

HIGH-DIMENSIONAL BAYESIAN FILTERING THROUGH DEEP DENSITY APPROXIMATION

KASPER BÅGMARK AND FILIP RYDIN

ABSTRACT. In this work, we systematically benchmark two recently developed deep density methods for nonlinear filtering. We model the filtering density of a discretely observed stochastic differential equation through the associated Fokker–Planck equation, coupled with Bayesian updates at discrete observation times. The two filters: the deep splitting filter and the deep backward stochastic differential equation filter, are both based on Feynman–Kac formulas, Euler–Maruyama discretizations and neural networks. The two methods are extended to logarithmic formulations providing sound, robust, and positivity-preserving density approximations in increasing state dimension. Comparing to the classical bootstrap particle filter and an ensemble Kalman filter, we benchmark the methods on numerous examples. In the low-dimensional examples the particle filters work well, but when we scale up to a partially observed 100-dimensional Lorenz-96 model, the particle-based methods fail and the logarithmic deep backward stochastic differential equation filter prevails. In terms of computational efficiency, the deep density methods reduce inference time by roughly two to five orders of magnitude relative to the particle-based filters.

1. INTRODUCTION

Filtering problems play a central role in applications, where hidden stochastic variables must be inferred from noisy and partial observations, with examples ranging from molecular dynamics and signal processing to finance and climate modeling [6, 11, 15, 16, 20, 23, 27, 28]. A mathematically rigorous description is provided by the filtering equations, which characterize the evolution of the conditional distribution of the hidden state given the observations.

Classical solution strategies include Kalman-type methods and particle filters. While successful in low dimensions, these approaches face severe challenges in high-dimensional and nonlinear settings: the curse of dimensionality causes particle methods to degenerate [10, 35, 36], and linearization-based approximations commonly utilized in Kalman filters lose accuracy [18, 26]. To overcome these limitations, recent research has explored probabilistic numerical methods that reformulate the filtering problem through associated Partial Differential Equations (PDEs) and Backward Stochastic Differential Equations (BSDEs). Pioneering work has shown that deep neural networks can be used to approximate high-dimensional PDEs and BSDEs with remarkable success [8, 17].

Building on this foundation, the deep splitting method and the deep BSDE method were recently extended and analyzed in the filtering context, leading to the Deep Splitting Filter (DSF) [4] and the deep BSDE Filter (BSDEF) [3]. In the present paper, we complement the initial numerical evaluations of the DSF and BSDEF by providing a systematic numerical comparison of the these two approaches. We also propose log-scale variants that solve the corresponding log-equations. Recasting the equations to log-scale yields positivity-preserving schemes for the density approximation and improves numerical stability in regimes where the filtering density becomes small. Across a broad suite of linear and nonlinear examples, spanning moderate to high dimensions, we investigate accuracy, robustness, and computational cost, and benchmark against the Ensemble Kalman Filter (EnKF), the Extended Kalman Filter (EKF), and a bootstrap Particle Filter (PF). Our aim is to examine the regimes in which these deep density approximations, particularly their log variants, offer practical advantages over classical alternatives. The code for the implementation is publicly available¹.

2020 *Mathematics Subject Classification.* 60G25, 60G35, 62F15, 62G07, 62M20, 65C30, 65M75, 68T07.

Key words and phrases. Filtering problem, high-dimensional, Fokker–Planck equation, backward stochastic differential equations, splitting, deep learning.

¹<https://github.com/bagmark/deep-density-filtering>

1.1. Contributions. This paper considers two density based neural approaches to nonlinear filtering, the deep splitting filter and the deep BSDE filter, together with their logarithmic counterparts (LogDSF, LogBSDEF). The log-formulations are derived from the log-transformed filtering equations (Theorem 2.1) and yield positivity-preserving density approximations together with an alternative objective function that avoids vanishing densities and gradients. All methods are implemented within a common simulation framework, combining Euler–Maruyama discretization for prediction, Bayes’ rule for updates, and normalization through importance sampling with unconditional-Gaussian or EKF-based proposals. Their performance is evaluated on a broad set of linear and nonlinear test cases: long-horizon Ornstein–Uhlenbeck processes in dimension $d = 10$, a nonlinearly observed Schlögl model, and the stochastic Lorenz-96 system with dimension increasing up to $d = 100$. Comparisons are made against classical baselines including the EKF, EnKF, and PF, using up to four different error metrics. This includes a first moment error, a mean absolute error, the Kullback–Leibler divergence, and the negative log-likelihood of the state.

To summarize, the main contributions of this work are:

- (1) A log-density reformulation of the filtering equations (Theorem 2.1) and the resulting LogDSF and LogBSDEF formulations, yielding positivity-preserving schemes for the density approximation and providing a theoretically grounded and numerically robust framework.
- (2) An extensive numerical comparison of DSF, BSDEF, and their logarithmic variants across low-, moderate-, and high-dimensional filtering problems, with baselines provided by the EKF, EnKF, and PF.
- (3) A demonstration that the LogBSDEF formulation remains accurate and stable in high-dimensional, nonlinear, and partially observed regimes, including satisfactory performance on the chaotic Lorenz-96 system with dimension $d = 100$.

1.2. Outline. Section 2 formulates the filtering problem and the associated PDE system, and establishes the log-density equation used throughout the paper. Section 3 reformulates the DSF and BSDEF frameworks together with their logarithmic variants, including their common probabilistic backbone through Feynman–Kac formulas. Section 4 summarizes the experimental setup, including evaluation metrics, with further details provided in the supplementary materials. Section 5 presents the main results, covering numerical experiments for one-dimensional toy examples, high-dimensional linear models, the Schlögl system, and the stochastic Lorenz-96 system. The section ends with a comparison of computational efficiency of the considered methods. Section 6 offer a conclusion and a brief discussion. Appendix A contains the proof of Theorem 2.1. The Supplementary Materials provide all details for reproducibility. It is organized into the following parts: Classical methods, Additional example: Linear spring-mass, Normalization, Model architectures, Training, Evaluation, Additional figures, The chemical Schlögl model.

1.3. Notation. Throughout this work, $T > 0$, is the time horizon and $K \in \mathbb{N}$ the number of observation times, with the uniform grid $t_k = k \frac{T}{K}$, for $k = 1, \dots, K$. The state space is \mathbb{R}^d and the observations are in $\mathbb{R}^{d'}$, where $d, d' \in \mathbb{N}$. For an observation sequence we write $o_{1:k} = (o_1, \dots, o_k) \in \mathbb{R}^{d' \times k}$ and denote the observation space $\mathbb{O}_k := \mathbb{R}^{d' \times k}$. The Euclidean norm is $\|\cdot\|$. For scalar fields $\varphi: \mathbb{R}^d \rightarrow \mathbb{R}$, $\nabla\varphi$ denotes the gradient; for vector-valued maps $g: \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$, Dg denotes the Jacobian. Function spaces follow standard conventions: $C^k(\mathbb{R}^d; \mathbb{R}^n)$ consists of all k -times continuously differentiable functions from \mathbb{R}^d to \mathbb{R}^n , $C^{1,2}([s, t] \times \mathbb{R}^d; \mathbb{R}^n)$ the functions which are once differentiable in time and twice differentiable in space, and $L^\infty(\mathbb{X}; \mathbb{Y})$ contains the essentially bounded measurable functions $\mathbb{X} \rightarrow \mathbb{Y}$. The Gaussian law with mean $m \in \mathbb{R}^d$ and covariance $Q \in \mathbb{R}^{d \times d}$, is $\mathcal{N}(m, Q)$, where we let $\mathcal{N}(x | m, Q)$ denote the Gaussian density evaluated at $x \in \mathbb{R}^d$.

2. THE FILTERING PROBLEM

Let $(\Omega, \mathcal{A}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$ be a complete filtered probability space with the filtration $\mathcal{F} := (\mathcal{F}_t)_{0 \leq t \leq T}$ being generated by the two independent d -dimensional Brownian motions B and W . Assuming sufficient regularity for the functions $\mu: \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, we have a unique strong

solution $S = (S_t)_{t \in [0, T]}$ to the Stochastic Differential Equation (SDE) [32]

$$(1) \quad S_t = S_0 + \int_0^t \mu(S_s) ds + \int_0^t \sigma(S_s) dB_s, \quad t \in [0, T],$$

where S_0 is distributed according to some initial distribution π_0 . We let $h: \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ be a measurement function, implicitly defining the stochastic observation process $O = (O_k)_{k=1}^K$ by

$$(2) \quad O_k = h(S_{t_k}) + V_k, \quad k = 1, \dots, K,$$

where $V_k \sim \mathcal{N}(0, R)$ with some $d' \times d'$ -covariance matrix R . From (2) we obtain the likelihood function $L(o, x) = \mathcal{N}(o \mid h(x), R)$, which will be used throughout this paper. The methods presented in Section 3 are, however, applicable to more general likelihood models beyond the Gaussian case. Within this setting, our objective is to infer the latent state S given the sequence of observations O . More precisely, for $k = 1, \dots, K$, and $o_{1:k} \in \mathbb{O}_k$, we want to find the conditional probability density p_k that satisfies, for a measurable set $B \subset \mathbb{R}^d$,

$$\mathbb{P}(S_{t_k} \in B \mid o_{1:k}) = \int_B p_k(x \mid o_{1:k}) dx.$$

The conditional density, commonly referred to as the filtering density, satisfies a system of PDEs. To introduce this system we first define $a = (a_{ij}) := \sigma \sigma^\top$ with entries a_{ij} , the infinitesimal generator A , defined by the coefficients of (1), and its adjoint A^* . For $\varphi \in C^2$, A and A^* are given by

$$A\varphi = \frac{1}{2} \sum_{i,j=1}^d a_{ij} \frac{\partial^2 \varphi}{\partial x_i \partial x_j} + \sum_{i=1}^d \mu_i \frac{\partial \varphi}{\partial x_i}, \quad A^*\varphi = \frac{1}{2} \sum_{i,j=1}^d \frac{\partial^2}{\partial x_i \partial x_j} (a_{ij} \varphi) - \sum_{i=1}^d \frac{\partial}{\partial x_i} (\mu_i \varphi).$$

Expanding the derivatives in A^* , we see that $(A^*\varphi)(x) = A\varphi(x) + f(x, \varphi(x), \nabla \varphi(x))$, where the function $f: \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$ is defined by

$$f(x, u, v) = \sum_{i,j=1}^d \frac{\partial a_{ij}(x)}{\partial x_i} v_j + \frac{1}{2} \sum_{i,j=1}^d \frac{\partial^2 a_{ij}(x)}{\partial x_i \partial x_j} u - \sum_{i=1}^d \frac{\partial \mu_i(x)}{\partial x_i} u - 2 \sum_{i=1}^d \mu_i(x) v_i.$$

Assuming sufficient regularity on μ and σ , we denote by $p = (p_k)_{k=0}^K$ the conditional density that solves the unnormalized filtering equations. These equations are initialized at $p_0(0) = \pi_0$, and are recursively given, for $k = 0, \dots, K-1$, $x \in \mathbb{R}^d$, $o_{1:k} \in \mathbb{O}_k$, and $t \in [t_k, t_{k+1}]$, by

$$(3) \quad p_k(t, x, o_{1:k}) = p_k(t_k, x, o_{1:k}) + \int_{t_k}^t \left(A p_k(s, x, o_{1:k}) + f(x, p_k(s, x, o_{1:k}), \nabla p_k(s, x, o_{1:k})) \right) ds,$$

$$(4) \quad p_{k+1}(t_{k+1}, x, o_{1:k+1}) = p_k(t_{k+1}, x, o_{1:k}) L(o_{k+1}, x).$$

The first equation (3) governs the deterministic evolution of the distribution of S between two observations, while (4) incorporates the observations through Bayes' formula. The solution $p = (p_k)_{k=0}^K$ consists of $K+1$ piecewise continuous conditional densities satisfying, $p_k(o_{1:k}) \in C^{1,2}([t_k, t_{k+1}] \times \mathbb{R}^d)$ for all $k = 0, \dots, K-1$ and $o_{1:k} \in \mathbb{O}_k$. This is shown under a parabolic Hörmander condition in [4].

Considering that h and R are assumed to be known, the update equation (4) is tractable since we omit the normalizing constant. However, it is sometimes beneficial to consider the normalized density as we discuss further in Section 4.1. It remains to solve the prediction equation (3), also known as the Fokker–Planck equation or the Kolmogorov forward equation. In Section 3 we recall two optimization-based methods that approximately solve the filtering equations, the deep splitting filter [4] and the deep BSDE filter [3]. The main objective of this work is comparing these methods that are constructed to mitigate the curse of dimensionality through Monte Carlo approximations and neural networks.

However, an intrinsic issue of solving for p directly in (3)–(4), is to handle the fact that probability density values diminish in increasing dimension. More precisely, with increasing state dimension d , the numerical values of p decrease exponentially. This can be seen by considering an initial Gaussian distribution $\pi_0 = \mathcal{N}(0, I)$ in a 100-dimensional space where the value of π_0 in the origin is approximately 10^{-40} . Multiplying this with a likelihood of a similar structure quickly becomes unstable if not handled correctly. Moreover, since the methods that we consider rely on iterative

gradient descent to optimize a set of parameters, they are further affected by the vanishing gradient phenomenon [34]. Both of these issues can be mitigated by directly modeling the log-density of the filtering equations, which leads to positivity-preserving schemes for the density approximation. We state these equations next and prove them in Appendix A.

Theorem 2.1. *Let p be the solution to (3)–(4). Define the log-transformed solution by $v = (v_k)_{k=0}^K$, $v_k: [t_k, t_{k+1}] \times \mathbb{R}^d \times \mathbb{O}_k \rightarrow \mathbb{R}$, such that $v_k = -\log p_k$, for $k = 0, \dots, K$. The log-density v , satisfying $v_0(0) = -\log(\pi_0)$, is, for $k = 0, \dots, K-1$, $x \in \mathbb{R}^d$, $o_{1:k} \in \mathbb{O}_k$, and $t \in [t_k, t_{k+1}]$, the solution to*

$$(5) \quad v_k(t, x, o_{1:k}) = v_k(t_k, x, o_{1:k}) + \int_{t_k}^t \left(Av_k(s, x, o_{1:k}) + f_{\log}(x, v_k(s, x, o_{1:k}), \nabla v_k(s, x, o_{1:k})) \right) ds,$$

$$(6) \quad v_{k+1}(t_{k+1}, x, o_{1:k+1}) = v_k(t_{k+1}, x, o_{1:k}) - \log(L(o_{k+1}, x)),$$

where f_{\log} is defined, for $x \in \mathbb{R}^d$, $u \in \mathbb{R}$, and $w \in \mathbb{R}^d$, by

$$f_{\log}(x, u, w) = -\frac{1}{2} \|\sigma(x)^\top w\|^2 - f(x, 1, -w).$$

We note two aspects about (5)–(6). The first and most important one is that the log-equation is also written in terms of the generator A , which means it is a type of nonlinear Kolmogorov backward equation. This makes both the deep splitting and the deep BSDE methods applicable. Secondly, the first order term f_{\log} is now independent of its second argument and nonlinear in its third argument.

This formulation also has an important practical implication: approximating the solution in terms of the log-density naturally leads to positivity-preserving density approximations when the filter is reconstructed by exponentiation. In the next section, we introduce numerical approximation schemes for both the density and log-density formulations.

