning, we use a standard multivariate Gaussian for the reference density η and $\beta_\pi = \beta_f = \beta_g = 0.4$ to construct the bridging density. For each $t \in \{5, 10, \dots, 50\}$, we report the Hellinger distance, the ESS of the path estimation (Alg. 4) for the joint density, and the ESS of the approximate posterior parameter density. The results are shown in Fig. 5. In the linear preconditioning case, we observe that the approximation ansatz $(\ell, r) = (33, 30)$ outperforms the approximation ansatz $(\ell, r) = (33, 15)$ at all time steps, which is expected. We also observe that the performance of the nonlinear preconditioning technique using the approximation ansatz $(\ell, r) = (33, 15)$ is comparable with that of linear preconditioning with a higher-rank approximation ansatz $(\ell, r) = (33, 30)$. For both $(\ell, r) = (33, 15)$ using nonlinear preconditioning and $(\ell, r) = (33, 30)$ using linear preconditioning, the Hellinger distance remains below 0.05 for all time steps, while the ESSs for the joint densities and the marginal parameter densities are respectively around 80% and 98% after 50 time steps. These results provide strong evidence of the reliability of the TT-based sequential estimation methods.

Comparison with SMC². We compare Alg. 2 with $(\ell, r) = (33, 30)$ and $(\ell, r) = (33, 15)$ to the SMC² method. We consider two setups for the SMC² method. One uses a particle size $N_\theta = 500$ for the parameters and another one uses an increased particle size

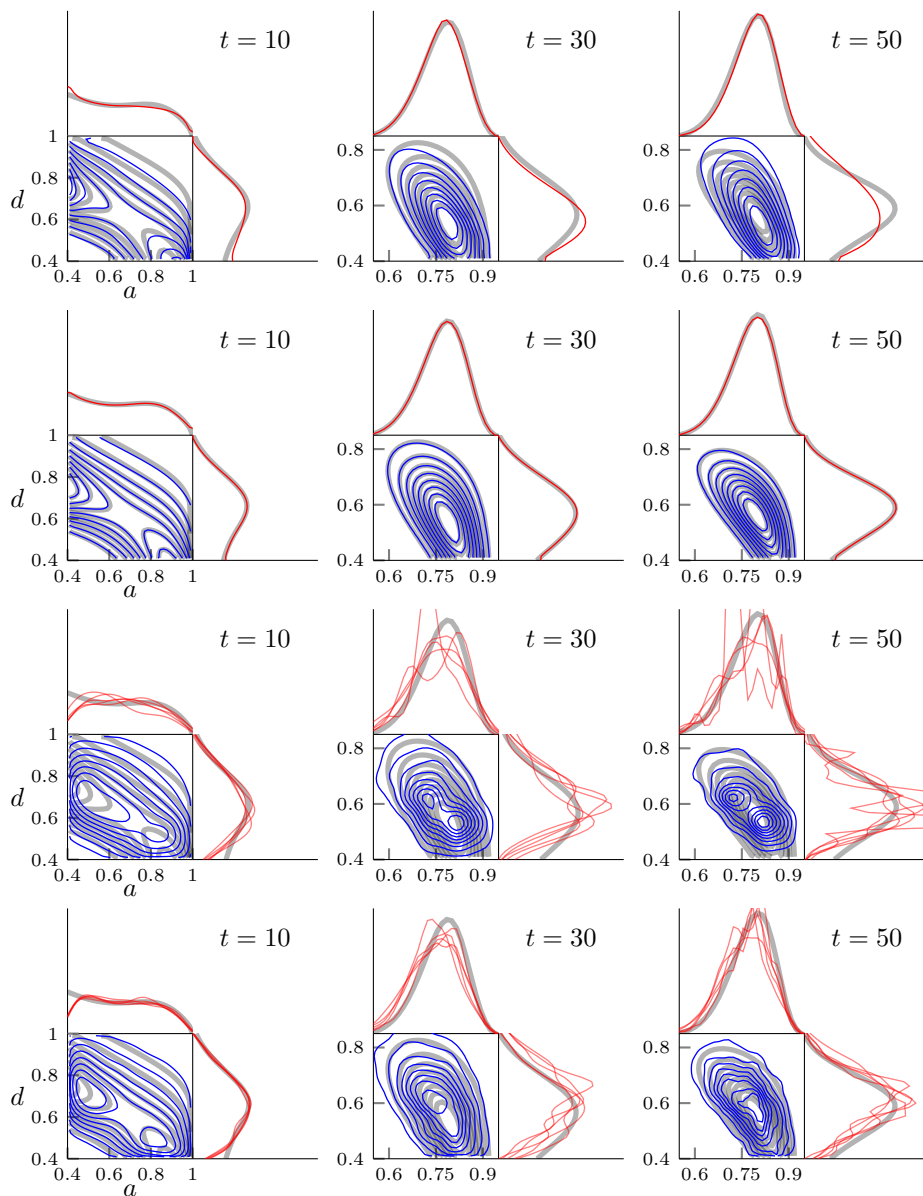


Figure 6: Linear Kalman filter with unknown parameters. The contours of posterior densities of $\theta = (a, d)$ at different time. From the top row to the bottom row, blue lines are results obtained by Alg. 2 with $(\ell, r) = (33, 15)$, Alg. 2 with $(\ell, r) = (33, 30)$, SMC² with 500 particles, and SMC² with 5000 particles, respectively. The marginal densities are shown by red curves on each side. The analytical solutions are plotted as thick gray curves at the background. For SMC², the marginal densities are results obtained from 5 batches, whereas the contours are results obtained by putting all batches together.

$N_\theta = 5000$. In all numerical experiments, the particle size for the states is initially set to $N_x = 100$ and adaptively increased so that the acceptance rate of the parameter particles remains a relatively high level (Chopin et al., 2013). In our setup, the computational cost

of running SMC² with $(N_\theta, N_x) = (500, 100)$ is about the same as running Alg. 2 with $(\ell, r) = (33, 30)$. Both take roughly 10 minutes using a 3.20 GHz Intel i7-8700 CPU. In Fig. 6, we present contours of the posterior parameter densities and their one-dimensional marginals at $t \in \{10, 30, 50\}$ estimated using various algorithms. For SMC², we follow the setup of Chopin et al. (2013) to run five independent batches of particles, and then apply the kernel density estimation (KDE) method (Silverman, 1986; Peter D, 1985) to estimate the joint density using the union of the five batches and the one-dimensional parameter marginal density of each batch. Both TT-based setups perform better than the SMC²-based setups. The approximate posterior parameter densities from TT with $(\ell, r) = (33, 30)$ accurately capture the true densities at all time steps, while those from SMC² do not have a consistent result across different batches, especially with $(N_\theta, N_x) = (500, 100)$.

6.2 Stochastic volatility

Our second numerical example uses the stochastic volatility model defined in Example 1. We test Alg. 2 and Alg. 4 on synthetic data sets and a real-world data set taken from the S&P 500 index.