3. DEEP DENSITY METHODS

In this section we outline the methods under consideration, namely the deep splitting filter and the deep BSDE filter. In addition to these previously studied approaches, we also propose adaptations based on the logarithmic formulation (5)–(6), which yield positivity-preserving schemes for the density approximation. Despite their different appearances, the derivations of these methods follow a very similar structure, and it is convenient to introduce a common probabilistic framework from which they can be developed. To this end, we define an auxiliary process $X: [0, T] \times \Omega \rightarrow \mathbb{R}^d$, that for all $t \in [0, T]$, \mathbb{P} -a.s., satisfies

$$X_t = X_0 + \int_0^t \mu(X_s) ds + \int_0^t \sigma(X_s) dW_s, \quad X_0 \sim q_0.$$

The particular choice of initial distribution q_0 is not essential for the derivation, but a natural choice is to set $q_0 = \pi_0$, which ensures that X equals S in distribution. In some examples, however, we will find it advantageous to select q_0 differently, for instance to increase the variance in the trajectories of X . Both methods are obtained by applying Itô's formula to $p_k(t_k - t, X_t)$, which yields two closely related but distinct Feynman–Kac representations (corresponding to the two approximation strategies).

Since the methods are based on simulation, and closed-form solutions are only available in some specific cases, we approximate the process X . We fix a uniform time step $\tau = \frac{T}{KN} = t_{k,n+1} - t_{k,n}$, $k = 0, \dots, K-1$, $n = 0, \dots, N-1$, and approximate X , by \mathcal{X}^k , on each interval $[t_k, t_{k+1}]$ with the Euler–Maruyama scheme, for $n = 0, \dots, N-1$, given by

$$(7) \quad \mathcal{X}_{n+1}^k = \mathcal{X}_n^k + \mu(\mathcal{X}_n^k)(t_{k,n+1} - t_{k,n}) + \sigma(\mathcal{X}_n^k)(W_{t_{k,n+1}} - W_{t_{k,n}}).$$

Successive intervals are linked by setting $\mathcal{X}_N^{k-1} = \mathcal{X}_0^k$, so that there is a one-to-one correspondence between the discretization times $t_{k,n}$ and the indices \mathcal{X}_n^k . This construction provides the numerical backbone on which both the deep splitting and deep BSDE filters are formulated.

3.1. Deep splitting filter. The original deep splitting methodology was developed for parabolic PDEs in [8], analyzed in [21], later extended to a class of SPDEs in [7] and to partial integro-differential equations in [19]. In [4] the method was applied to the filtering equations with suitable adaptations, and a strong convergence order of one was proven under a parabolic Hörmander condition with smooth and bounded coefficients. To remain coherent with the deep BSDE framework, we adopt a slightly modified, but equivalent, notation. The idea is to approximate the conditional predictive density through a recursive optimization problem. We define, for a fixed discretization time step τ , the operator G^τ acting on a test function $\varphi \in C^1(\mathbb{R}^d; \mathbb{R})$ by

$$(G^\tau \varphi)(x) = \varphi(x) + \tau f(x, \varphi(x), \nabla \varphi(x)), \quad x \in \mathbb{R}^d.$$

For every $k = 0, \dots, K-1$ and $n = 0, \dots, N-1$, the approximation $\phi_{k,n+1}$ is then defined, for $k = 0, \dots, K-1$, $n = 0, \dots, N-1$, as the solution of the recursive optimization problem

$$(8) \quad \min_{\phi \in L^\infty(\mathbb{O}_k; C(\mathbb{R}^d; \mathbb{R}))} \mathbb{E} \left[\left| \phi(\mathcal{X}_n^k, O_{1:k}) - G^\tau \tilde{g}_{k,n}(\mathcal{X}_{n+1}^k, O_{1:k}) \right|^2 \right],$$

where the function $\tilde{g}_{k,n}$ encodes the recursion and is defined by

$$\tilde{g}_{k,n}(x, o_{1:k}) = \begin{cases} \phi_{k,n}(x, o_{1:k}), & n \geq 1, \\ \phi_{k-1,N}(x, o_{1:k-1})L(o_k, x), & n = 0, k \geq 1, \\ \pi_0(x), & n = 0, k = 0. \end{cases}$$

That is, for prediction steps we use the previous approximation $\phi_{k,n}$, while at observation times ($n = 0$) we incorporate the likelihood L through Bayes' formula as in (4). This recursion yields an approximation of the conditional predictive density p at the discretization times $t_{k,n}$, in the sense that $\phi_{k,n}(x, o_{1:k}) \approx p_k(t_{k,n}, x, o_{1:k})$ for all $k = 0, \dots, K-1$, $n = 1, \dots, N$, $o_{1:k} \in \mathbb{O}_k$ and $x \in \mathbb{R}^d$. To this end, we define the approximative filter $\tilde{p} = (\tilde{p}_k)_{k=1}^K$, of the exact filter $(p_k(t_k))_{k=1}^K$, by

$$\tilde{p}_k(x, o_{1:k}) = \phi_{k-1,N}(x, o_{1:k-1})L(o_k, x), \quad o_{1:k} \in \mathbb{O}_k, x \in \mathbb{R}^d.$$

In the logarithmic formulation of the method, the operator G^τ is modified to act on $\varphi \in C^1(\mathbb{R}^d; \mathbb{R})$ by

$$(G^\tau \varphi)(x) = \varphi(x) + \tau f_{\log}(x, \varphi(x), \nabla \varphi(x)), \quad x \in \mathbb{R}^d,$$

while the recursion $\tilde{g}_{k,n}$, for $x \in \mathbb{R}^d$ and $o_{1:k} \in \mathbb{O}_k$, becomes

$$\tilde{g}_{k,n}(x, o_{1:k}) = \begin{cases} \phi_{k,n}(x, o_{1:k}), & n \geq 1, \\ \phi_{k-1,N}(x, o_{1:k-1}) - \log(L(o_k, x)), & n = 0, k \geq 1, \\ -\log(\pi_0(x)), & n = 0, k = 0. \end{cases}$$

We represent ϕ in (8) by a neural network, approximate the expectation using Monte Carlo samples of (\mathcal{X}, O) , and employ stochastic gradient descent to locate the minimizer. Details on the network architectures and optimization settings are provided in Supplementary Materials SM5.

In summary, the deep splitting method reformulates the filtering equations into a sequence of recursive optimization problems, where each step approximates the action of the operator G^τ on the predictive densities. The recursion alternates between propagation by the Euler–Maruyama scheme and correction through Bayes' formula at observation times. In this way, the approximations $(\phi_{k,n})_{k=0, n=1}^{K-1, N}$ provide a discretized representation of the conditional predictive densities, while the terminal quantities $(\tilde{p}_k)_{k=1}^K$ yield an explicit approximation of the (unnormalized) filter itself.

3.2. Deep BSDE filter. In contrast to the deep splitting scheme outlined above, the deep BSDE formulation leads to a single optimization problem per observation time. The method was introduced in [17] for high-dimensional semilinear PDEs, and has been extended in various directions [1, 9, 13, 25]; see the survey [24] for an overview. In the context of nonlinear filtering, the approach was adapted in [3] by coupling deep BSDE formulations for the prediction step with Bayes' formula for the update, resulting in an iterative alternation between propagation and correction. The approximation error was bounded, under an Hörmander condition, by an estimate that combines an a priori 1/2-order discretization error with an a posteriori term given by the attained objective value.

For each $k = 0, \dots, K - 1$, the optimization problem is designed such that its solution ϕ_k approximates the conditional predictive density $p_k(t_{k+1})$. We solve

$$(9) \quad \min_{\substack{\phi \in L^\infty(\mathbb{O}_k; C(\mathbb{R}^d; \mathbb{R})) \\ (\bar{v}_n)_{n=0}^{N-1} \in L^\infty(\mathbb{O}_k; C(\mathbb{R}^d; \mathbb{R}^d))^N}} \mathbb{E} \left[\left| \mathcal{Y}_{k,N}^{O_{1:k}} - \bar{g}_k(\mathcal{X}_N^k, O_{1:k}) \right|^2 \right],$$

where the process $(\mathcal{Y}_n^{O_{1:k}})_{n=0}^N$ is defined recursively, with $n = 0, \dots, N - 1$, by an Euler–Maruyama discretization of the BSDE

$$\begin{aligned} \mathcal{Y}_0^{O_{1:k}} &= \phi(\mathcal{X}_0^k, O_{1:k}), \\ \mathcal{Y}_{n+1}^{O_{1:k}} &= \mathcal{Y}_n^{O_{1:k}} - f(\mathcal{X}_n^k, \mathcal{Y}_n^{O_{1:k}}, \bar{v}_n(\mathcal{X}_n^k, O_{1:k}))(t_{k,n+1} - t_{k,n}) \\ &\quad + \bar{v}_n(\mathcal{X}_n^k, O_{1:k})^\top \sigma(\mathcal{X}_n^k)(W_{t_{k,n+1}} - W_{t_{k,n}}). \end{aligned}$$

Here the auxiliary functions $(\bar{v}_\ell)_{\ell=0}^{N-1}$, multiplied with σ , act as the approximations to the Z -components in the continuous-time BSDE formulation. The terminal condition \bar{g}_k incorporates the update step via Bayes' formula, and is defined from the previous optimum ϕ_{k-1} , for $x \in \mathbb{R}^d$ and $o_{1:k} \in \mathbb{O}_k$, by

$$\bar{g}_k(x, o_{1:k}) = \begin{cases} \phi_{k-1}(x, o_{1:k-1})L(o_k, x), & k \geq 1, \\ \pi_0(x), & k = 0. \end{cases}$$

We define the deep BSDE filter approximation $\hat{p} = (\hat{p}_k)_{k=1}^K$, of the exact filter $(p_k(t_k))_{k=1}^K$, by

$$\hat{p}_k(x, o_{1:k}) = \phi_{k-1}(x, o_{1:k-1})L(o_k, x), \quad o_{1:k} \in \mathbb{O}_k, x \in \mathbb{R}^d,$$

where $(\phi_k)_{k=0}^{K-1}$ are the obtained solutions of the optimization problems (9). Similar to the deep splitting approach, we also outline the version approximating the log-density (5)–(6). In the logarithmic version we switch f for f_{\log} and define \bar{g}_k , for $x \in \mathbb{R}^d$ and $o_{1:k} \in \mathbb{O}_k$, by

$$\bar{g}_k(x, o_{1:k}) = \begin{cases} \phi_{k-1}(x, o_{1:k-1}) - \log(L(o_k, x)), & k \geq 1, \\ -\log(\pi_0(x)), & k = 0. \end{cases}$$

As in the deep splitting method, $(\phi, (\bar{v}_n)_{n=0}^{N-1})$ in (9) are represented by deep neural networks, the expectation is estimated by Monte Carlo samples of (\mathcal{X}, O) , and the minimizer is found using stochastic gradient descent. Further details on architectures and optimization are given in Supplementary Materials SM5.

In summary, the deep BSDE method approximates the prediction step by discretizing a backward stochastic differential equation associated with the nonlinear filtering PDE. The optimization simultaneously learns the predictive density ϕ_k and the auxiliary processes (\bar{v}_n) , while the update step is enforced through the terminal condition \bar{g}_k . Thus the method condenses the filtering problem into one global optimization per observation time, in contrast to the stepwise recursion of the deep splitting scheme.

4. EXPERIMENTAL SETUP

In this section we briefly outline and summarize details regarding implementation and evaluation of the experiments that we report in Section 5.

4.1. Normalizing constant. The filtering equations (3)–(4), and likewise (5)–(6), yield solutions in the form of unnormalized densities (or log versions thereof). Theoretically, and in the corresponding analysis of the methods [3, 4], this unnormalized setting is sufficient. However, for large numbers of observation times K , it is often beneficial to apply an approximate normalization to enhance numerical stability. This occurs in the update step (4) when multiplying the predicted density with the likelihood, often with very small overlap of probability mass between the two, resulting in a much smaller total probability mass. Without normalization, the resulting density deteriorates numerically when the procedure is repeated many times.

In one dimension it is straightforward to perform a quadrature approximation of the normalization constant. However, as we are mostly interested in scaling to higher dimensions, we tackle the issue by importance sampling. A natural choice of importance distribution is a variance-scaled

extended Kalman filter; in addition, we also consider a cheaper observation-independent wide-tailed Gaussian proposal. Even though in many examples we expect the approximate Kalman filter to perform poorly, it still gives an idea of the domain of interest, and increasing the variance yields effective samples in the domain of the true filter. See Supplementary Materials SM3 for details.

4.2. Training. During training, the aim is to approximately solve the optimization problems (8) and (9) for DSF and BSDEF respectively. If nothing else is stated we have used Fully Connected neural Networks (FCN) to parameterize the networks in the optimization problems. We employ the deep approximative methods for both the original formulation and the log-formulation, yielding four methods of interest, denoted DSF, BSDEF, LogDSF, and LogBSDEF. The details about architectures and training can be found in Supplementary Materials SM5, which includes hyperparameters such as the number of discretizations steps N that is used.

4.3. Classical methods. In Section 5 we benchmark against three classical baselines: the extended Kalman filter, the ensemble Kalman filter, and the bootstrap particle filter. The EKF linearizes the nonlinear model using a first-order Taylor expansion and then applies the Kalman filter, yielding a computationally cheap baseline. The EnKF and PF propagate ensembles/particles forward; achieving a fixed accuracy typically requires an exponentially growing number of ensemble members/particles as the state dimension increases, making these methods computationally expensive. See Supplementary Materials SM1 for details on the three methods.

4.4. Error metrics and evaluation. To capture the differences in the performance of the methods, we investigate moment-based errors and uncertainty-based errors. The first type describes how well the methods manage to capture the mean or how well the computed mean estimates the true trajectory. The second type describes how well we capture the uncertainty contained in the overall density, including the tails of the distribution.

The first moment-based metric is a First Moment Error (FME), which measures the error between the true conditional mean, associated with the solution to (3)–(4), and the approximate mean. The second moment-based metric is a Mean Absolute Error (MAE), which measures the difference between the true trajectory S and the approximative mean. More precisely, we simulate M independent trajectories $(S_{t_k}^{(m)}, O_k^{(m)})_{k=1}^K$, $m = 1, \dots, M$, from (1)–(2). For each $O_k^{(m)}$ we evaluate the true conditional mean $\mu_k^{(m)}$, as well as an approximate mean $\hat{\mu}_k^{(m)}$ from each considered approximate filter. We define the metrics, for each $k = 1, \dots, K$, and $\hat{\mu}_k$, by

$$\begin{aligned} \text{FME}(\hat{\mu}_k) &= \mathbb{E}_{O_k} \left[\|\mu_k - \hat{\mu}_k\| \right] \approx \frac{1}{M} \sum_{m=1}^M \|\mu_k^{(m)} - \hat{\mu}_k^{(m)}\|, \\ \text{MAE}(\hat{\mu}_k) &= \mathbb{E}_{S_{t_k}, O_k} \left[\|S_{t_k} - \hat{\mu}_k\| \right] \approx \frac{1}{M} \sum_{m=1}^M \|S_{t_k}^{(m)} - \hat{\mu}_k^{(m)}\|. \end{aligned}$$

While $\text{FME}(\mu_k) = 0$ by definition, $\text{MAE}(\mu_k) \neq 0$ remains useful as a scale: the mean of the true filter, μ_k , is the $L^2(\Omega, \mathfrak{G}(O_{1:k}); \mathbb{R}^d)$ -optimal estimator (the conditional expectation), so its MAE provides the natural reference level and lower bound. We therefore define the relative MAE (rMAE),

$$\text{rMAE}(\hat{\mu}_k) = \frac{\text{MAE}(\hat{\mu}_k) - \text{MAE}(\mu_k)}{\text{MAE}(\mu_k)},$$

which equals 0 for the optimal estimator and is smaller-is-better in general. We report this metric in percentage to highlight how much larger the error of the approximation $\hat{\mu}_k$ is compared to the optimal estimator μ_k .