Synthetic data. We fix $\sigma = 1$ and aim to estimate parameters $\boldsymbol{\theta} = (\gamma, \beta)$. We consider the sequential estimation problems for $T = 1000$ steps. Here $\gamma = 0.6$ and $\beta = 0.4$ are used to generate synthetic data. We set the prior density of the initial state to be $p(\mathbf{x}|\boldsymbol{\theta}) := \mathcal{N}(\mathbf{x}; 0, 1/(1 - \gamma^2))$ and the prior density of the parameters to be $p(\boldsymbol{\theta}) := \text{uniform}(\boldsymbol{\theta}; [0.1, 0.9]^2)$. Since the analytical marginal densities are no longer available in this example, we can only evaluate the 1002-dimensional joint posterior density of the parameters and all the states $(\boldsymbol{\theta}, \mathbf{x}_{0:1000})$, and then estimate the corresponding joint ESS as a measure of the approximation errors.




Figure 7: Stochastic volatility model with synthetic data. The ESS for the joint posterior density versus time, computed using different maximum TT ranks $r \in \{5, 10, 20\}$.

Similarly to the linear Kalman filter case, we use the piecewise Lagrange polynomial with the number of degrees of freedom ℓ to be 33 as the basis function, linear preconditioning, and five ALS iterations for building TT decompositions in Alg. 2. The accompanying path estimation algorithm (Alg. 4) is used to generate weighted samples from the resulting TT-based approximations. We set the maximum TT rank to be $r \in \{5, 10, 20\}$ and show the change of ESS over time in Fig. 7. We observe that with $r = 5$, the TT-based approximation yields a reliable result with ESS above 20% even after 1000 steps, while with $r = 10$ it provides nearly 100% ESS at almost all time steps. Fig. 8 presents the trajectories of the states estimated from the approximate posterior state path density $\hat{p}(\mathbf{x}_{0:1000}|\mathbf{y}_{1:1000})$ using rank $r = 10$. We observe the true states are well followed by the trajectories estimated




Figure 8: Stochastic volatility model with synthetic data. The thick brown line represents the path of the true state and crosses are the observations. Estimated using the approximate posterior state path density $\hat{p}(\mathbf{x}_{0:1000}|\mathbf{y}_{1:1000})$ defined in (25) with a maximum TT rank $r = 10$, the black line and the shaded region represent the median path and the credible interval bounded between 5% and 95% percentiles, respectively.




Figure 9: Stochastic volatility model with synthetic data. The brown line and the shaded region in the horizontal (t, γ) -plane show the evolution of the median and the credible interval of the posterior estimates of γ bounded between 5% and 95% percentiles. The density profiles in the vertical axis show the approximate posterior parameter density $\hat{p}(\gamma|\mathbf{y}_{1:t})$ at different time.

from the approximations. In Fig. 9, we also show the approximate filtering densities of the parameter γ from the TT-based approximations, that is, $\hat{p}(\gamma|\mathbf{y}_{1:t})$ over time. We observe that the densities get concentrated around the true value $\gamma = 0.6$ (for generating the data) with an increasing amount of data observed over time.

Next, we compare the TT-based method with $r = 10$ to SMC² with 5 independent batches, each batch with $N_\theta = 1000$ parameter particles and initially $N_x = 100$ state particles. Fig. 10 shows the contours and the marginals of the posterior parameter densities at $t \in \{400, 700, 1000\}$ obtained by the TT-based method and SMC². Both methods yield densities that concentrate around the true values of the parameters used to generate the data, that is, $(\gamma, \beta) = (0.6, 0.4)$. However, the computation time of the TT-based method is about one hour using a 3.20 GHz Intel i7-8700 CPU, while that for one batch of SMC² method is over eight hours.

S&P 500 index. We then apply our algorithms to the daily returns of the S&P 500 index. The historical data of the stock prices are obtained from Yahoo Finance¹. The data set contains 1009 observations from 31st of December 2019 to 29th of December 2023, excluding weekends and holidays. We compute the continuously compounded daily returns $\{\mathbf{y}_t\}_{t=1}^T$ from the stock prices, where $T = 1008$.

1. <https://au.finance.yahoo.com/>

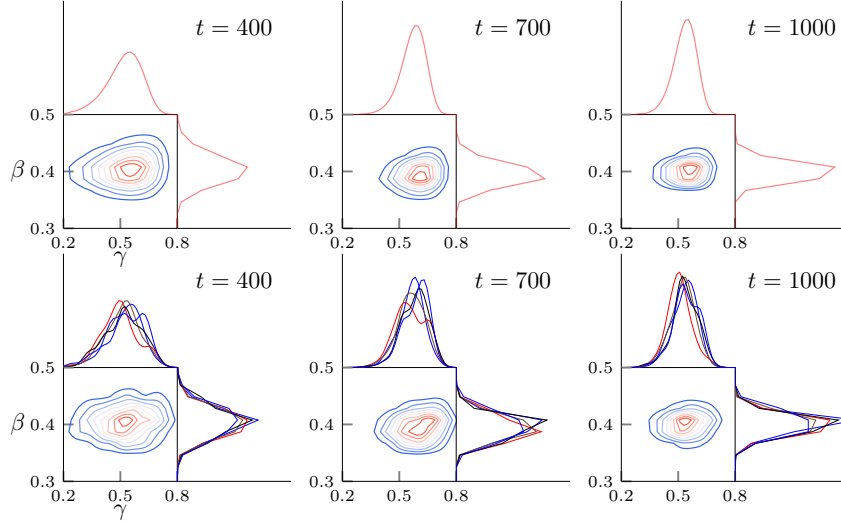


Figure 10: Stochastic volatility model with synthetic data. The contours of posterior densities of $\theta = (\gamma, \beta)$ at different time. The top and bottom rows show results obtained by Alg. 2 and SMC² respectively. The marginal densities are shown on each side. For SMC², the marginal densities shown are results obtained from 5 batches with 1000 particles in each batch, whereas the contours are results obtained by putting all batches together.

We aim to estimate the states $\{\mathbf{X}_t\}_{t=0}^T$ and all the three parameters $\Theta = (\gamma, \sigma, \beta)$ from the observed daily returns $\{\mathbf{y}_t\}_{t=1}^T$ of the S&P 500 index. Following the common practice in econometrics (Zhang and Maxwell, 2008; Yu et al., 2006), we set the priors for Θ and \mathbf{X}_0 as:

- $(\gamma + 1)/2 \sim \text{Beta}(\omega_1, \omega_2)$, with $\omega_1 = 20, \omega_2 = 1.5$.
- $\sigma^2 \sim \text{IG}(\zeta/2, S_\sigma/2)$, with $\zeta = 2, S_\sigma = 0.01$, where IG stands for the inverse Gamma distribution.
- $\log(\beta)|\sigma \sim \mathcal{N}(\beta_0, \sigma^2/q_0)$, with $\beta_0 = 0, q_0 = 0.8$.
- $\mathbf{X}_0|\gamma, \sigma \sim \mathcal{N}(0, \sigma^2/(1 - \gamma^2))$.

When building TT approximations, we transform the parameters and the states in the stochastic volatility model to

$$\begin{aligned} \Theta' &= (\Phi^{-1}(\gamma), \log(\sigma), \log(\beta)/\sigma), \\ \mathbf{X}'_t &= \mathbf{X}_t/\sigma, \end{aligned}$$

where Φ denotes the distribution function of the standard Gaussian random variable. The transformed parameters and states take values in an unbounded domain, so that the linear preconditioning techniques developed in Sec. 5.2 can be applied. After obtaining the approximate densities of (Θ', \mathbf{X}'_t) using our TT-based methods, we transform them back to the original variables (Θ, \mathbf{X}_t) to present the estimation results.

We set the maximum TT rank to be $r \in \{5, 10, 20\}$ and maintain the other TT configurations used in the synthetic case. Fig. 11 shows the change of ESS over time. The




Figure 11: Stochastic volatility model with the S&P500 returns. The ESS for the joint posterior density at different time, computed using maximum TT ranks $r \in \{5, 10, 20\}$.




Figure 12: Stochastic volatility model with the S&P500 returns. The observed returns (top) and the estimated trajectories of squared volatilities (bottom). In the bottom plot, the black line and the shaded region represent the median path and the credible interval bounded between 5% and 95% percentiles, respectively. These estimation results are obtained using the approximate posterior state path density $\hat{p}(\mathbf{x}_{0:1008}|\mathbf{y}_{1:1008})$ defined in (25) with a maximum TT rank $r = 20$.




Figure 13: Stochastic volatility model with the S&P500 returns. The contours of posterior densities of $\theta = (\gamma, \sigma, \beta)$ at $T = 1008$, estimated by Alg. 2. The marginal densities are shown on each side.

ESS clearly increases with the TT rank, reaching around 70% after 1008 steps for $r = 20$.

Fig. 12 shows the observed returns of the S&P 500 index during the selected period and the trajectories of the squared volatilities estimated from the approximate posterior state path density $\hat{p}(\mathbf{x}_{0:1008}|\mathbf{y}_{1:1008})$ using rank $r = 20$. We observe that the highly volatile returns around $t = 300$ are well captured by the peak of the estimated volatilities, and the increasing variability of the returns after $t = 800$ is also captured by the increase in the estimated volatilities in the corresponding timeframe. In Fig. 13, we also present the approximate posterior densities of the parameters $\theta = \{\gamma, \sigma, \beta\}$ from the TT-based approximations $\hat{p}(\theta|\mathbf{y}_{1:1008})$. We observe a highly nonlinear pattern between β and the other two parameters. These results demonstrate that our TT-based sequential estimation methods are effective for real-world data in this case.

6.3 Susceptible-infectious-removed model

We then apply our TT-based methods to a high-dimensional compartmental susceptible-infectious-removed (SIR) model to demonstrate its dimension scalability. The SIR model describes the numbers of susceptible, infectious and removed individuals across $J \in \mathbb{N}$ spatially dependent demographic compartments, denoted by $S_j(t)$, $I_j(t)$ and $R_j(t)$, respectively. The interaction between them are modelled by the following coupled ODEs:

$$\begin{cases} \frac{dS_j}{dt} = -\kappa_j S_j I_j + \frac{1}{2} \sum_{i \in \mathcal{I}_j} (S_i - S_j) \\ \frac{dI_j}{dt} = \kappa_j S_j I_j - \nu_j I_j + \frac{1}{2} \sum_{i \in \mathcal{I}_j} (I_i - I_j) \\ \frac{dR_j}{dt} = \nu_j I_j + \frac{1}{2} \sum_{i \in \mathcal{I}_j} (R_i - R_j) \end{cases}, \quad (37)$$

where κ_j and ν_j are the parameters representing the infection and recovery rates, and \mathcal{I}_j is the index set containing all neighbors of the j -th compartment. We fix the parameters $\kappa_j = 0.1$ and $\nu_j = 18$ for $j = 1, 2, \dots, J$, and focus on the inference for S_j , I_j and R_j .

In the SIR model, it is sufficient to infer any two of the three random processes in each compartment, and the third process is known given the other two. Here we infer the processes S_j and I_j , and set the state to be

$$\mathbf{x} = (S_1, I_1, \dots, S_J, I_J) \in \mathbb{R}^{2J}.$$

Then, the dynamic for the state follows

$$\frac{d\mathbf{x}}{dt} = G(\mathbf{x}),$$

where $G(\mathbf{x})$ is defined by the right-hand side of the first two equations in (37). To formulate the state-space model, we discretize the dynamical system at discrete time points $t = k\Delta t$, where $\Delta t = 0.02$, and perturb the states using Gaussian noises $\varepsilon_k^{(x)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{2J})$. Thus, conditioned on the state \mathbf{X}_{k-1} at time $(k-1)\Delta t$, the state \mathbf{X}_k at time $k\Delta t$ is given by

$$\mathbf{X}_k = \mathbf{X}_{k-1} + \int_{(k-1)\Delta t}^{k\Delta t} G(\mathbf{x}(t)) dt + \varepsilon_k^{(x)}.$$

We use the explicit fourth-order Runge-Kutta formula with constant time step sizes $\Delta t = 0.005$ to numerically solve the above time integration problem. For the observation process, we assume only $I_j(t)$ is observable at each time step, and the observations are perturbed by Gaussian noises $\varepsilon_k^{(y)} \sim \mathcal{N}(\mathbf{0}, 100 \mathbf{I}_J)$. This defines the observation process

$$\mathbf{Y}_{k,j} = \mathbf{X}_{k,2j} + \varepsilon_{k,j}^{(y)}, \quad k = 1, 2, \dots, T, \quad j = 1, 2, \dots, J.$$




Figure 14: Compartment connectivity graph of the Austrian states (Cui et al., 2024).

We consider the same realistic setting as in Cui et al. (2024), with $J = 9$ compartments following the Austrian state adjacency map shown in Fig 14. We generate the synthetic data for $T = 20$ time steps, with fixed inhomogeneous initial states $S_j(0) = 485 + j$ and $I_j(0) = 15 - j$, for $j = 1, 2, \dots, 9$. Subsequently, we impose a Gaussian prior on the initial states $p(\mathbf{x}_0) := \mathcal{N}(\mathbf{x}_0; \boldsymbol{\mu}_0, \mathbf{I}_{18})$ where $\boldsymbol{\mu}_0 = (S_1(0), I_1(0), \dots, S_9(0), I_9(0))$. We aim to estimate the 18-dimensional states from the partial observations, which poses a challenging high-dimensional inference problem.




Figure 15: SIR model. The ESS for the joint posterior density versus time, computed using maximum TT ranks $r \in \{10, 20, 40\}$.

Similarly to the previous experiments, we employ Alg. 2 with linear preconditioning, the piecewise Lagrange polynomial with the number of degrees of freedom ℓ to be 33, and five ALS iterations. Nevertheless, we increase the maximum TT rank to $r \in \{10, 20, 40\}$ to accommodate the larger state dimensions in this experiment. Fig. 15 shows the change of ESS over time. We observe that the large state dimension in this example requires TT




Figure 16: SIR model. The trajectories of the number of susceptible and infectious individuals in Vorarlberg (left) and Styria (right). The thick brown curves are the true state trajectories. The black curves represent the median posterior state path obtained by the approximate density $\hat{p}(\mathbf{x}_{0:20}|\mathbf{y}_{1:20})$, and the shaded region is the credible interval bounded between 5% and 95% percentiles of the approximate smoothing densities.

ranks above $r = 20$ to achieve accurate results. In Fig. 16, we present the trajectories of the first and middle dimensions of the state (Vorarlberg and Styria), estimated from the approximate posterior state path density $\hat{p}(\mathbf{x}_{0:20}|\mathbf{y}_{1:20})$ using rank $r = 40$. We observe that the true dynamics are closely followed by the estimated trajectories, despite S_j being unobservable and the observations for I_j being far from the true dynamics. These results demonstrate a strong dimension scalability of our TT-based sequential estimation methods in this case.