For the uncertainty-based metrics, which quantify distributional discrepancy, we use the forward Kullback–Leibler divergence when a reliable reference is available. Recall that p_k solves (3)–(4),

and let \widehat{p}_k denote the density of an approximation. For each $k = 1, \dots, K$, and \widehat{p}_k , we define

$$\begin{aligned} \text{KLD}(\widehat{p}_k) &= \mathbb{E}_{O_{1:k}} \left[D_{\text{KL}}(p_k(t_k, \cdot, O_{1:k}) \parallel \widehat{p}_k(\cdot, O_{1:k})) \right] \\ &= \mathbb{E}_{O_{1:k}} \left[\int_{\mathbb{R}^d} \log \left(\frac{p_k(t_k, x, O_{1:k})}{\widehat{p}_k(x, O_{1:k})} \right) p_k(t_k, x, O_{1:k}) \, dx \right]. \end{aligned}$$

In practice, the approximation of the Kullback–Leibler divergence relies on approximating the integral over \mathbb{R}^d , with J Monte Carlo samples $(Z^{(j,m)})_{j=1}^J$ from the reference distribution $p_k(t_k, x, O_{1:k}^{(m)})$ for every $m = 1, \dots, M$, according to

$$\text{KLD}(\widehat{p}_k) \approx \frac{1}{M} \sum_{m=1}^M \frac{1}{J} \sum_{j=1}^J \log \left(\frac{p_k(t_k, Z^{(j,m)}, O_{1:k}^{(m)})}{\widehat{p}_k(Z^{(j,m)}, O_{1:k}^{(m)})} \right).$$

In the linear-Gaussian setting, where the Kalman filter provides an analytical solution, this is straightforward, since we can sample Z directly from the Gaussian posterior p_k . In the low-dimensional nonlinear setting, where we use a bootstrap particle filter as the reference, we approximate sampling from p_k by drawing Z from the empirical particle distribution at time t_k , using the normalized particle weights as sampling probabilities. This procedure ensures that the Monte Carlo estimate of the Kullback–Leibler divergence remains consistent with the reference solution in both the linear and low-dimensional nonlinear cases.

However, obtaining a reliable reference solution for the full density is generally infeasible in the high-dimensional nonlinear setting. In such cases, instead we benchmark the uncertainty-based performance of an approximation through the Negative Log-Likelihood (NLL) of the ground truth state. The NLL measures how much probability mass the filter assigns to the realized state, and its expectation corresponds to the cross-entropy between the true and approximate posteriors, differing from the KL divergence only by an additive constant [14]. Hence, smaller NLL values indicate better alignment with the true filtering distribution. This metric is not zero for the optimal filter, but we rather seek as small values as possible. The metric is defined, for each $k = 1, \dots, K$, and \widehat{p}_k , by

$$\text{NLL}(\widehat{p}_k) = \mathbb{E}_{S_{t_k}, O_{1:k}} \left[-\log(\widehat{p}_k(S_{t_k}, O_{1:k})) \right] \approx -\frac{1}{M} \sum_{m=1}^M \log(\widehat{p}_k(S_{t_k}^{(m)}, O_{1:k}^{(m)})).$$

In Supplementary Materials SM6 we report all parameters, normalization methods and references used for each example from Section 5.

5. EXPERIMENTS

In this section we benchmark the proposed deep filters, DSF and BSDEF, and their logarithmic variants across problems of increasing difficulty, using the methodology of Section 3 and comparing to PF, EKF, and EnKF. While convergence results, as $N \rightarrow \infty$, are available for the deep density methods under strong regularity conditions, these conditions are often not met in the examples considered here. Nonetheless, the methodology of Section 3 yields well-defined algorithms, and we use the experiments to quantify their empirical performance and robustness compared to the classical approaches. We begin, in Section 5.1, with two one-dimensional toy problems with reliable references: a linear Ornstein–Uhlenbeck process and a nonlinear bistable system. In Section 5.2, robustness and applicability in linear high-dimensional settings are assessed through a long-horizon Ornstein–Uhlenbeck process and a high-dimensional counterpart, both serving as simple canonical tests. Section 5.3 continues with the Schlögl model providing a first step toward physically grounded dynamics. In Section 5.4 we test the models on a stochastic Lorenz-96 system, using high-accuracy PF surrogates when feasible. Performance is reported over time via MAE, FME, and KLD, with particular attention to the stability and accuracy gains of the logarithmic formulations. The final Section 5.5 analyzes computational scaling and inference-time costs. Normalization, architectural, training, and evaluation details appear in Supplementary Materials SM3–SM6.

We systematically compare deep density based methods with classical approaches by varying the particle count in PFs and the ensemble size in EnKFs. To keep the presentation focused, we report configurations that highlight performance differences and omit less informative variants.

5.1. Toy examples. We consider two simple one-dimensional equations, one linear and one nonlinear. In both of the examples we set $T = 1$, $K = 10$, $h(x) = x$, $R = 1$, $\sigma(x) = 1$ and $\pi_0 = \mathcal{N}(0, 1)$. The linear equation has a mean reverting drift $\mu(x) = -x$ and the corresponding SDE is solved by the Ornstein–Uhlenbeck process. In this case, the reference solution to (3)–(4) is provided by the Kalman filter. The drift in the nonlinear equation, $\mu(x) = -\frac{2}{5}(5x - x^3)$, is the negative gradient of a double-well potential. Hence the unconditional distribution of the solutions becomes bimodal and we refer to this process as a bistable process. This example, but with discrete dynamics, was previously studied in [5]. In this case we use a particle filter with 10^6 particles and 128 intermediate time steps between observations to sufficiently well approximate the true filter.

In Figure 1 we see the MAE, FME and KLD evaluated on the two examples. We compare the performance to particle filters with a range between 10^4 and 10^6 particles with one intermediate time step, and in the nonlinear case also to the EKF. Both the DSF and the BSDEF perform well but we note that the logarithmic versions, particularly the LogBSDEF, achieve a lower error in almost all metrics. It should be stressed here that the computational time for estimation and density evaluation are about two and five orders of magnitudes faster, respectively, see Section 5.5.

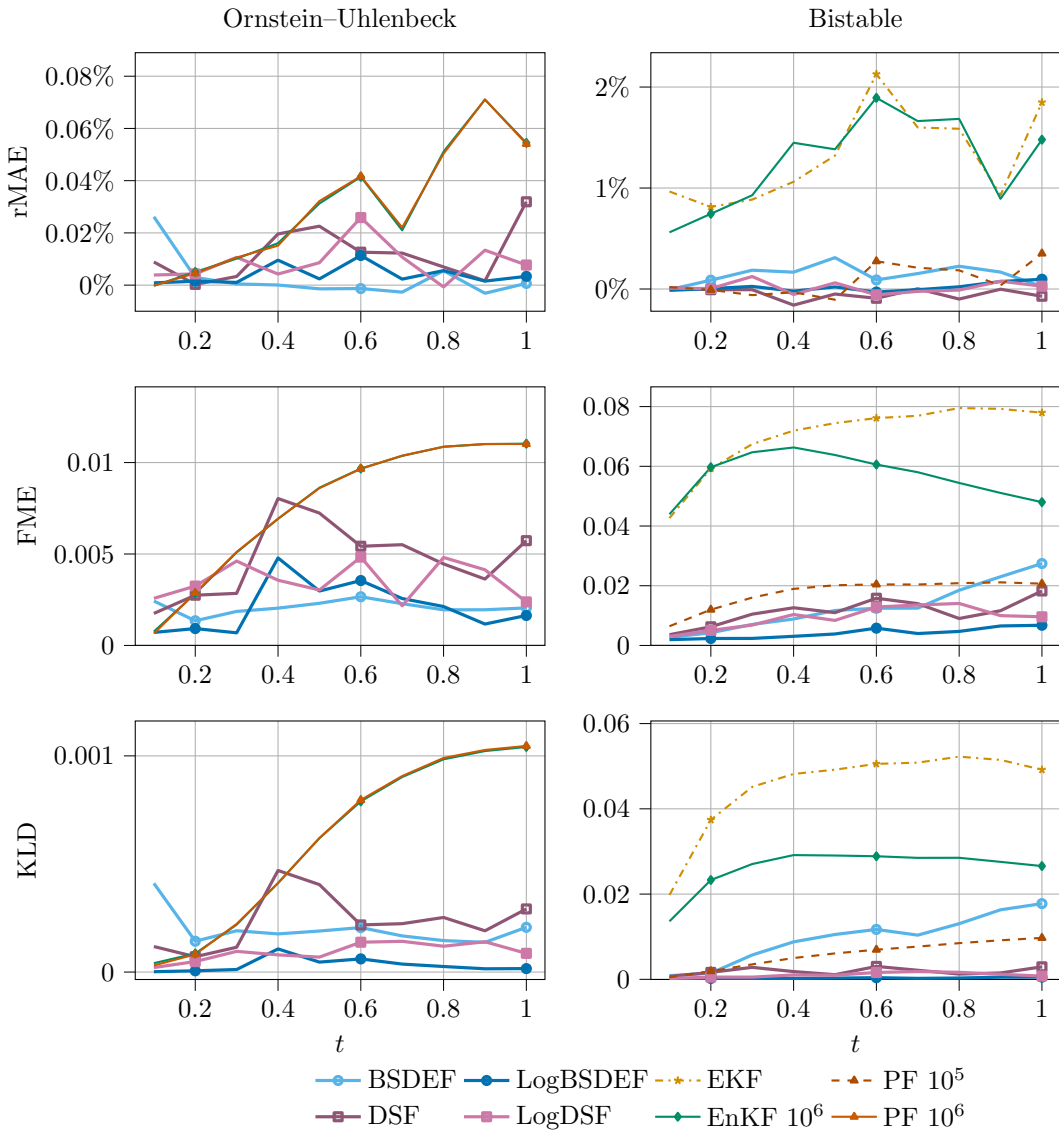


Figure 1. On the left and right panels the results for the Ornstein–Uhlenbeck process and the bistable process are depicted, respectively. From top to bottom the rMAE, FME, and KLD metrics are illustrated.

5.2. High-dimensional Ornstein–Uhlenbeck. To bridge the toy problems and large-scale tests, we study a controlled linear benchmark that isolates temporal and dimensional effects. The Ornstein–Uhlenbeck process serves this purpose, since it admits an analytic Kalman reference and its unconditional density is close to stationarity at $t = 0$ for $\pi_0 = \mathcal{N}(0, I)$. We consider a long-horizon example with $T = 10$, $K = 100$, and $d = 10$, and a high-dimensional example with $T = 1$, $K = 10$, and $d = 100$. In both examples we have $d' = d$, $h(x) = x$, $R = I$, $\mu(x) = -x$, $\sigma(x) = I$, and $\pi_0 = \mathcal{N}(0, I)$. Furthermore, in the long-horizon example we compare two different architectures, namely in addition to the standard FCN, we also employ a Long Short-Term Memory (LSTM)-adaptation, for the LogBSDEF, suitable for long-horizon problems. See Supplementary Materials SM4 for details on the architectures.

In Figure 2 we evaluate the deep learning based methods as well as PFs and EnKFs with different number of particles. Looking at the long-horizon process we see, in particular, that the two LogBSDEF models perform at least as good as the particle filter for the rMAE and FME, and for the KLD it outperforms the classical approximations by a large margin. Moreover, the approximation quality remains stable over time and by introducing the LSTM architecture we make LogBSDEF perform best out of all benchmarked methods. In Supplementary Materials SM7 we demonstrate the methods on a single trajectory by plotting trajectories and marginal densities, where we also see a satisfactory performance of the LogDSF.

For the 100-dimensional Ornstein–Uhlenbeck example we clearly see that the particles filters give quite poor approximations even when using 10^6 particles. The LogBSDEF gives satisfactory results, where the averaged performance is on par with the ensemble Kalman filters, which are expected to yield excellent performance in this linear setting. Remark that we only include the results for LogDSF and LogBSDEF, since the non-logarithmic variants are unstable during training due to the previously mentioned numerical problems with high-dimensional probability densities. Similarly, we found that the LogDSF is unstable in the 100-dimensional example and thus omit it there. In the figures below, whenever a method is excluded, it means that it failed to train properly.

The computational gain from using the LogBSDEF compared to the underperforming PF with 10^6 particles, is even higher than for the one-dimensional example, namely three orders of magnitude faster, see Section 5.5. From this point on we refrain from commenting on more results on computational efficiency. It depends on dimension and not the problem at hand.

5.3. Schlögl model. The nonlinear SDE considered in this example originates from the classical Schlögl model [33, 37, 38], a canonical benchmark for bistable reaction systems. The model describes the concentration process S of a single reacting species subject to autocatalytic reactions and exchange with two external reservoirs. Using the chemical Langevin approximation [22], one can derive the dynamics satisfying (1) with $\sigma: \mathbb{R} \rightarrow \mathbb{R}^{1 \times 4}$ and a four-dimensional Brownian motion $B = (B^{(i)})_{i=1}^4$ corresponding to the four reaction channels. The dynamics are governed by the drift and diffusion coefficients

$$\begin{aligned}\mu(x) &= (b_1(x) - b_2(x) + b_3(x) + b_4(x)), \\ \sigma(x) &= \left(\sqrt{b_1(x)}, -\sqrt{b_2(x)}, \sqrt{b_3(x)}, -\sqrt{b_4(x)} \right),\end{aligned}$$

where

$$b_1(x) = \frac{\theta_1 A}{2} x(x-1), \quad b_2(x) = \frac{\theta_2}{6} x(x-1)(x-2), \quad b_3(x) = \theta_3 B, \quad b_4(x) = \theta_4 x.$$

The reaction rate parameters are fixed to $\theta = (3 \times 10^{-7}, 10^{-4}, 10^{-3}, 3.5)$, $A = 10^5$, and $B = 2 \times 10^5$. This setup produces a bistable system with two distinct equilibrium states separated by a potential barrier. In the Euler–Maruyama discretization (7), the four Brownian channels can be merged into a single effective noise term, producing a scheme, identical in distribution, that reduces the computational cost. A more detailed presentation of the model is given in Supplementary Materials SM8.