6.4 Predator-prey model

Finally, we consider the predator-prey model, which is a time-invariant dynamical system in the form of

$$\begin{cases} \frac{dP}{dt} = rP\left(1 - \frac{P}{K}\right) - s\left(\frac{PQ}{a + P}\right) \\ \frac{dQ}{dt} = u\left(\frac{PQ}{a + P} - vQ\right) \end{cases}, \quad (38)$$

where P and Q denote the prey population and the predator population, respectively. There are six parameters, $\boldsymbol{\theta} = (r, K, a, s, u, v)$, that control the behavior of the system. To formulate the state-space model, we represent P and Q as $\mathbf{x} = (P, Q)$. This way, the dynamical system in (38) can be equivalently expressed as

$$\frac{d\mathbf{x}}{dt} = G(\mathbf{x}; \boldsymbol{\theta}),$$

where $G(\mathbf{x}; \boldsymbol{\theta})$ is defined by the right-hand side of (38). Similarly to the previous SIR model, we discretize the dynamical system at discrete time points $t = k\Delta t$, where $\Delta t = 2$, and

perturb the states using Gaussian noises $\varepsilon_k^{(x)} \sim \mathcal{N}(\mathbf{0}, 4\mathbf{I}_2)$. This way, the state transition process is given by

$$\mathbf{X}_k = \mathbf{X}_{k-1} + \int_{(k-1)\Delta t}^{k\Delta t} G(\mathbf{x}(t); \boldsymbol{\theta}) dt + \varepsilon_k^{(x)}.$$

We still use the explicit fourth-order Runge-Kutta formula with constant time step sizes $\Delta t = 0.1$ to numerically solve the above time integration problem. We observe \mathbf{x} at the same sequence of discrete time points, and the observations are perturbed by Gaussian noises $\varepsilon_k^{(y)} \sim \mathcal{N}(\mathbf{0}, 4\mathbf{I}_2)$. This defines the observation process

$$\mathbf{Y}_k = \mathbf{X}_k + \varepsilon_k^{(y)}.$$

We generate synthetic observations using $\boldsymbol{\theta} = (0.6, 114, 25, 0.3, 0.5, 0.5)$ and $\mathbf{x}_0 = [50, 5]^\top$ up to the terminal time $T = 20$. We impose a uniform prior on the parameters, $p(\boldsymbol{\theta}) := \text{uniform}(\mathbf{x}; [\mathbf{a}, \mathbf{b}])$, where $[\mathbf{a}, \mathbf{b}] \equiv [a_1, b_1] \times \cdots \times [a_6, b_6]$ is a hypercube defined by $\mathbf{a} = (0.1, 110, 20, 0.1, 0, 0)$ and $\mathbf{b} = (1.1, 130, 30, 1.1, 1, 1)$. Furthermore, we impose a Gaussian prior on the initial states $p(\mathbf{x}_0) := \mathcal{N}(\mathbf{x}_0; \boldsymbol{\mu}_0, \mathbf{I}_2)$ where $\boldsymbol{\mu}_0 = (50, 5)$.

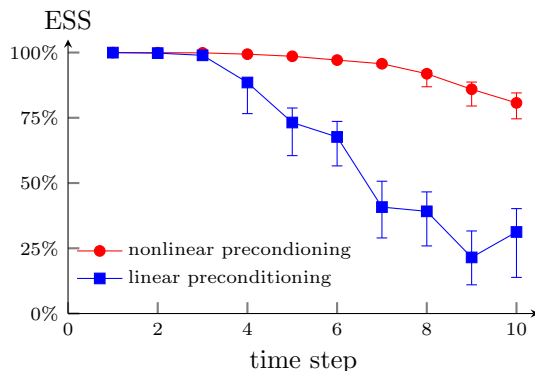


Figure 17: Predator-prey model. The ESS for the joint posterior density versus time, computed using different preconditioning techniques.




Figure 18: Predator-prey model. The thick brown curves are the true trajectories, discontinuity indicates the state noises at discrete time steps, and blue crosses are the observations. The black curves represent the median posterior state path obtained by the approximate density $\hat{p}(\mathbf{x}_{0:20}|\mathbf{y}_{1:20})$, and the shaded region is the credible interval bounded between 5% and 95% percentiles of the approximate smoothing densities.

We first employ Alg. 2 with linear preconditioning, the piecewise Lagrange polynomial with the number of degrees of freedom ℓ to be 33, a maximum TT rank $r = 20$, and five ALS iterations. The accompanying path estimation algorithm (Alg. 4) is used to generate weighted samples from the resulting TT-based approximations. In Fig. 17, the ESSs for the joint posterior density of the parameters and all the states are represented by square markers in blue. In this case, the ESS decreases rapidly. Instead of brutally increasing the rank to enhance the approximation power of TT, we keep the same TT rank and apply the nonlinear preconditioning technique. Here we use the bridging density (34) with $\beta_\pi = \beta_f = \beta_g = 0.4$ and a standard multivariate Gaussian reference density. With nonlinear preconditioning, the ESSs are shown by red circles in Fig. 17. We observe that nonlinear preconditioning significantly increases the ESS, where it remains at about 40% after 20 steps. In Fig. 18, we also present the trajectories of the states estimated from the approximate posterior state path density $\hat{p}(\mathbf{x}_{0:20}|\mathbf{y}_{1:20})$ with nonlinear preconditioning and the true state trajectory. The curves between the two adjacent states \mathbf{x}_{k-1} and \mathbf{x}_k are the solution of the dynamical system (38) with initial condition $\mathbf{x} = \mathbf{x}_{k-1}$. We observe that the true trajectories are well followed by the approximations.

7. Conclusion and further extensions

We present new TT-based methods to sequentially learn the posterior distributions of states and parameters of state-space models. Our main innovation is the recursive approximation of posterior densities using TT rather than relying on (weighted) particle representations. Using the square-root approximation technique (cf. Section 3), we derive conditional KR rearrangements that map reference random variables to the approximated posterior random variables. As a result, this defines the particle filtering, path estimation and particle smoothing algorithms accompanying the recursive TT-based approximations.

This work opens the door to many future research directions. For example, one can integrate our TT-based methods into SMC methods (Chopin et al., 2013; Crisan and Miguez, 2018), in which TT-based methods can potentially provide more efficient particle filters and MCMC proposal kernels for SMC. The KR rearrangements offered by our methods also open the door to integrating other structured particle sets—for instance, quasi Monte Carlo points, see Dick et al. (2013) and references therein—into sequential learning algorithms to improve the rate of convergence of estimators. This integration can significantly complement the Hilbert space-filling curve technique used by the sequential quasi Monte Carlo method (Gerber and Chopin, 2015).

Further research is needed to understand the rank structure of TT decompositions in state-space models—particularly on the impact of the Markov property—which may provide *a posteriori* error bounds to certify the algorithms presented here. The function approximation perspective of the work presented here is not limited to the TT decomposition. Other function approximation methods based on separable bases—for example, Gaussian processes and radial basis functions (Wendland, 2004; Williams and Rasmussen, 2006) with product-form kernels, multivariate wavelet and Fourier bases (Daubechies, 1992; de Boor et al., 1993; DeVore and Popov, 1988; Mallat, 1989), and sparse grids and spectral polynomials (Bungartz and Griebel, 2004; Shen et al., 2011; Xiu and Karniadakis, 2002)—can also be used to approximate posterior densities and compute their marginal densities. See

Cui et al. (2023) for the preliminary investigations in stationary inference problems. This offers a pathway to designing problem-specific function approximation methods for sequential learning. For instance, state transition processes governed by certain partial differential equations may yield a hyperbolic-cross structure (Dũng et al., 2018) suitable for sparse spectral polynomials.

Dimension scalability is a limitation of the TT decomposition and general transport-map methods. Although the complexity of building the TT decomposition can be dimension independent for functions equipped with rapid decaying weights (Griebel and Harbrecht, 2023)—which is the case for many high- or infinite-dimensional problems (Stuart, 2010)—it is still computationally prohibitive to operate with the apparent state and parameter dimensions for very high-dimensional systems such as weather forecasts. A possibility of reducing the number of variables in the approximation problem is to use the conditional Gaussian structure of certain state-space models (Chen et al., 2022, 2018; Chen and Majda, 2018) to marginalize high-dimensional conditional Gaussian variables analytically. Then, one only needs to apply TT-based approximations to the rest of the non-Gaussian variables. For a broader range of state-space models, there may exist approximate conditional Gaussian structures and similar structures that allow for analytical marginalization. In ongoing work, we are extending likelihood-informed subspace methods (Cui et al., 2014; Cui and Tong, 2022; Zahm et al., 2022) to state-space models to identify these intrinsic structures.

Acknowledgments

We would like to thank Sergey Dolgov, Rob Scheichl and Olivier Zahm for many fruitful discussions about tensor trains and sparse polynomials. This work was supported in part by the Australian Research Council under the grant DP210103092.

Appendix A. f -divergence of marginal random variables

Commonly used statistical divergences such as the Kullback–Leibler divergence, the total variation distance and the squared Hellinger distance are instances of the f -divergence. Briefly, given a convex function $f(\cdot)$, the f -divergence is defined as

$$D_f(p\|q) = \int f\left(\frac{p(\mathbf{x})}{q(\mathbf{x})}\right)q(\mathbf{x})d\mathbf{x}.$$

The following lemma bounds the f -divergence of marginal random variables by the corresponding f -divergence of joint random variables.

Lemma 13 *Let $p_1(\mathbf{x}_1, \mathbf{x}_2)$ and $p_2(\mathbf{x}_1, \mathbf{x}_2)$ be two joint densities, and $\bar{p}_1(\mathbf{x}_1)$ and $\bar{p}_2(\mathbf{x}_1)$ be the marginal densities of p_1 and p_2 , respectively. The f -divergence of \bar{p}_1 from \bar{p}_2 is bounded from above by that of p_1 from p_2 , that is, $D_f(\bar{p}_1\|\bar{p}_2) \leq D_f(p_1\|p_2)$.*

Proof Since the function f is convex, applying Jensen’s inequality, we have

$$D_f(p_1\|p_2) = \mathbb{E}_{p_2(\mathbf{x}_1, \mathbf{x}_2)} \left[f\left(\frac{p_1(\mathbf{X}_1, \mathbf{X}_2)}{p_2(\mathbf{X}_1, \mathbf{X}_2)}\right) \right]$$

$$\begin{aligned}
 &= \int \left(\int f \left(\frac{p_1(\mathbf{x}_1, \mathbf{x}_2)}{p_2(\mathbf{x}_1, \mathbf{x}_2)} \right) \frac{p_2(\mathbf{x}_1, \mathbf{x}_2)}{\bar{p}_2(\mathbf{x}_1)} d\mathbf{x}_2 \right) \bar{p}_2(\mathbf{x}_1) d\mathbf{x}_1 \\
 &\geq \int f \left(\int \frac{p_1(\mathbf{x}_1, \mathbf{x}_2)}{p_2(\mathbf{x}_1, \mathbf{x}_2)} \frac{p_2(\mathbf{x}_1, \mathbf{x}_2)}{\bar{p}_2(\mathbf{x}_1)} d\mathbf{x}_2 \right) \bar{p}_2(\mathbf{x}_1) d\mathbf{x}_1 \\
 &= \int f \left(\int \frac{p_1(\mathbf{x}_1, \mathbf{x}_2)}{\bar{p}_2(\mathbf{x}_1)} d\mathbf{x}_2 \right) \bar{p}_2(\mathbf{x}_1) d\mathbf{x}_1 \\
 &= \int f \left(\frac{\bar{p}_1(\mathbf{x}_1)}{\bar{p}_2(\mathbf{x}_1)} \right) \bar{p}_2(\mathbf{x}_1) d\mathbf{x}_1 = D_f(\bar{p}_1 \| \bar{p}_2),
 \end{aligned}$$

which concludes the proof. \blacksquare

When the densities are only known up to some constants, the L^2 distance between the square root of the unnormalized marginal densities is also bounded by the L^2 distance between their unnormalized joint densities. This is shown in the following lemma.

Lemma 14 *Let $\pi_1(\mathbf{x}_1, \mathbf{x}_2)$ and $\pi_2(\mathbf{x}_1, \mathbf{x}_2)$ be two unnormalized densities. Then, we denote their corresponding unnormalized marginal densities by $\bar{\pi}_1(\mathbf{x}_1) = \int \pi_1(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_2$ and $\bar{\pi}_2(\mathbf{x}_1) = \int \pi_2(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_2$. The L^2 distance between $\sqrt{\bar{\pi}_1}$ and $\sqrt{\bar{\pi}_2}$ is bounded from above by the L^2 distance between $\sqrt{\pi_1}$ and $\sqrt{\pi_2}$, that is, $\|\sqrt{\bar{\pi}_1} - \sqrt{\bar{\pi}_2}\|_{L^2} \leq \|\sqrt{\pi_1} - \sqrt{\pi_2}\|_{L^2}$.*

Proof The L^2 distance between $\sqrt{\pi_1}$ and $\sqrt{\pi_2}$ can be expressed as

$$\begin{aligned}
 \|\sqrt{\pi_1} - \sqrt{\pi_2}\|_{L^2}^2 &= \int \int (\sqrt{\pi_1(\mathbf{x}_1, \mathbf{x}_2)} - \sqrt{\pi_2(\mathbf{x}_1, \mathbf{x}_2)})^2 d\mathbf{x}_1 d\mathbf{x}_2 \\
 &= \int \int \left(\sqrt{\frac{\pi_1(\mathbf{x}_1, \mathbf{x}_2)}{\pi_2(\mathbf{x}_1, \mathbf{x}_2)}} - 1 \right)^2 \pi_2(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_1 d\mathbf{x}_2 \\
 &= \int \int \left(\sqrt{\frac{\pi_1(\mathbf{x}_1, \mathbf{x}_2)}{\pi_2(\mathbf{x}_1, \mathbf{x}_2)}} - 1 \right)^2 \frac{\pi_2(\mathbf{x}_1, \mathbf{x}_2)}{\bar{\pi}_2(\mathbf{x}_1)} d\mathbf{x}_2 \bar{\pi}_2(\mathbf{x}_1) d\mathbf{x}_1.
 \end{aligned}$$

Since $f(r) = (\sqrt{r} - 1)^2, r \geq 0$ is a convex function, applying Jensen's inequality to the inner integral over \mathbf{x}_2 , the rest of the proof follows the same steps of Lemma 13. \blacksquare

Lemma 14 is useful to establish the bound for the divergence between two unnormalized marginal densities given the divergence bound between their joint densities. However, Lemma 14 implicitly assumes that both unnormalized joint functions (densities) are non-negative by construction, which is generally not true when π_2 is a function approximation of π_1 . The following lemma does not restrict the approximate functions to be non-negative, and gives a similar result under the L^1 distance, which is analogous to the total variation distance for normalized densities.