We fix $d' = 1$, $T = 1$, $K = 10$, $\pi_0 = \frac{\mathcal{N}(150, 10^2) + \mathcal{N}(350, 10^2)}{2}$, $h(x) = \log(1+x)$, and $R = 0.5$ for the experiment. In the derivation of the deep filtering methods of Section 3 we introduced X to be initialized from a density q_0 . In this example we set q_0 different from π_0 to account for the strong drift of the process. More precisely, we set $q_0 = \frac{\mathcal{N}(150, 25^2) + \mathcal{N}(375, 60^2)}{2}$ to better match the generated samples of X_0 with the distribution of S_{t_1} . See [2, Section 3.4] for an extensive explanation of

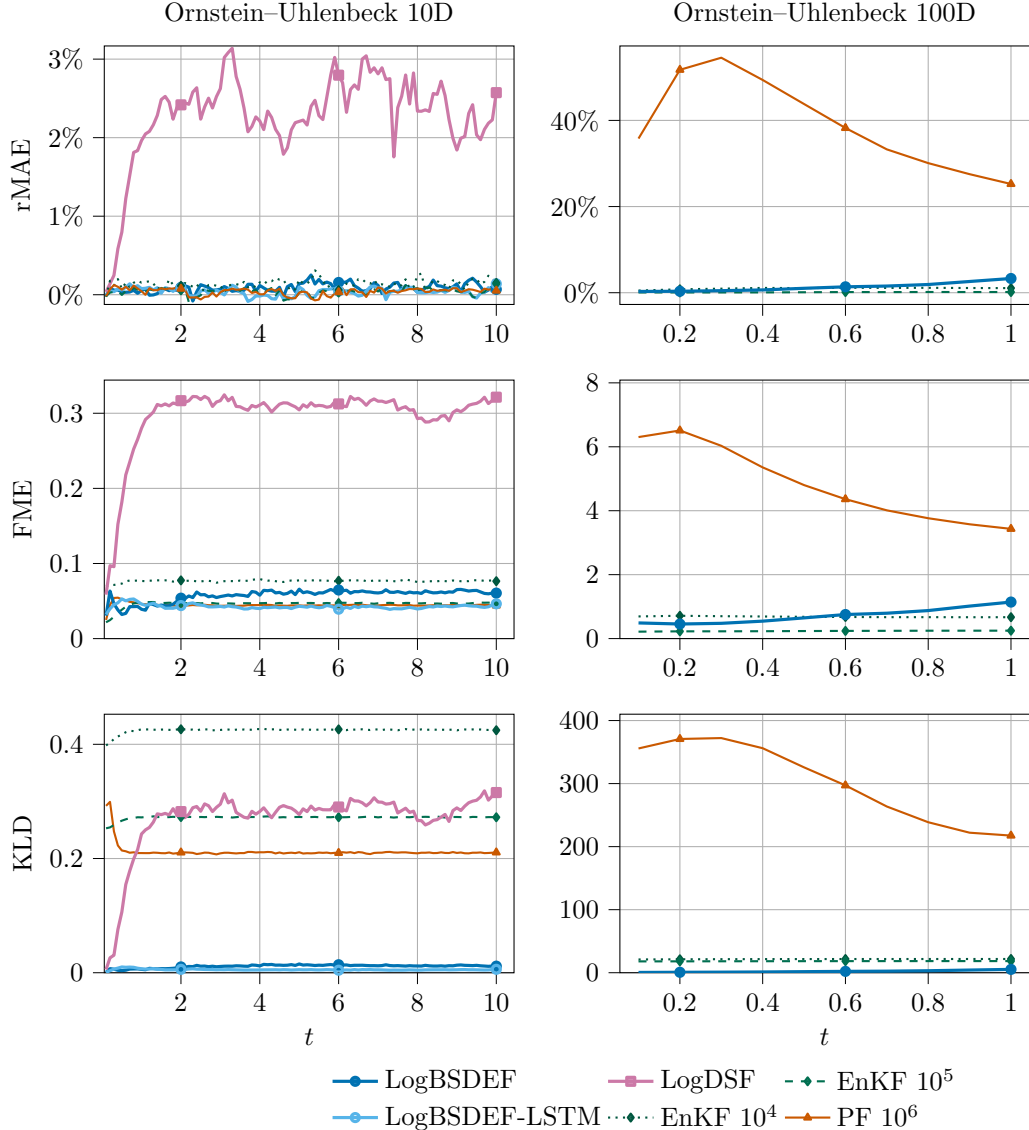


Figure 2. On the left and right panels the results for the long-horizon 10-dimensional Ornstein–Uhlenbeck process and the short-horizon 100-dimensional Ornstein–Uhlenbeck process are depicted, respectively. From top to bottom the rMAE, FME, and KLD metrics are illustrated.

a very similar problem of matching densities. We use a particle filter with 10^6 particles and 128 intermediate time steps as a reference solution.

In the experiment we employ both the LogBSDEF and the LogDSF, but the latter does not converge for the architectures and discretizations that we used, clearly indicating the difficulty of this problem. In the left column of Figure 3 we show the metrics over time, where we compare the LogBSDEF method to particle filters, ensemble Kalman filters, and the extended Kalman filter. Clearly the nonlinearity is too strong for the EKF and EnKFs to yield sufficiently good results. Looking at the FME, our proposed LogBSDEF method performs about as well as the particle filter with 10^5 particles initially but the error accumulates over time. Similar to the other examples our method excels with respect to the KLD, where performance is mostly on par with the particle filter with 10^5 particles.

5.4. Lorenz-96. In our final example, we tackle a strongly nonlinear, high-dimensional system, precisely the regime where classical methods succumb to the curse of dimensionality. The Lorenz-96 model is a high-dimensional chaotic dynamical system originally introduced in [31] as a testbed

for numerical weather prediction. It captures essential features of spatio-temporal chaos, while remaining simple enough for controlled experiments. In its deterministic form, the d -dimensional Lorenz-96 system satisfies (1) with diffusion coefficient $\sigma(x) = 0$ and drift coefficient

$$\mu(x)_i = (x_{i+1} - x_{i-2})x_{i-1} - x_i + F, \quad i = 1, \dots, d,$$

where the indices are taken modulo d , that is $x_i = x_{i+d}$. The parameter $F > 0$ acts as a constant external forcing, typically chosen in the range $F \in [4, 16]$ to produce chaotic dynamics. To account for unresolved processes or model error, we consider a stochastic extension of the Lorenz-96 model by adding additive noise in the form of a Brownian motion with $\sigma(x) = \sigma I$, with $\sigma \in \mathbb{R}$.

The Lorenz-96 SDE serves as a standard testbed for assessing data assimilation techniques, due to its balance of chaotic behavior, high dimensionality, and computational tractability [39]. It is extensively studied in the context of spatio-temporal chaos (e.g., fractal dimension scaling) [29], bifurcation structure and wave propagation [30], as well as various data assimilation and neural network hybrid methods [12].

We fix $d = 4$, $d' = 4$, $h(x) = x$, $T = 1$, $K = 10$, $\pi_0 = \mathcal{N}((F, F, F, F), I)$, $\sigma = I$, $R = 2I$ and $F = 8$. In this strongly nonlinear setting we do not have a good reference solution for increasing dimensions, but in the case of $d = 4$, we found the particle filter with 10^6 particles enhanced with 128 intermediate time steps to be sufficiently good. We tried the LogBSDEF and the LogDSF, and compared their performance to the PF, EKF, and EnKF. In the experiments we tried a multitude of settings to make the LogDSF stable, but it consistently failed to converge. In Supplementary Materials SM5 we detail the settings we used for the LogBSDEF.

In the right column of Figure 3 we demonstrate the results of the rMAE, FME, and KLD. The standard approximative particle filter diverges after some timesteps, where especially the KLD explodes. In the evaluation of the KLD, we clipped the values at $|\log(10^{-200})| \approx 460$, for numerical reasons, which is seen by the flattened curve of the PF. We see that our method performs superior to the EnKF with 10^6 ensemble members in all metrics, but we can also note that none of the methods reach close to the MAE-minimum that the reference solutions obtain, indicating the severe difficulty of the problem. In Supplementary Materials SM7 we illustrate for a single trajectory, the state S and corresponding filter means, projected on the x_1x_2 - and x_3x_4 -planes, respectively, to give an idea about the chaoticity of the system.

We continue by increasing the state dimension to $d = [4, 10, 20, 40, 100]$. In addition, we only have partial observations with $d' = [4, 5, 5, 10, 25]$, where we observe every, every second or every fourth position, respectively. More precisely, the measurement function is defined by

$$(h(x))_i = \begin{cases} x_i, & (d, d') = (4, 4), \\ x_{2i}, & (d, d') = (10, 5), \\ x_{4i}, & (d, d') = (20, 5), (40, 10), (100, 25), \end{cases}$$

for $i = 1, \dots, d'$. The remaining parameters are the same as in the $d = 4$ case. Here it is challenging to obtain reference solutions and instead we measure how the error increases by dimension in both the MAE and NLL metrics. Optimally, the performance is stable with increasing state dimension. We accumulate the scores for all the time steps for each dimension in both metrics, $\text{MAE} = \frac{1}{K} \sum_{k=1}^K \text{MAE}(\hat{p}_k)$, $\text{NLL} = \frac{1}{K} \sum_{k=1}^K \text{NLL}(\hat{p}_k)$ and report the metrics in Figure 4 on log-log axes. Since the NLL grows approximately linearly with the state dimension d and the MAE with $d^{\frac{1}{2}}$, we plot two reference slopes for each metric. In this example we ideally wanted to use more particles for the PFs and ensembles for the EnKFs, but we were limited by memory consumption. The classical methods scale poorly in this regard, especially when recovering the density function, see more on the computational efficiency in Section 5.5. We can see that the PF underperforms severely in almost all dimensions. Furthermore, the performance of the LogBSDEF and the EnKF is mostly on par for up to about $d = 20$, and then the EnKF starts diverging as well. The robust performance of the LogBSDEF demonstrates the potential of deep density methods to extend nonlinear filtering to high-dimensional settings.

5.5. Computational efficiency. Finally, we study the inference time of the methods and how it scales with the state dimension. We display the time required to i) estimate one whole state trajectory by the first moment for each observation time and ii) evaluate each filtering density in

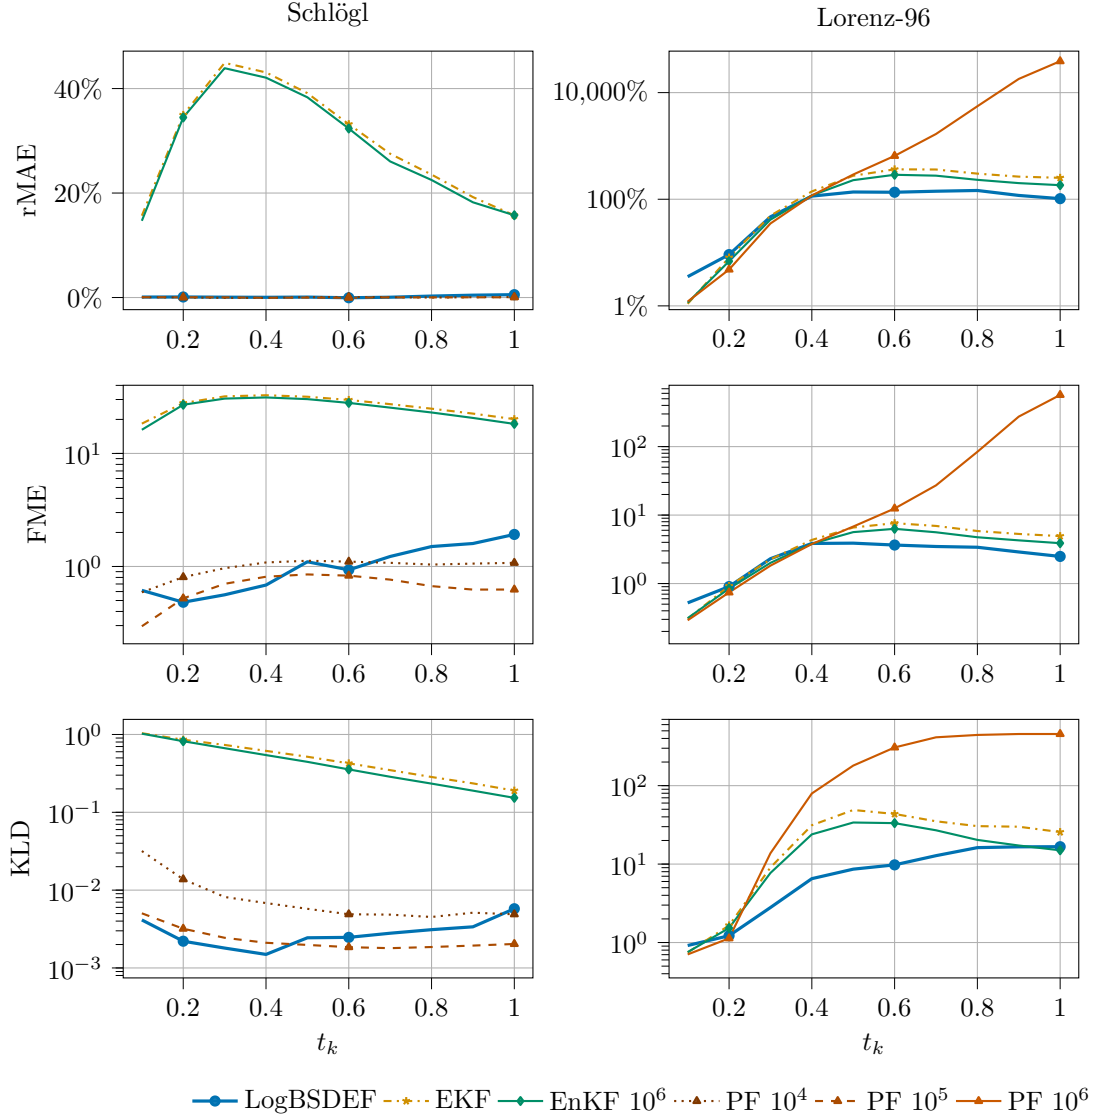


Figure 3. On the left and right panels the results for the Schlögl model and the 4-dimensional Lorenz-96 model are depicted, respectively. From top to bottom the rMAE, FME, and KLD metrics are illustrated.

1000 points. The deep learning-based filters were tested using a single NVIDIA A40 GPU with 48 GB of VRAM. The PF and EnKF were run using an Intel(R) Xeon(R) Platinum 8358 CPU with 32 cores.

We remark that the runtime for all the deep density methods primarily depends on the dimension of the problem (which more or less determines the size of the neural network) as well as the sampling method and the number of samples used to compute the first moment or normalizing constant. Hence, we show the runtime when fixing $T = 1$, $K = 10$ and $d = d'$ for the Ornstein–Uhlenbeck case, and argue that the other problems have more or less the same computational cost. In this setting we show the inference times in Figure 5. As can be seen, BSDEF (whose runtime is equal to that of DSF, LogBSDEF, and LogDSF) scales well with the state dimension, with runtimes remaining rather constant regardless of the sampling method (abbreviated by I-EKF for the EKF importance distribution and I-G for the observation independent Gaussian importance distribution, see Supplementary Materials SM3 for details). In comparison, the benchmark filters explode in estimation time with higher d and their density calculation time is consistently many magnitudes higher due to the expensive kernel density estimate that we utilize to obtain densities (see Supplementary Materials SM6 for details).




Figure 4. The averaged MAE and NLL metrics evaluated over state dimension $d = [4, 10, 20, 40, 100]$. At each time step the NLL value is capped at $|\log(10^{-200})| \approx 460$.

In Figure 6 we report the total time to perform training, state estimation, and density evaluation (with 1000 points per observation time) on two examples, as a function of the number of observation sequences. We include the best-performing deep density based methods alongside a particle filter and an ensemble Kalman filter. For the 100-dimensional Lorenz-96 case, the classical methods do not yield accurate estimates. Their curves are shown only to illustrate computational cost. The trade-off is clear: deep density based methods incur a one time training cost but offer low per-sample inference time, whereas particle-based methods avoid training but have higher per-sample cost. In terms of computation time, the deep density based methods become advantageous once more than about 1300 (bistable) or 430 (Lorenz-96) sequences are processed.