Lemma 15 *Let $\pi_1(\mathbf{x}_1, \mathbf{x}_2)$ be an unnormalized density and $\bar{\pi}_1(\mathbf{x}_1) = \int \pi_1(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_2$ be its marginal density. Suppose $\pi_1(\mathbf{x}_1, \mathbf{x}_2)$ has an approximation $\pi_2(\mathbf{x}_1, \mathbf{x}_2)$, which does not preserve non-negativity, and let $\bar{\pi}_2(\mathbf{x}_1) = \int \pi_2(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_2$ be the corresponding marginal function. The L^1 distance between $\bar{\pi}_1$ and $\bar{\pi}_2$ is bounded from above by the L^1 distance between π_1 and π_2 , that is, $\|\bar{\pi}_1 - \bar{\pi}_2\|_{L^1} \leq \|\pi_1 - \pi_2\|_{L^1}$.*

Proof The L^1 distance between π_1 and π_2 can be expressed as

$$\begin{aligned} \|\pi_1 - \pi_2\|_{L_1} &= \int \int |\pi_1(\mathbf{x}_1, \mathbf{x}_2) - \pi_2(\mathbf{x}_1, \mathbf{x}_2)| d\mathbf{x}_1 d\mathbf{x}_2 \\ &= \int \int \left| \frac{\pi_2(\mathbf{x}_1, \mathbf{x}_2)}{\pi_1(\mathbf{x}_1, \mathbf{x}_2)} - 1 \right| \pi_1(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_1 d\mathbf{x}_2 \\ &= \int \int \left| \frac{\pi_2(\mathbf{x}_1, \mathbf{x}_2)}{\pi_1(\mathbf{x}_1, \mathbf{x}_2)} - 1 \right| \frac{\pi_1(\mathbf{x}_1, \mathbf{x}_2)}{\bar{\pi}_1(\mathbf{x}_1)} d\mathbf{x}_2 \bar{\pi}_1(\mathbf{x}_1) d\mathbf{x}_1. \end{aligned}$$

Since $f(r) = |r - 1|$ is a convex function, applying Jensen's inequality to the inner integral over \mathbf{x}_2 , the rest of the proof follows the same steps of Lemma 13. \blacksquare

Appendix B. Posterior densities of the linear Kalman filter with unknown parameters

We recall that the linear Kalman filter specifies the state and observation processes as

$$\begin{cases} \mathbf{X}_t - \mu &= b(\mathbf{X}_{t-1} - \mu) + a\varepsilon_t^{(x)} \\ \mathbf{Y}_t &= \mathbf{C}\mathbf{X}_t + d\varepsilon_t^{(y)} \end{cases}, \quad (39)$$

where $\varepsilon_t^{(x)}$ and $\varepsilon_t^{(y)}$ follow respectively independent Gaussian distributions $\mathcal{N}(0, \mathbf{l}_m)$ and $\mathcal{N}(0, \mathbf{l}_n)$, and $\mathbf{C} \in \mathbb{R}^{n \times m}$ is the observation matrix. Here \mathbf{l}_m represents an $m \times m$ identity matrix. The prior density of the initial state \mathbf{X}_0 is given by $p(\mathbf{x}_0 | \mu) := \mathcal{N}(\mathbf{x}_0; \mu \mathbf{1}_m, \mathbf{l}_m)$, where $\mathbf{1}_m$ is an m -dimensional vector filled with ones. Without loss of generality, we set μ to be zero to simplify notation used in the derivation. This way, the parameters controlling the model are collected as $\boldsymbol{\theta} = (a, b, d)$. It is well known that the filtering density $(\mathbf{X}_t | \boldsymbol{\Theta} = \boldsymbol{\theta}, \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ and the smoothing density $(\mathbf{X}_k | \boldsymbol{\Theta} = \boldsymbol{\theta}, \mathbf{Y}_{1:t} = \mathbf{y}_{1:t})$ for any $k < t$ are Gaussian conditioned on $\boldsymbol{\Theta} = \boldsymbol{\theta}$. See Särkkä (2013) and reference therein for details. Using this property, here we derive the analytical form of the posterior parameter density $p(\boldsymbol{\theta} | \mathbf{y}_{1:t})$.

As a starting point, we collect all the quantities from time 0 to t in vectorized forms,

$$\begin{aligned} \bar{\mathbf{X}} &= [\mathbf{X}_0; \dots; \mathbf{X}_t] \in \mathbb{R}^{m(t+1)}, \\ \bar{\mathbf{Y}} &= [\mathbf{Y}_1; \dots; \mathbf{Y}_t] \in \mathbb{R}^{nt}, \\ \bar{\varepsilon}^{(x)} &= [\varepsilon_0^{(x)}; \dots; \varepsilon_t^{(x)}] \in \mathbb{R}^{m(t+1)}, \\ \bar{\varepsilon}^{(y)} &= [\varepsilon_1^{(y)}; \dots; \varepsilon_t^{(y)}] \in \mathbb{R}^{nt}, \end{aligned}$$

where $\varepsilon_0^{(x)} \sim \mathcal{N}(0, \mathbf{l}_m)$ is independent of the other state noises $\varepsilon_k^{(x)}$, $k \geq 1$. Corresponding to the collected states $\bar{\mathbf{X}}$ and the collected observations $\bar{\mathbf{Y}}$, we define the joint state transition matrix $\mathbf{A}_\theta \in \mathbb{R}^{m(t+1) \times m(t+1)}$ and the joint observation matrix $\mathbf{H} \in \mathbb{R}^{nt \times m(t+1)}$ as

$$\mathbf{A}_\theta = \begin{bmatrix} a \mathbf{l}_m & & & & & \\ -b \mathbf{l}_m & \mathbf{l}_m & & & & \\ & -b \mathbf{l}_m & \mathbf{l}_m & & & \\ & & \ddots & \ddots & & \\ & & & -b \mathbf{l}_m & \mathbf{l}_m & \end{bmatrix} \quad \text{and} \quad \mathbf{H} = \begin{bmatrix} 0 & \mathbf{C} & & & & \\ & 0 & \mathbf{C} & & & \\ & & 0 & \mathbf{C} & & \\ & & & \ddots & \ddots & \\ & & & & 0 & \mathbf{C} \end{bmatrix},$$

respectively, where $\mathbf{0}$ denotes a matrix of size $n \times m$ filled with zeros. Then, the state and observation processes defined in (39) can be equivalently expressed as

$$\bar{\mathbf{X}} = a \mathbf{A}_\theta^{-1} \bar{\varepsilon}^{(x)}, \quad \bar{\mathbf{Y}} = \mathbf{H} \bar{\mathbf{X}} + d \bar{\varepsilon}^{(y)}. \quad (40)$$