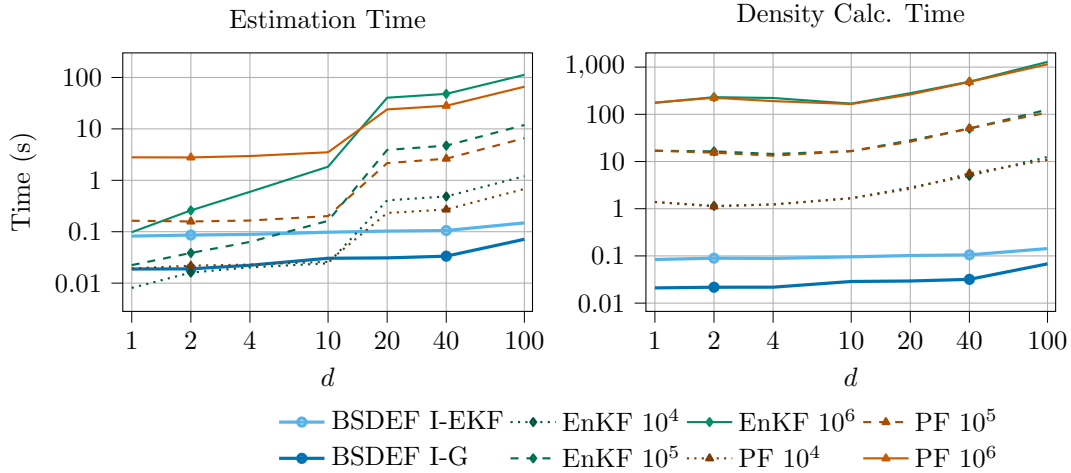


Figure 5. On the left, we display the average time for estimating one whole trajectory in the Ornstein-Uhlenbeck case. On the right, we display the average time for evaluating filtering densities for all observation times in 1000 points. Note that this includes the time to obtain normalization constants for the BSDEF. In both plots, the time it takes to propagate particles for the EnKF and PF is included. For BSDEF, we fix the number of samples to 10^4 used to compute the first moment and normalizing constant, although it is normally possible to use fewer samples without sacrificing much performance.




Figure 6. The computational time, from initializing each method to evaluation of 1000 spatial points in each observation time, over increasing number of sequences. The intersection between the PFs and deep density methods’ computational time occurs at 1300 samples for the bistable example and at 430 samples for the 100-dimensional Lorenz-96 example.

6. CONCLUSION AND DISCUSSION

This work benchmarked deep filtering methods, the Deep Splitting Filter (DSF) and the deep Backward Stochastic Differential Equation Filter (BSDEF), together with their logarithmic variants for nonlinear state estimation in high-dimensional systems. By formulating prediction and update probabilistically, the methods can be trained and allow efficient evaluation.

Across all testbeds, a clear pattern emerges. On linear problems with reliable ground truth, the methods recover high accuracy and remain stable in long horizons and high dimensions. The logarithmic formulations improve numerical robustness and consistently reduce error, with LogBSDEF giving the strongest results. In addition, the log-scale formulation yields positivity-preserving density approximations, eliminating inconsistent negative density estimates and contributing to stable training and evaluation in high dimensions. On strongly nonlinear dynamics, including the Schlögl model and stochastic Lorenz-96 system, the deep density filters remain numerically stable and often competitive, whereas particle and ensemble methods tend to degrade or become costly when the underlying dimension increases. The LogBSDEF exhibits only mild growth in inference time with dimension, supporting its use in high-dimensional settings, unlike particle filters and the ensemble Kalman filter, which inherently incur the curse of dimensionality. The difference in computational time is about two to five orders of magnitude.

Several limitations remain. Performance depends on network architecture and sampling choices, and the selection of the initialization density q_0 can influence stability in stiff regimes. Our experiments use Gaussian observation noise and fixed time discretizations; extending to alternative likelihoods and adaptive stepping is left for future work. Finally, while the methods can achieve high accuracy and stability across many settings, they require careful hyperparameter tuning which remains a practical obstacle. Supplementary Materials SM5 documents the configurations, which confirms the variety of parameters used.

At the same time, the experiments with the LSTM-based model in the long-horizon setting indicate that more advanced neural network architectures can help improve performance. However, we tried a few other architectures that did not make it into the paper, among them transformer-based and other LSTM architectures, which, for these examples and configurations, did not improve performance. Further work in this direction is needed in order to find optimal architectures.

Acknowledgements. The authors would like to thank Adam Andersson, Stig Larsson, and Ruben Seyer for valuable input and discussions on this manuscript. The work of K.B. was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The work of F.R. was funded by the Swedish Electromobility

Centre (SEC) and partially supported by WASP. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at Chalmers e-Commons partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

APPENDIX A. PROOF OF THEOREM 2.1

For simplicity, we hide t , x , o , and k in the notation and write $p = p_k(t, x, o_{1:k})$. The initial condition follows by insertion and simplification. It remains to derive the log-transformed equation from the original Fokker–Planck equation. By taking the time derivative in $v = -\log p$, we obtain

$$(10) \quad \frac{\partial v}{\partial t} = -\frac{1}{p} \frac{\partial p}{\partial t},$$

or equivalently

$$(11) \quad \frac{\partial p}{\partial t} = -p \frac{\partial v}{\partial t}.$$

Analogous relations hold for the spatial derivatives. We consider Ap and substitute $\frac{\partial p}{\partial x_i}$ by (11) to get

$$Ap = \frac{1}{2} \sum_{i,j=1}^d a_{ij} \frac{\partial^2 p}{\partial x_i \partial x_j} + \sum_{i=1}^d \mu_i \frac{\partial p}{\partial x_i} = -\frac{1}{2} \sum_{i,j=1}^d a_{ij} \frac{\partial}{\partial x_i} \left(p \frac{\partial v}{\partial x_j} \right) - \sum_{i=1}^d \mu_i p \frac{\partial v}{\partial x_i}.$$

Next, by the product rule and (11), we obtain

$$(12) \quad \begin{aligned} Ap &= \frac{1}{2} \sum_{i,j=1}^d a_{ij} p \left(\frac{\partial v}{\partial x_i} \frac{\partial v}{\partial x_j} \right) - \frac{1}{2} \sum_{i,j=1}^d a_{ij} p \left(\frac{\partial^2 v}{\partial x_i \partial x_j} \right) - \sum_{i=1}^d \mu_i p \frac{\partial v}{\partial x_i} \\ &= p \left(\frac{1}{2} \|\sigma^\top \nabla v\|^2 - Av \right). \end{aligned}$$

By inserting (12) into the differentiated form of the Fokker–Planck equation, we get

$$(13) \quad \frac{\partial}{\partial t} p = Ap + f(p, \nabla p) = p \left(\frac{1}{2} \|\sigma^\top \nabla v\|^2 - Av \right) + f(p, \nabla p).$$

We continue by inserting (13) into (10) to obtain

$$\begin{aligned} \frac{\partial}{\partial t} v &= -\frac{1}{p} \left(p \left(\frac{1}{2} \|\sigma^\top \nabla v\|^2 - Av \right) + f(p, \nabla p) \right) \\ &= Av - \frac{1}{2} \|\sigma^\top \nabla v\|^2 - \frac{1}{p} f(p, \nabla p). \end{aligned}$$

Since f is linear, we obtain

$$\frac{\partial}{\partial t} v = Av - \frac{1}{2} \|\sigma^\top \nabla v\|^2 - f(1, -\nabla v).$$

The proof is finished by recalling the definition of f_{\log} from the theorem statement.

REFERENCES

- [1] K. Andersson, A. Andersson, and C. W. Oosterlee. The deep multi-FBSDE method: a robust deep learning method for coupled FBSDEs. *J. Sci. Comput.* 106.3 (2026), p. 77.
- [2] K. Bågmark, A. Andersson, and S. Larsson. An energy-based deep splitting method for the nonlinear filtering problem. *Partial Differ. Equ. Appl.* 4 (2023).
- [3] K. Bågmark, A. Andersson, and S. Larsson. Nonlinear filtering based on density approximation and deep BSDE prediction. *arXiv:2508.10630* (2025).
- [4] K. Bågmark, A. Andersson, S. Larsson, and F. Rydin. A convergent scheme for the Bayesian filtering problem based on the Fokker–Planck equation and deep splitting. *arXiv:2409.14585* (2024).
- [5] F. Bao, Z. Zhang, and G. Zhang. A score-based filter for nonlinear data assimilation. *J. Comput. Phys.* 514 (2024), Paper No. 113207, 16.
- [6] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, 2001.
- [7] C. Beck, S. Becker, P. Cheridito, A. Jentzen, and A. Neufeld. Deep learning based numerical approximation algorithms for stochastic partial differential equations and high-dimensional nonlinear filtering problems. *arXiv:2012.01194* (2020).
- [8] C. Beck, S. Becker, P. Cheridito, A. Jentzen, and A. Neufeld. Deep splitting method for parabolic PDEs. *SIAM J. Sci. Comput.* 43 (2021), A3135–A3154.

- [9] C. Beck, W. E, and A. Jentzen. Machine learning approximation algorithms for high-dimensional fully non-linear partial differential equations and second-order backward stochastic differential equations. *J. Nonlinear Sci.* 29.4 (2019), pp. 1563–1619.
- [10] P. Bickel, B. Li, and T. Bengtsson. Sharp failure rates for the bootstrap particle filter in high dimensions. *Pushing the limits of contemporary statistics: contributions in honor of Jayanta K. Ghosh*. Vol. 3. Inst. Math. Stat. (IMS) Collect. Inst. Math. Statist., Beachwood, OH, 2008, pp. 318–329.
- [11] S. S. Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House Publishers, 1999.
- [12] J. Brajard, A. Carrassi, M. Bocquet, and L. Bertino. Combining data assimilation and machine learning to emulate a dynamical model from sparse and noisy observations: a case study with the Lorenz 96 model. *J. Comput. Sci.* 44 (2020), pp. 101171, 11.
- [13] Q. Chan-Wai-Nam, J. Mikael, and X. Warin. Machine learning for semi linear PDEs. *J. Sci. Comput.* 79.3 (2019), pp. 1667–1712.
- [14] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. 2nd. Wiley-Interscience, 2006.
- [15] N. Cui, L. Hong, and J. R. Layne. A comparison of nonlinear filtering approaches with an application to ground target tracking. *Signal Processing* 85 (2005), pp. 1469–1492.
- [16] L. Duc, T. Kuroda, K. Saito, and T. Fujita. Ensemble Kalman filter data assimilation and storm surge experiments of tropical cyclone nargis. *Tellus A* 67 (2015), p. 25941.
- [17] W. E, J. Han, and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat* 5 (Nov. 2017), pp. 349–380.
- [18] M. Ehrendorfer. A review of issues in ensemble-based Kalman filtering. *Meteorol. Z.* 16 (2007).
- [19] R. Frey and V. Köck. Convergence analysis of the deep splitting scheme: the case of partial integro-differential equations and the associated forward backward SDEs with jumps. *SIAM J. Sci. Comput.* 47.1 (2025), A527–A552.
- [20] G. Galanis, P. Louka, P. Katsafados, I. Pytharoulis, and G. Kallos. Applications of Kalman filters based on non-linear functions to numerical weather predictions. *Ann. Geophys* 24 (2006), pp. 1–10.
- [21] M. Germain, H. Pham, and X. Warin. Approximation error analysis of some deep backward schemes for nonlinear PDEs. *SIAM J. Sci. Comput.* 44.1 (2022), A28–A56.
- [22] D. T. Gillespie. The chemical Langevin equation. *J. Chem. Phys.* 113.1 (2000), pp. 297–306.
- [23] I. R. Goodman, R. P. S. Mahler, and H. T. Nguyen. *Mathematics of Data Fusion*. Kluwer Academic Publishers Group, Dordrecht, 1997, pp. xii+507.
- [24] J. Han, A. Jentzen, and W. E. A brief review of the deep BSDE method for solving high-dimensional partial differential equations. *arXiv:2505.17032* (2025).
- [25] J. Han and J. Long. Convergence of the deep BSDE method for coupled FBSDEs. *Probab. Uncertain. Quant. Risk* 5 (2020), Paper No. 5, 33.
- [26] X. Han and X. Li. An evaluation of the nonlinear/non-Gaussian filters for the sequential data assimilation. *Remote Sens. Environ.* 112.4 (2008), pp. 1434–1449.
- [27] M. S. Johannes and N. G. Polson. MCMC methods for continuous-time financial econometrics. *Handbook of Financial Econometrics*. Elsevier, 2009, pp. 1–72.
- [28] K. Kamino, N. Kadakia, F. Avgidis, Z.-X. Liu, K. Aoki, T. S. Shimizu, and T. Emonet. Optimal inference of molecular interaction dynamics in FRET microscopy. *Proc. Natl. Acad. Sci. U.S.A.* 120.15, e2211807120 (2023).
- [29] A. Karimi and M. R. Paul. Extensive chaos in the Lorenz-96 model. *Chaos* 20.4, 043105 (2010).
- [30] D. L. van Kekem and A. E. Sterk. Symmetries in the Lorenz-96 model. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.* 29.1 (2019), pp. 1950008, 18.
- [31] E. N. Lorenz. Predictability: a problem partly solved. *Seminar on Predictability, Vol. I*. ECMWF, Reading, Berkshire, UK, 1996, pp. 1–18.
- [32] B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer Science & Business Media, 2003.
- [33] F. Schlögl. Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift für physik* 253.2 (1972), pp. 147–161.
- [34] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein. Deep information propagation. *Proc. Int. Conf. Learn. Represent.* 2017.
- [35] C. Snyder. Particle filters, the “optimal” proposal and high-dimensional systems. *Proceedings of the ECMWF Seminar on Data Assimilation for atmosphere and ocean*. 2011, pp. 1–10.
- [36] C. Snyder, T. Bengtsson, and M. Morzfeld. Performance bounds for particle filters using the optimal proposal. *Mon. Weather Rev.* 143 (2015), pp. 4750–4761.
- [37] M. Vellela and H. Qian. Stochastic dynamics and non-equilibrium thermodynamics of a bistable chemical system: the Schlögl model revisited. *J. R. Soc. Interface* 6.39 (2009), pp. 925–940.
- [38] M. Vlysidis and Y. N. Kaznessis. Solving stochastic reaction networks with maximum entropy Lagrange multipliers. *Entropy* 20.9 (2018), p. 678.
- [39] D. S. Wilks. Effects of stochastic parametrizations in the Lorenz-96 system. *Q. J. R. Meteorol. Soc.* 131 (2005), pp. 389–407.

KASPER BÅGMARK, DEPARTMENT OF MATHEMATICAL SCIENCES, CHALMERS UNIVERSITY OF TECHNOLOGY AND UNIVERSITY OF GOTHENBURG, SE-412 96 GOTHENBURG, SWEDEN

Email address: `bagmark@chalmers.se`

FILIP RYDIN, DEPARTMENT OF ELECTRICAL ENGINEERING, CHALMERS UNIVERSITY OF TECHNOLOGY, SE-412 96 GOTHENBURG, SWEDEN

Email address: `filipry@chalmers.se`

SUPPLEMENTARY MATERIALS: HIGH-DIMENSIONAL BAYESIAN FILTERING THROUGH DEEP DENSITY APPROXIMATION

KASPER BÅGMARK AND FILIP RYDIN

SM1. CLASSICAL METHODS

In this section we recall the classical methods that we ran our comparisons against. This includes a bootstrap particle filter, the extended Kalman filter, and an ensemble Kalman filter.