The joint posterior density of the state path $\mathbf{X}_{0:t} \equiv \bar{\mathbf{X}}$ and the parameters Θ takes the form

$$p(\bar{\mathbf{x}}, \boldsymbol{\theta} | \bar{\mathbf{y}}) = \frac{1}{p(\bar{\mathbf{y}})} p(\bar{\mathbf{y}} | \bar{\mathbf{x}}, \boldsymbol{\theta}) p(\bar{\mathbf{x}} | \boldsymbol{\theta}) p(\boldsymbol{\theta}).$$

Our goal is to derive the posterior parameter density

$$p(\boldsymbol{\theta} | \bar{\mathbf{y}}) = \int p(\bar{\mathbf{x}}, \boldsymbol{\theta} | \bar{\mathbf{y}}) d\bar{\mathbf{x}} = \frac{1}{p(\bar{\mathbf{y}})} p(\boldsymbol{\theta}) p(\bar{\mathbf{y}} | \boldsymbol{\theta}), \quad (41)$$

where $p(\bar{\mathbf{y}} | \boldsymbol{\theta})$ is the marginal likelihood defined as

$$p(\bar{\mathbf{y}} | \boldsymbol{\theta}) = \int p(\bar{\mathbf{y}} | \bar{\mathbf{x}}, \boldsymbol{\theta}) p(\bar{\mathbf{x}} | \boldsymbol{\theta}) d\bar{\mathbf{x}}.$$

Combining the two equations in (40), we obtain

$$\bar{\mathbf{Y}} = a \mathbf{H} \mathbf{A}_\theta^{-1} \bar{\varepsilon}^{(x)} + d \bar{\varepsilon}^{(y)}.$$

Note that $\bar{\varepsilon}^{(x)}$ and $\bar{\varepsilon}^{(y)}$ are independent Gaussian random variables. Thus, conditioned on the parameters $\Theta = \boldsymbol{\theta}$, the observations $\bar{\mathbf{Y}}$ follow a multivariate Gaussian distribution:

$$(\bar{\mathbf{Y}} | \Theta = \boldsymbol{\theta}) \sim \mathcal{N}(\bar{\mathbf{y}}; 0, \bar{\Sigma}_\theta) \quad \text{where} \quad \bar{\Sigma}_\theta = a^2 \mathbf{H} \mathbf{A}_\theta^{-1} \mathbf{A}_\theta^{-\top} \mathbf{H}^\top + d^2 \mathbf{I}_{nt}.$$

To be more specific, we have

$$p(\bar{\mathbf{y}} | \boldsymbol{\theta}) = (2\pi)^{-nt/2} \det(\bar{\Sigma}_\theta)^{-1/2} \exp\left(-\frac{1}{2} \bar{\mathbf{y}}^\top \bar{\Sigma}_\theta^{-1} \bar{\mathbf{y}}\right).$$

The explicit form of the posterior parameter density follows from (41).

References

- Brian D O Anderson and John B Moore. *Optimal filtering*. Courier Corporation, 2012.
- Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342, 2010.
- Alexandros Beskos, Ajay Jasra, Nikolas Kantas, and Alexandre Thiery. On the convergence of adaptive sequential Monte Carlo methods. *Annals of Applied Probability*, 26(2):1111–1146, 2016.
- Daniele Bigoni, Allan P Engsig-Karup, and Youssef M Marzouk. Spectral tensor-train decomposition. *SIAM Journal on Scientific Computing*, 38(4):A2405–A2439, 2016.

- Yoram Bresler. Two-filter formulae for discrete-time non-linear Bayesian smoothing. *International Journal of Control*, 43(2):629–641, 1986.
- Mark Briers, Arnaud Doucet, and Simon Maskell. Smoothing algorithms for state-space models. *Annals of the Institute of Statistical Mathematics*, 62:61–89, 2010.
- Hans-Joachim Bungartz and Michael Griebel. Sparse grids. *Acta Numerica*, 13:147–269, 2004.
- Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in Hidden Markov Models*. Springer Science & Business Media, 2006.
- James Carpenter, Peter Clifford, and Paul Fearnhead. Improved particle filter for nonlinear problems. *IEE Proceedings-Radar, Sonar and Navigation*, 146(1):2–7, 1999.
- Nan Chen and Andrew J Majda. Conditional Gaussian systems for multiscale nonlinear stochastic systems: Prediction, state estimation and uncertainty quantification. *Entropy*, 20(7):509, 2018.
- Nan Chen, Andrew J Majda, and Xin T Tong. Rigorous analysis for efficient statistically accurate algorithms for solving Fokker–Planck equations in large dimensions. *SIAM/ASA Journal on Uncertainty Quantification*, 6(3):1198–1223, 2018.
- Nan Chen, Yingda Li, and Honghu Liu. Conditional Gaussian nonlinear system: A fast preconditioner and a cheap surrogate model for complex nonlinear systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 32(5):053122, 2022.
- Nicolas Chopin, Pierre E Jacob, and Omiros Papaspiliopoulos. SMC2: an efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(3):397–426, 2013.
- Alexandre J Chorin and Xuemin Tu. Implicit sampling for particle filters. *Proceedings of the National Academy of Sciences*, 106(41):17249–17254, 2009.
- Gregory M Constantine and Thomas H Savits. A multivariate Faa di Bruno formula with applications. *Transactions of the American Mathematical Society*, 348(2):503–520, 1996.
- Dan Crisan and Joaquin Miguez. Nested particle filters for online parameter estimation in discrete-time state-space Markov models. *Bernoulli*, 24(4A):3039–3086, 2018.
- Tiangang Cui and Sergey Dolgov. Deep composition of tensor-trains using squared inverse Rosenblatt transports. *Foundations of Computational Mathematics*, 22(6):1863–1922, 2022.
- Tiangang Cui and Xin T Tong. A unified performance analysis of likelihood-informed subspace methods. *Bernoulli*, 28(4):2788–2815, 2022.
- Tiangang Cui, James Martin, Youssef M. Marzouk, Antti Solonen, and Alessio Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30:114015, 2014.

- Tiangang Cui, Sergey Dolgov, and Olivier Zahm. Self-reinforced polynomial approximation methods for concentrated probability densities. *arXiv preprint: 2303.02554*, 2023.
- Tiangang Cui, Sergey Dolgov, and Robert Scheichl. Deep importance sampling using tensor trains with application to a priori and a posteriori rare events. *SIAM Journal on Scientific Computing*, 46(1):C1–C29, 2024.
- Ingrid Daubechies. *Ten lectures on wavelets*. SIAM, 1992.
- Carl de Boor, Ronald A DeVore, and Amos Ron. On the construction of multivariate (pre) wavelets. *Constructive approximation*, 9(2):123–166, 1993.
- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. On adaptive resampling strategies for sequential Monte Carlo methods. *Bernoulli*, 18(1):252–278, 2012.
- Ronald A DeVore and Vasil A Popov. Interpolation of besov spaces. *Transactions of the American Mathematical Society*, 305(1):397–414, 1988.
- Josef Dick, Frances Y Kuo, and Ian H Sloan. High-dimensional integration: the quasi-Monte Carlo way. *Acta Numerica*, 22:133–288, 2013.
- Sergey V Dolgov and Dmitry V Savostyanov. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM Journal on Scientific Computing*, 36(5):A2248–A2271, 2014.
- Arnaud Doucet and Adam M Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12(656-704):3, 2009.
- Dinh Dũng, Vladimir Temlyakov, and Tino Ullrich. *Hyperbolic cross approximation*. Springer, 2018.
- Garry A Einicke and Langford B White. Robust extended Kalman filtering. *IEEE Transactions on Signal Processing*, 47(9):2596–2599, 1999.
- Geir Evensen. The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4):343–367, 2003.
- Geir Evensen, Femke C Vossepoel, and Peter Jan van Leeuwen. *Data assimilation fundamentals: A unified formulation of the state and parameter estimation problem*. Springer Nature, 2022.
- Paul Fearnhead, David Wyncoll, and Jonathan Tawn. A sequential smoothing algorithm with linear computational cost. *Biometrika*, 97(2):447–464, 2010.
- Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical Sciences*, pages 163–185, 1998.
- Mathieu Gerber and Nicolas Chopin. Sequential quasi Monte Carlo. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 77(3):509–579, 2015.