SM1.1. Particle filter. The classical bootstrap particle filter, a type of sequential Monte Carlo method, was introduced in [9], and provides an approximation of the nonlinear filter by propagating and reweighting an ensemble of particles. The method is based on $M \in \mathbb{N}$ weighted particles $(\tilde{X}_k^{(i)}, V_k^{(i)})_{i=1}^M$, with $\tilde{X}_k^{(i)}$ denoting the particle and $V_k^{(i)}$ its weight, approximating the posterior distribution at each observation time t_k , $k = 1, \dots, K$.

The filter is initialized by sampling each $\tilde{X}_0^{(i)}$, $i = 1, \dots, M$, according to π_0 , and setting $V_0^{(i)} = \frac{1}{M}$ for $i = 1, \dots, M$. Next, given $(\tilde{X}_{k-1}^{(i)}, V_{k-1}^{(i)})_{i=1}^M$ that approximates $p_{k-1}(t_{k-1})$ with an empirical distribution, the bootstrap particle filter proceeds by:

- (1) **Prediction:** Propagate each particle forward according to F approximating the dynamics of the signal process S , from (1), by

$$\hat{X}_k^{(i)} = F(\tilde{X}_{k-1}^{(i)}, \Delta W_k^{(i)}), \quad i = 1, \dots, M.$$

Here $\Delta W_k^{(i)}$ is the Brownian motion increment for the i 'th particle at time t_k . For instance we can let F be the forward operator given by the Euler–Maruyama discretization defined from (7).

- (2) **Weighting:** Compute importance weights proportional to the likelihood of the observation O_k given the predicted particles:

$$\tilde{V}_k^{(i)} = V_{k-1}^{(i)} L(O_k, \hat{X}_k^{(i)}), \quad i = 1, \dots, M.$$

Normalize weights:

$$\bar{V}_k^{(i)} = \frac{\tilde{V}_k^{(i)}}{\sum_{j=1}^M \tilde{V}_k^{(j)}}, \quad i = 1, \dots, M.$$

- (3) **Resampling:** Generate a new set of particles $(\tilde{X}_k^{(i)}, V_k^{(i)})_{i=1}^M$ by sampling with replacement from $\{\hat{X}_k^{(i)}\}_{i=1}^M$ according to the new normalized weights $\{\bar{V}_k^{(i)}\}_{i=1}^M$, and reset weights to $V_k^{(i)} = 1/M$.

Under standard regularity assumptions (bounded likelihood, nondegenerate resampling), the empirical measure

$$\bar{p}_k^M = \frac{1}{M} \sum_{i=1}^M \delta_{\tilde{X}_k^{(i)}}$$

is a consistent estimator of the filtering distribution in the sense that, for any bounded test function φ ,

$$\bar{p}_k^M(\varphi) = \frac{1}{M} \sum_{i=1}^M \varphi(\tilde{X}_k^{(i)}) \xrightarrow[M \rightarrow \infty]{a.s.} \mathbb{E}[\varphi(S_{t_k}) \mid O_{1:k}].$$

Moreover, a central limit theorem holds with \sqrt{M} -rate convergence and an explicit asymptotic variance, see [3, 4, 5]. The bootstrap particle filter thus provides a sequential Monte Carlo approximation of the nonlinear filter that converges almost surely as $M \rightarrow \infty$, but suffers from the curse of dimensionality, as its variance grows exponentially with the dimension of the state space [1].

SM1.2. Extended Kalman filter. For completeness we briefly recall the main steps of the continuous-discrete extended Kalman filter, which provides a Gaussian approximation of the filtering distribution. The EKF propagates the conditional mean and covariance through a local linearization of the dynamics and a Kalman-like update at observation times. We let $(m_k, P_k)_{k=0}^K$ denote the mean and covariance parameterizing the EKF at each time t_k . If π_0 is a Gaussian density, we simply initiate m_0 and P_0 with the true parameters from π_0 . The method recursively approximates the prediction step (3) and the update step (4).

- (1) **Prediction:** Given the mean and covariance (m_{k-1}, P_{k-1}) at t_{k-1} , we integrate the moment equations, initialized at $\hat{m}(t_{k-1}) = m_{k-1}$ and $\hat{P}(t_{k-1}) = P_{k-1}$, over $t \in [t_{k-1}, t_k]$

$$\frac{d\hat{m}(t)}{dt} = \mu(\hat{m}(t)), \quad \frac{d\hat{P}(t)}{dt} = A(t)\hat{P}(t) + \hat{P}(t)A(t)^\top + \sigma(\hat{m}(t))\sigma(\hat{m}(t))^\top,$$

where $A(t) := D\mu(\hat{m}(t))$ is the Jacobian of the drift μ evaluated at $\hat{m}(t)$. We denote the moment predictions at the terminal time by

$$m_{k|k-1} = \hat{m}(t_k), \quad P_{k|k-1} = \hat{P}(t_k).$$

- (2) **Update:** Let $H_k := Dh(m_{k|k-1})$ denote the Jacobian of the observation function h evaluated at the predicted mean. Then the Kalman gain is

$$K_k = P_{k|k-1}H_k^\top (H_k P_{k|k-1}H_k^\top + R)^{-1},$$

where we recall that R is the observation noise covariance. The posterior mean and covariance become

$$m_k = m_{k|k-1} + K_k(O_k - h(m_{k|k-1})), \quad P_k = (I - K_k H_k)P_{k|k-1}.$$

The EKF thus produces a Gaussian $\mathcal{N}(x | m_k, P_k)$ approximating the filter $p_k(t_k, x)$. This distribution can be used as a closed-form approximation of the filter, but also as a proposal distribution for Monte Carlo methods such as importance sampling. This is utilized in the high-dimensional examples and further explained in Supplementary Materials SM3.

SM1.3. Ensemble Kalman filter. The ensemble Kalman filter [6] is a Monte Carlo implementation of the Kalman filter methodology, designed to make Kalman-type filtering feasible for high-dimensional, possibly nonlinear models. Rather than propagating a mean and covariance matrix, the EnKF maintains an ensemble of M particles $\{\tilde{X}_{k-1}^{(i)}\}_{i=1}^M$ that approximates the filtering distribution at time t_{k-1} . The method alternates between a prediction step, where each ensemble member is advanced forward in time, and an update step, where information from the new observation O_k is assimilated. We present here the stochastic or perturbed-observations variant of the EnKF, which preserves the posterior covariance structure in expectation [2, 11].

- (1) **Prediction:** Each ensemble member is propagated according to some forward operator F that approximates the dynamics of S , for example using the Euler–Maruyama scheme (7), to form the forecast ensemble

$$\hat{X}_k^{(i)} = F(\tilde{X}_{k-1}^{(i)}, \Delta W_k^{(i)}), \quad i = 1, \dots, M.$$

Here $\Delta W_k^{(i)}$ is the Brownian motion increment for the i 'th ensemble member at time t_k . The forecast sample mean and covariance are then computed as

$$\bar{x}_k = \frac{1}{M} \sum_{i=1}^M \hat{X}_k^{(i)}, \quad P_k^x = \frac{1}{M-1} \sum_{i=1}^M (\hat{X}_k^{(i)} - \bar{x}_k)(\hat{X}_k^{(i)} - \bar{x}_k)^\top.$$

- (2) **Update:** Draw independent perturbations $\varepsilon_k^{(i)} \sim \mathcal{N}(0, R)$ and define the predicted observations, and their sample mean, by

$$\hat{Y}_k^{(i)} = h(\hat{X}_k^{(i)}) + \varepsilon_k^{(i)}, \quad \bar{y}_k = \frac{1}{M} \sum_{i=1}^M \hat{Y}_k^{(i)}.$$

From the forecast ensemble $(\widehat{X}_k^{(i)}, \widehat{Y}_k^{(i)})_{i=1}^M$ we have the covariance matrices

$$P_k^{xy} = \frac{1}{M-1} \sum_{i=1}^M (\widehat{X}_k^{(i)} - \bar{x}_k) (\widehat{Y}_k^{(i)} - \bar{y}_k)^\top,$$

$$P_k^{yy} = \frac{1}{M-1} \sum_{i=1}^M (\widehat{Y}_k^{(i)} - \bar{y}_k) (\widehat{Y}_k^{(i)} - \bar{y}_k)^\top.$$

The Kalman gain is then

$$K_k = P_k^{xy} (P_k^{yy})^{-1}.$$

Finally, the updated ensemble is obtained by shifting each forecast member according to

$$\widetilde{X}_k^{(i)} = \widehat{X}_k^{(i)} + K_k (O_k - \widehat{Y}_k^{(i)}), \quad i = 1, \dots, M.$$

In expectation, this ensemble has the correct posterior mean and covariance under the linear-Gaussian model assumptions [2]. Iterating this procedure yields an approximate recursive filter.

We remark that the stochastic EnKF, outlined here, converges to the exact Kalman filter as $M \rightarrow \infty$ in the linear Gaussian setting [7]. In nonlinear and non-Gaussian problems, it remains a useful approximation but may underestimate uncertainty or produce biased posteriors [11]. In this work, we use the stochastic EnKF as an approximative Gaussian baseline to benchmark the proposed deep density based filters.

SM2. ADDITIONAL EXAMPLE: LINEAR SPRING-MASS

We now consider a linear spring-mass chain with $r = \frac{d}{2}$ masses connected in series by springs and dampers, disturbed by random forcing modelled by the diffusion term in (1). We define the system by the constant diffusion coefficient, $\sigma(x) = I_{d \times d}$, and the drift coefficient $\mu(x) = Ax$, where

$$A = \begin{pmatrix} 0 & I_{r \times r} \\ A_{21} & A_{22} \end{pmatrix}.$$

The stiffness and damping matrices $A_{21}, A_{22} \in \mathbb{R}^{r \times r}$ are determined by masses m_1, \dots, m_r , spring constants k_1, \dots, k_{r+1} , and damping coefficients c_1, \dots, c_{r+1} via the standard tridiagonal/diagonal structure:

$$(A_{21})_{ii} = -\frac{k_i + k_{i+1}}{m_i}, \quad (A_{21})_{i,i+1} = \frac{k_{i+1}}{m_i}, \quad (A_{21})_{i+1,i} = \frac{k_{i+1}}{m_{i+1}},$$

$$(A_{22})_{ii} = -\frac{c_i + c_{i+1}}{m_i}, \quad (A_{22})_{ij} = 0 \quad (i \neq j).$$

We use heterogeneous parameters sampled independently for each component, distributed as $m_i \sim \text{Unif}[0.8, 1.2]$, $k_i \sim \text{Unif}[0.8, 1.2]$, and $c_i \sim \text{Unif}[0.15, 0.25]$. This sampling is done once, before running the experiments, and should afterwards be seen as deterministic constants. We let $d' = r$ and define the observation process through the measurement function $h(x) = Hx$ with $H = [I_{r \times r} \quad 0_{r \times r}]$, that is, relative positions are observed but not velocities. The initial distribution is $\pi_0 = \mathcal{N}(0, I_{d \times d})$. We fix $T = 1$, $K = 10$, and $R = I_{r \times r}$. This defines a linear Gaussian state-space model, where the solution S is an Ornstein–Uhlenbeck process, and the solution to the filtering equations is provided by the Kalman filter.

We report results for $d \in \{10, 100\}$ in Figure SM1. In this setting we employ the EnKF with 10^5 ensembles, which in some metrics matches the performance of the LogBSDEF and in others outperforms it. For the 10-dimensional problem we see decent performance for the LogDSF but evidently the LogBSDEF outperforms it across all metrics, and in the 100-dimensional case the optimization for the LogDSF does not converge properly and is thus omitted.

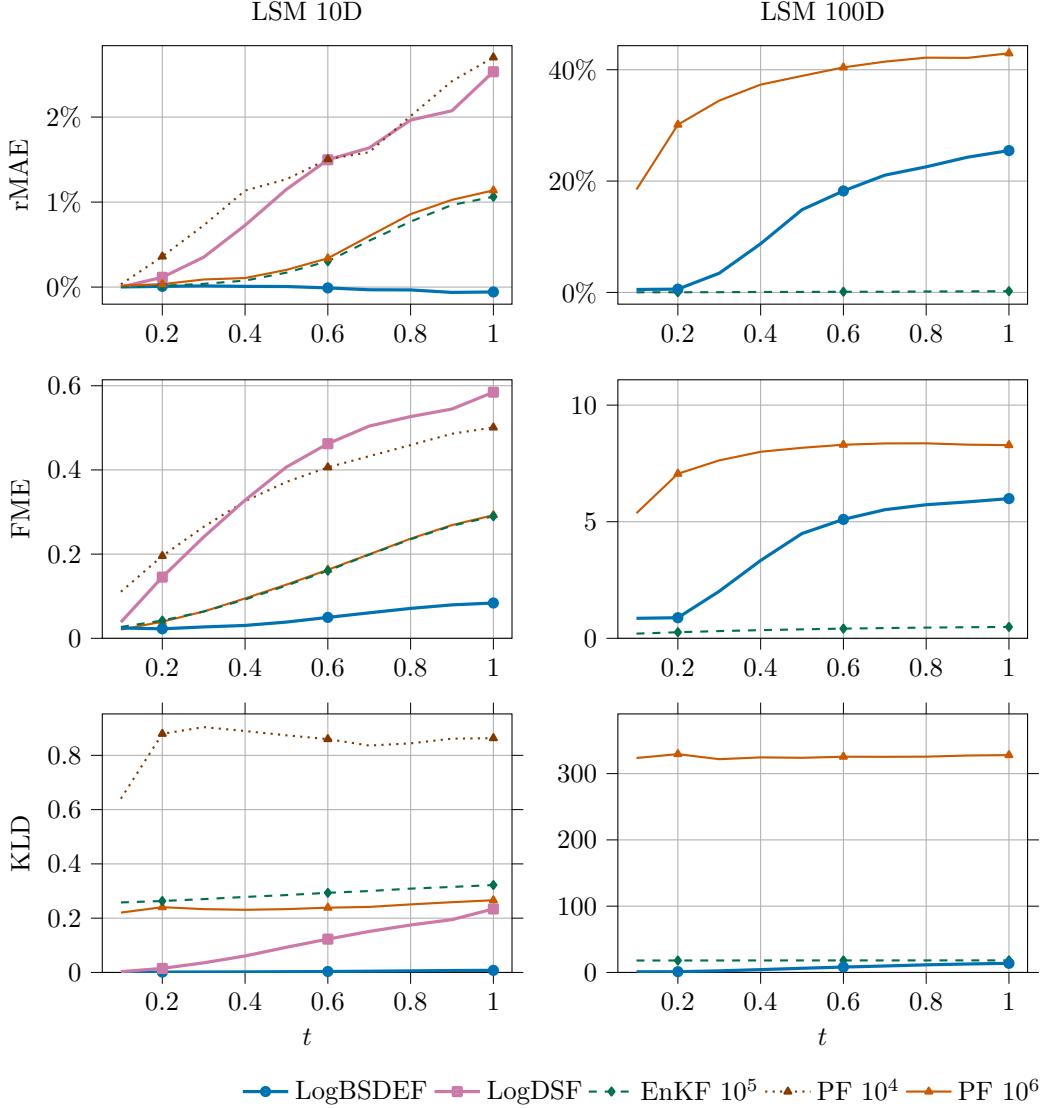


Figure SM1. On the left and right panels the results for the 10-dimensional and 100-dimensional linear spring-mass models are depicted respectively. From top to bottom the rMAE, FME, and KLD metrics are illustrated.