- Walter R Gilks and Carlo Berzuini. Following a moving target—Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(1):127–146, 2001.
- Simon J Godsill, Arnaud Doucet, and Mike West. Monte Carlo smoothing for nonlinear time series. *Journal of the American Statistical Association*, 99(465):156–168, 2004.
- Neil J Gordon, David J Salmond, and Adrian FM Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEEE Proceedings F (Radar and Signal Processing)*, volume 140, pages 107–113. IET, 1993.
- Alex Gorodetsky, Sertac Karaman, and Youssef M Marzouk. A continuous analogue of the tensor-train decomposition. *Computer Methods in Applied Mechanics and Engineering*, 347:59–84, 2019.
- Georg A Gottwald and Sebastian Reich. Combining machine learning and data assimilation to forecast dynamical systems from noisy partial observations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(10):101103, 2021.
- Michael Griebel and Helmut Harbrecht. Analysis of tensor approximation schemes for continuous functions. *Foundations of Computational Mathematics*, 23:219–240, 2023.
- Wolfgang Hackbusch. *Tensor spaces and numerical tensor calculus*, volume 42. Springer Science & Business Media, 2012.
- Edward Herbst and Frank Schorfheide. Tempered particle filtering. *Journal of Econometrics*, 210(1):26–44, 2019.
- Truong-Vinh Hoang, Sebastian Krumscheid, Hermann G Matthies, and Raúl Tempone. Machine learning-based conditional mean filter: a generalization of the ensemble Kalman filter for nonlinear data assimilation. *arXiv preprint: 2106.07908*, 2021.
- Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960.
- Nicholas Kantas, Arnaud Doucet, Sumeetpal Sindhu Singh, and Jan Marian Maciejowski. An overview of sequential Monte Carlo methods for parameter estimation in general state-space models. *IFAC Proceedings Volumes*, 42(10):774–785, 2009.
- 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(1):464–489, 2014.
- Genshiro Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.
- Herbert Knothe et al. Contributions to the theory of convex bodies. *The Michigan Mathematical Journal*, 4(1):39–52, 1957.
- Jun S Liu, Rong Chen, and Tanya Logvinenko. A theoretical framework for sequential importance sampling with resampling. In *Sequential Monte Carlo methods in practice*, pages 225–246. Springer, 2001.

- Stephane G Mallat. Multiresolution approximations and wavelet orthonormal bases of $L^2(\mathbb{R})$. *Transactions of the American Mathematical Society*, 315(1):69–87, 1989.
- Simon Maskell and Neil Gordon. A tutorial on particle filters for on-line nonlinear/non-Gaussian Bayesian tracking. *IEEE Target Tracking: Algorithms and Applications (Ref. No. 2001/174)*, pages 2–1, 2002.
- Matthias Morzfeld, Xuemin Tu, Ethan Atkins, and Alexandre J Chorin. A random map implementation of implicit filters. *Journal of Computational Physics*, 231(4):2049–2066, 2012.
- Ivan Oseledets and Eugene Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications*, 432(1):70–88, 2010.
- Ivan Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.
- Hill Peter D. Kernel estimation of a distribution function. *Communications in Statistics-Theory and Methods*, 14(3):605–620, 1985.
- Michael K Pitt and Neil Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599, 1999.
- Michael K Pitt and Neil Shephard. Auxiliary variable based particle filters. *Sequential Monte Carlo methods in practice*, pages 273–293, 2001.
- Sebastian Reich. A nonparametric ensemble transform method for Bayesian inference. *SIAM Journal on Scientific Computing*, 35(4):A2013–A2024, 2013.
- Sebastian Reich and Colin Cotter. *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press, 2015.
- Paul B Rohrbach, Sergey Dolgov, Lars Grasedyck, and Robert Scheichl. Rank bounds for approximating Gaussian densities in the tensor-train format. *SIAM/ASA Journal on Uncertainty Quantification*, 10(3):1191–1224, 2022.
- Murray Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472, 1952.
- Simo Särkkä. *Bayesian filtering and smoothing*. Number 3. Cambridge University Press, 2013.
- Simo Särkkä, Jouni Hartikainen, Isambi S Mbalawata, and Heikki Haario. Posterior inference on parameters of stochastic differential equations via non-linear Gaussian filtering and adaptive mcmc. *Statistics and Computing*, 25(2):427–437, 2015.
- Jie Shen, Tao Tang, and Li-Lian Wang. *Spectral methods: algorithms, analysis and applications*, volume 41. Springer Science & Business Media, 2011.
- Bernard W Silverman. *Density estimation for statistics and data analysis*, volume 26. CRC press, 1986.

- Ian H Sloan and Henryk Woźniakowski. *When are quasi-Monte Carlo algorithms efficient for high dimensional integrals?*, *Journal of Complexity*, 14(1):1–33, 1998.
- Chris Snyder, Thomas Bengtsson, Peter Bickel, and Jeff Anderson. Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12):4629–4640, 2008.
- Alessio Spantini, Daniele Bigoni, and Youssef Marzouk. Inference via low-dimensional couplings. *The Journal of Machine Learning Research*, 19(1):2639–2709, 2018.
- Alessio Spantini, Ricardo Baptista, and Youssef Marzouk. Coupling techniques for nonlinear ensemble filtering. *SIAM Review*, 64(4):921–953, 2022.
- Andrew M Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Dongbin Xiu and George Em Karniadakis. The wiener–askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2):619–644, 2002.
- Jun Yu, Zhenlin Yang and Xibin Zhang. A class of nonlinear stochastic volatility models and its implications for pricing currency options. *Computational Statistics & Data Analysis*, 51(4):2218–2231, 2006.
- Olivier Zahm, Tiangang Cui, Kody Law, Alessio Spantini, and Youssef Marzouk. Certified dimension reduction in nonlinear Bayesian inverse problems. *Mathematics of Computation*, 91(336):1789–1835, 2022.
- Benjamin Zanger, Tiangang Cui, Martin Schreiber, and Olivier Zahm. Sequential transport maps using SoS density estimation and α -divergences. *arXiv preprint*, arXiv:2402.17943, 2024.
- Xibin Zhang and Maxwell L King. Box-Cox stochastic volatility models with heavy-tails and correlated errors. *Journal of Empirical Finance*, 15(3):549–566, 2008.