SM3. NORMALIZATION

We use three different normalization techniques. The first one, quadrature, is only suitable for the one-dimensional examples but is also quite effective in this setting. By defining a connected domain $D = [l, r] \subset \mathbb{R}$ so that

$$\mathbb{P}(S_t \in D, t \in [0, T]) > 1 - \epsilon,$$

for a sufficiently small ϵ , we employ simple Riemann quadrature, based on evenly spaced quadrature points $x^{(i)} \in D$, $i = 1, \dots, I$. This leads to the following quadrature approximation of the normalization constant:

$$\widehat{Z}_k^{(m)} = \frac{|r-l|}{I} \sum_{i=1}^I \widehat{p}_k(x^{(i)} | O_{1:k}^{(m)}).$$

For higher-dimensional problems, quadrature quickly becomes infeasible due to the curse of dimensionality. In this case, we use importance sampling at each observation time. Given a proposal

distribution q and $I \in \mathbb{N}$ proposal samples $\{x^{(i)}\}_{i=1}^I$ from q , we compute the Monte Carlo estimator

$$\widehat{Z}_k^{(m)} = \frac{1}{I} \sum_{i=1}^I \frac{\widehat{p}_k(x^{(i)} | O_{1:k}^{(m)})}{q(x^{(i)} | O_{1:k}^{(m)})}.$$

We consider two proposal choices:

- **EKF-based proposal:** We approximate the filtering distribution with an extended Kalman filter, using its mean and covariance for a Gaussian proposal q . This results in observation dependent proposals and requires drawing a fresh set of samples for each observation chain, typically yielding low-variance estimates.
- **Wide-tailed Gaussian proposal:** As a cheaper yet more robust alternative, we use a Gaussian proposal with inflated variance and mean equal to the unconditional mean of S_{t_k} (approximated once from simulated trajectories). Since this proposal is independent of the observation sequence, the same set of proposal samples can be reused for all observation chains, significantly reducing computational cost. However, it provides less relevant samples than the EKF-based proposal.

SM4. MODEL ARCHITECTURES

To parameterize $(\phi_{k,n})_{k=0,n=1}^{K-1,N}$ in (8) for the DSF and $(\phi_k, (\bar{v}_{k,n})_{n=0}^{N-1})_{k=0}^{K-1}$ in (9) for the BSDEF we utilize fully connected neural network architectures. In both cases, $\phi_k: \mathbb{R}^{d' \times k} \times \mathbb{R}^d \rightarrow \mathbb{R}$ represents a density. Thus we utilize the same architecture, visualized in Figure SM2, for both methods. We remark that for all time steps, regardless of the number of available observations, the input dimension is held constant and equal to $d + d'(K - 1)$. This allows us to initialize each network with the same weights as the one trained in the previous step, improving convergence and stability across all experiments. To this end, the observations are concatenated with the state value x and remaining unavailable observations are set to 0. The resulting network is rather standard, with L fully connected hidden layers of constant width and the Rectified Linear Unit (ReLU) activation. The final layer has output dimension 1 and for the logarithmic variants we use linear final activation, while the standard DSF and BSDEF have exponential activation to ensure positive output. For BSDEF and its logarithmic version, the networks $(\bar{v}_{k,n})_{n=0}^{N-1}$ have a similar architecture, where the main difference is the output layer, which has dimension d and always linear activation. For these networks representing gradients, we also typically set a lower layer width.




Figure SM2. The standard FCN architecture used in the implementation. In the figure, x denotes the state value in which the density is evaluated and $o_{1:k}$ denotes the available observations. The input is padded with zeros so that it has a constant dimension with respect to k .

In the Ornstein–Uhlenbeck example with 100 observations, we also test a model more suited for capturing long-horizon temporal dependencies, visualized in Figure SM3. The model first encodes the observation sequence (started with a token of zeros) with an LSTM network [10] and then concatenates the final cell state c and hidden state h with an embedding of the point x . This intermediate representation is then decoded using a standard FCN. Since we only apply the LSTM-based architecture for the LogBSDEF, the output layer has linear activation and either 1 or d neurons depending on whether the model represents ϕ or v . Note that the LSTM encoder is independent of x , which means it only has to be run once when evaluating multiple state values for

the same observation sequence. Moreover, to prevent an explosion of the number of parameters, all LSTM encoders for $(v_{k,n})_{n=0}^{N-1}$, with a fixed k , share weights.

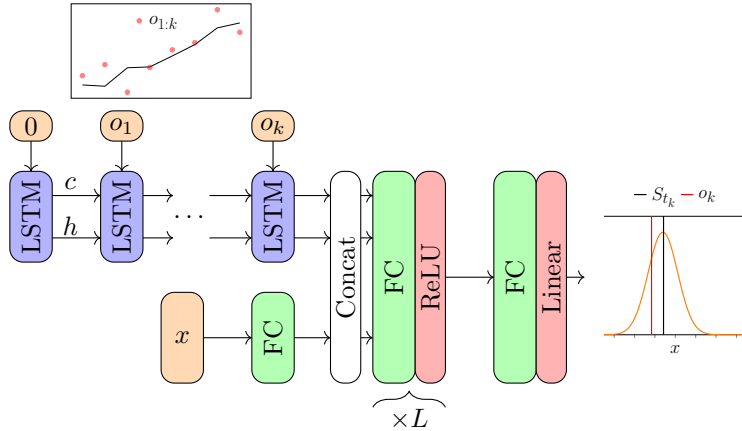


Figure SM3. The LSTM-based architecture tested in the long-horizon Ornstein–Uhlenbeck example in 10 dimensions. A token zero input $0 \in \mathbb{R}^{d'}$ is used in the first LSTM cell, while subsequent ones take as input the observation chain $o_{1:k}$.

SM5. TRAINING

SM5.1. Training the deep density methods. For DSF and LogDSF, we train in epochs over a fixed dataset: we pre-generate 20 000 mini-batches and iterate over them for a chosen number of epochs. Because the first step after an update typically requires a larger adjustment, we use a longer schedule for the update network (commonly 100 epochs) and a shorter one for the prediction network (commonly 10 epochs). This is effective since each newly introduced network is initialized from the previous one, see Supplementary Materials SM4.

For BSDEF and LogBSDEF, we instead generate each mini-batch on the fly. This reduces the statistical error compared to reusing the same data but is more computationally demanding. In the limit of sufficiently many samples, both approaches are comparable; however, on-the-fly sampling is particularly costly for DSF/LogDSF because they optimize N objective terms per observation, compared to a single objective for BSDEF/LogBSDEF under identical time discretizations.

A practical consideration is that DSF/LogDSF often use fewer time steps N than BSDE-based models, which we also found to produce more stable results. Throughout this work we choose hyperparameters per (model, example) to achieve the best accuracy. The reported settings below reflect those choices rather than an exhaustive sweep.

This is especially important to note when considering the omitted failed runs of the DSF and LogDSF for some of the high-dimensional or nonlinear examples. Exhaustive attempts were performed to find hyperparameters yielding convergence of the method. However, in some cases we did not find such parameters and this should be further investigated.

We employ early stopping by monitoring the loss (averaged per 200 iterations) and stop training if it does not decrease with a patience of P evaluations. More precisely, we start a counter $p = 0$ that increments by one every time the loss increases, and sets it to $p = 0$ when the loss decreases again. Coupled to this is the learning rate where we either use a constant rate or a cosine-annealing rate. Our application of the cosine-annealing consist of decreasing the learning rate from η_{\max} to η_{\min} over a cycle of length C following a cosine-curve:

$$\eta_c = \eta_{\min} + \frac{1}{2}(\eta_{\max} - \eta_{\min}) \left(1 + \cos\left(\pi \frac{c}{C}\right)\right),$$

where c increases every iteration when $p > \frac{P}{2}$. The hyperparameters are set to $P = 50$ and $C = 80$.

SM5.2. Training parameters. Here we report the parameters affecting the training of the models, used for each model and experiment employed in Section 5 and Supplementary Materials SM2. We recall that the BSDEF-type models (BSDEF, LogBSDEF and its LSTM variant) have one ϕ -network and N v -networks, while the DSF-type models (DSF, LogDSF) only have one ϕ -network. In Table SM1 we report the widths of the networks by Width_ϕ and Width_v , respectively. In all the examples except for the LSTM networks we use three hidden layers. We abbreviate the normalization methods of Supplementary Materials SM3 with I-EKF and I-G for the EKF-based and wide-tailed Gaussian proposals respectively.

Regarding the LSTM-based model in the long-horizon OU example of Section 5.2, it has 4 hidden layers in the FCN decoder with constant width 512 for the ϕ -network and 128 for the v -networks. The embedding layer for the state value x has the same width as the hidden layers in the decoder and the LSTM encoder has hidden dimension 256 and 128 in the ϕ - and v -cases respectively. The LSTM encoder also has 2 extra cells without any input between each observation cell. The other hyperparameters (N , B , LR, Normalization method) are the same as the FCN for the same example, see Table SM1.

Table SM1. Training setup across models and examples. Models are abbreviated B - BSDEF, LB - LogBSDEF, D - DSF, LD - LogDSF. Columns: time steps N , widths (W_ϕ , W_v) for the ϕ - and v -networks (3 hidden layers fixed), batch size B , LR schedule (start \rightarrow end), and normalization (method and number of normalized trajectories per batch).

Model	Example	N	W_ϕ	W_v	B	LR schedule	Norm
<i>BSDEF models</i>							
B	OU (1d)	64	128	32	512	Const 10^{-4}	(Quad, 64)
B	Bistable	64	128	32	512	Const 10^{-4}	(Quad, 64)
LB	OU (1d)	64	128	32	512	Const 10^{-4}	(Quad, 64)
LB	Bistable	64	128	32	512	Const 10^{-4}	(Quad, 64)
LB	Schlögl	128	256	64	4096	Cos $10^{-4} \rightarrow 10^{-6}$	(Quad, 128)
LB	OU (10d)	16	512	256	1024	Cos $10^{-4} \rightarrow 10^{-5}$	(I-EKF, 64)
LB	OU (100d)	32	1024	512	2048	Cos $10^{-4} \rightarrow 10^{-5}$	(I-EKF, 64)
LB	LSM (10d)	32	512	256	2048	Cos $10^{-4} \rightarrow 10^{-5}$	(I-EKF, 64)
LB	LSM (100d)	32	2048	512	4096	Cos $10^{-4} \rightarrow 10^{-5}$	(I-EKF, 512)
LB	L96 (4d)	32	256	64	2048	Const 10^{-4}	(I-G, 64)
LB	L96 (10d)	32	512	256	2048	Const 10^{-4}	(I-G, 64)
LB	L96 (20d)	32	512	256	2048	Const 10^{-4}	(I-G, 64)
LB	L96 (40d)	32	512	256	2048	Const 10^{-4}	(I-G, 64)
LB	L96 (100d)	32	1024	512	2048	Const 10^{-4}	(I-G, 64)
<i>DSF models</i>							
D	OU (1d)	16	128	N/A	2048	Const 10^{-4}	(Quad, 64)
D	Bistable	16	128	N/A	2048	Const 10^{-4}	(Quad, 64)
LD	OU (1d)	16	128	N/A	2048	Const 10^{-4}	(Quad, 64)
LD	Bistable	16	128	N/A	2048	Const 10^{-4}	(Quad, 64)
LD	OU (10d)	16	512	N/A	1024	Const 10^{-4}	(I-EKF, 64)
LD	LSM (10d)	32	512	N/A	2048	Const 10^{-4}	(I-EKF, 64)

SM6. EVALUATION

In this section we briefly detail the parameters used for evaluating the metrics of Section 4. To evaluate the probability density of the particle filters and ensemble Kalman filters, both for the benchmark methods and the reference solution, we employ Gaussian kernel density estimators, implemented using `scipy.stats.gaussian_kde` [12]. This function applies Scott's rule for bandwidth

selection by default [14, 15], and we perform no additional tuning of the bandwidth parameter. In Table SM2 we report the evaluation parameters used for all the experiments.

Table SM2. Evaluation setup across examples. Columns: number of Monte Carlo evaluation samples M , normalization method, reference solution, reference discretization steps N_{ref} . For the Lorenz-96 high-dimensional example we do not have access to a reference solution and only evaluate the MAE and NLL metrics.



Example	M	Normalization method	Reference	N_{ref}
OU (1d)	10^4	Quad	KF	512
Bistable	10^4	Quad	PF 10^6	128
OU (10d)	10^3	I-EKF	KF	128
OU (100d)	10^3	I-EKF	KF	128
LSM (10d)	10^3	I-EKF	KF	128
LSM (100d)	10^3	I-EKF	KF	128
Schlögl	10^4	Quad	PF 10^6	128
Lorenz-96 (4d)	10^3	I-G	PF 10^6	128
<i>Only MAE and NLL</i>				
Lorenz-96 (4d)	10^4	I-G	N/A	N/A
Lorenz-96 (10d)	10^4	I-G	N/A	N/A
Lorenz-96 (20d)	10^4	I-G	N/A	N/A
Lorenz-96 (40d)	10^4	I-G	N/A	N/A
Lorenz-96 (100d)	10^4	I-G	N/A	N/A

SM7. ADDITIONAL FIGURES

In this section we present qualitative results: example trajectories and density comparisons. Figure SM4 shows a trajectory and the final-time density for the long-horizon 10-dimensional Ornstein–Uhlenbeck model with some different methods. Figures SM5 and SM6 analogously show trajectories and densities for the 10-dimensional linear spring-mass model. Finally, Figure SM7 shows one trajectory of the four-dimensional Lorenz-96 model in the x_1x_2 - and x_3x_4 -marginal planes.

SM8. THE CHEMICAL SCHLÖGL MODEL

The Schlögl model [13, 16, 17] is a prototypical bistable chemical reaction network involving a single reacting species S coupled to two external reservoirs A and B . Its reactions are



Here, the notation $X \xrightarrow{\theta_n} Y$ means that configuration X is transformed into configuration Y with rate constant θ_n according to mass-action kinetics.

Reactions in (SM8.1) describe an autocatalytic feedback loop: two molecules of S react with a molecule from reservoir A to produce a third S , while the reverse reaction consumes three S to regenerate A . Reactions in (SM8.2) model inflow and outflow of S through coupling with reservoir B , maintaining a constant supply and removal of the species.

Starting from the chemical master equation and applying the chemical Langevin approximation [8], one obtains the SDE, given by the dynamics detailed in Section 5.3, governing the concentration of S . Each reaction channel contributes an independent noise source $(B^{(i)})_{i=1}^4$, corresponding to intrinsic stochastic fluctuations of the reaction events. The nonlinear drift encodes autocatalytic growth and saturation, while the multiplicative diffusion terms capture random reaction noise.

For parameter values

$$\theta = (3 \times 10^{-7}, 10^{-4}, 10^{-3}, 3.5), \quad A = 10^5, \quad B = 2 \times 10^5,$$

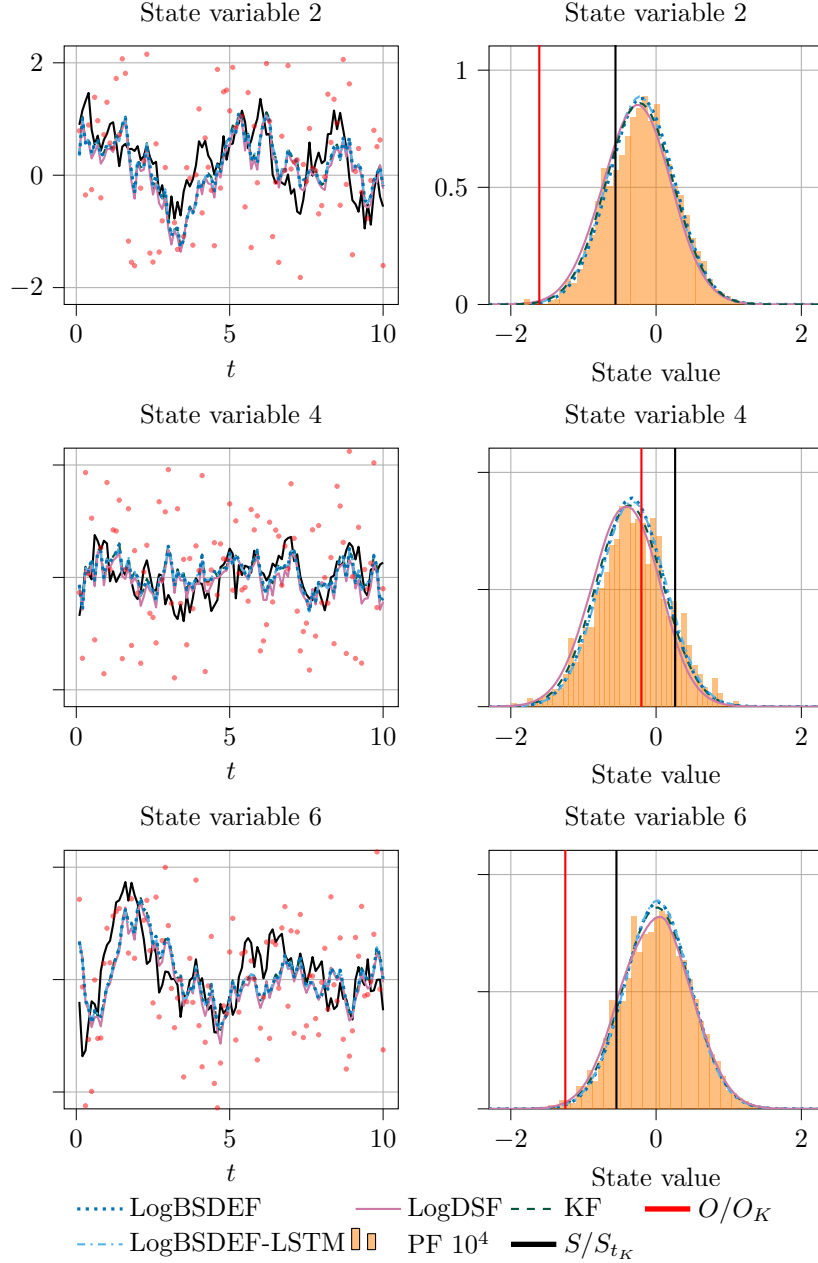


Figure SM4. One trajectory for the long-horizon Ornstein–Uhlenbeck problem in 10 dimensions. The left column shows the sample path with corresponding filter mean estimates. The right column displays the filtering densities at the final time $T = 10$.

the deterministic mean-field limit exhibits two stable steady states, separated by an unstable equilibrium. Random fluctuations can drive rare transitions between these states, a hallmark of bistable systems. This bistability makes the Schlögl model a canonical benchmark in stochastic chemical kinetics and nonlinear filtering of reaction networks.

REFERENCES

- [1] P. Bickel, B. Li, and T. Bengtsson. Sharp failure rates for the bootstrap particle filter in high dimensions. *Pushing the limits of contemporary statistics: contributions in honor of Jayanta K. Ghosh*. Vol. 3. Inst. Math. Stat. (IMS) Collect. Inst. Math. Statist., Beachwood, OH, 2008, pp. 318–329.
- [2] G. Burgers, P. J. van Leeuwen, and G. Evensen. Analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.* 126.6 (1998), pp. 1719–1724.

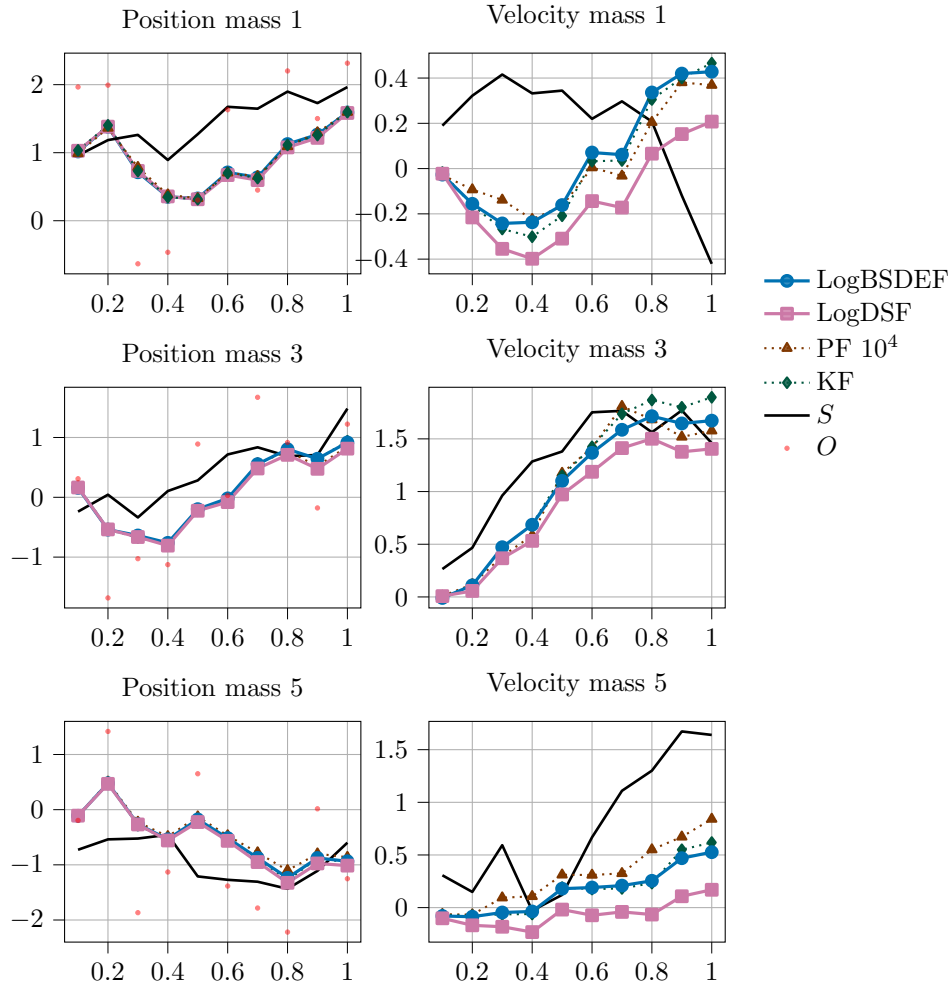


Figure SM5. One trajectory and filter estimates for selected components in the 10-dimensional linear spring-mass example.

- [3] N. Chopin. Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Ann. Statist.* 32.6 (2004), pp. 2385–2411.
- [4] D. Crisan and A. Doucet. A survey of convergence results on particle filtering methods for practitioners. *IEEE Trans. Signal Process.* 50.3 (2002), pp. 736–746.
- [5] P. Del Moral. *Feynman-Kac formulae: Genealogical and interacting particle systems with applications*. Springer-Verlag, New York, 2004, pp. xviii+555.
- [6] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* 99.C5 (1994), pp. 10143–10162.
- [7] G. Evensen. The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.* 53.4 (2003), pp. 343–367.
- [8] D. T. Gillespie. The chemical Langevin equation. *J. Chem. Phys.* 113.1 (2000), pp. 297–306.
- [9] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEEE Proceedings F (Radar and Signal Processing)* 140.2 (1993), pp. 107–113.
- [10] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation* 9.8 (1997), pp. 1735–1780.
- [11] M. Katzfuss, J. R. Stroud, and C. K. Wikle. Understanding the ensemble Kalman filter. *Amer. Statist.* 70.4 (2016), pp. 350–357.
- [12] P. Virtanen, et al. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods* 17.3 (2020), pp. 261–272.
- [13] F. Schlögl. Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift für physik* 253.2 (1972), pp. 147–161.
- [14] D. W. Scott. *Multivariate density estimation*. John Wiley & Sons, Inc., New York, 1992, pp. xiv+317.
- [15] B. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, 1986.
- [16] M. Vellela and H. Qian. Stochastic dynamics and non-equilibrium thermodynamics of a bistable chemical system: the Schlögl model revisited. *J. R. Soc. Interface* 6.39 (2009), pp. 925–940.
- [17] M. Vlysidis and Y. N. Kaznessis. Solving stochastic reaction networks with maximum entropy Lagrange multipliers. *Entropy* 20.9 (2018), p. 678.

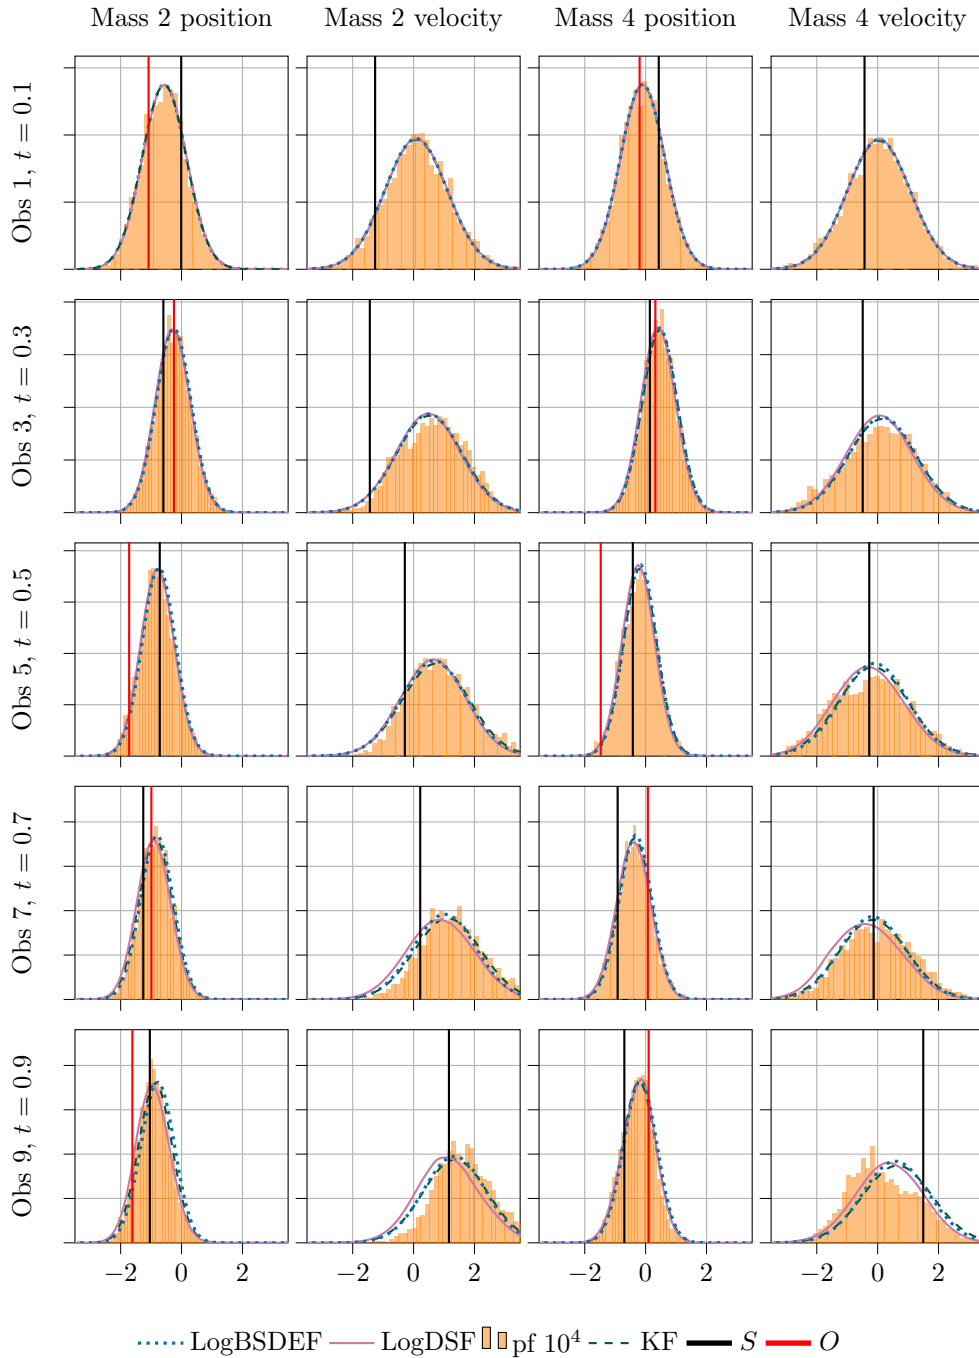


Figure SM6. Marginal densities over time in the 10-dimensional linear spring-mass example for selected components. Recall that only positions are observed.

KASPER BÄGMARK, DEPARTMENT OF MATHEMATICAL SCIENCES, CHALMERS UNIVERSITY OF TECHNOLOGY AND UNIVERSITY OF GOTHENBURG, SE-412 96 GOTHENBURG, SWEDEN

Email address: bagmark@chalmers.se

FILIP RYDIN, DEPARTMENT OF ELECTRICAL ENGINEERING, CHALMERS UNIVERSITY OF TECHNOLOGY, SE-412 96 GOTHENBURG, SWEDEN

Email address: filipry@chalmers.se

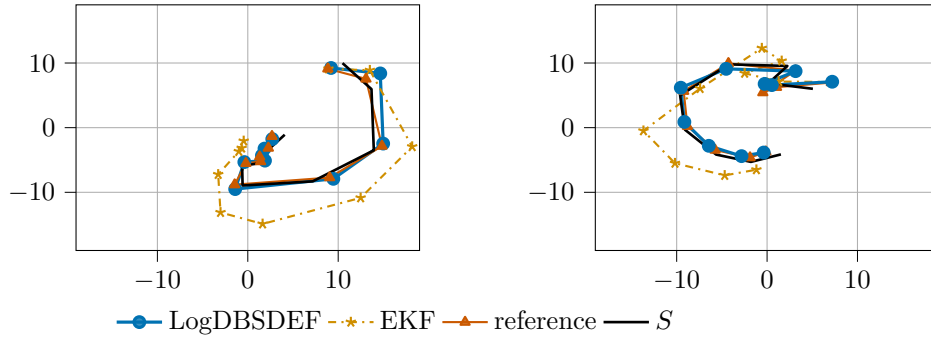


Figure SM7. One realization of the state process S , starting at $S_0 \sim \mathcal{N}((8, 8, 8, 8), I)$, in the four-dimensional Lorenz-96 example together with the corresponding filtering means from the different methods. The left and right panels illustrate the first two and last two variables respectively.