

Variational Gaussian filtering via Wasserstein gradient flows

Adrien Corenflos*, Hany Abdulsamad†

Department of Electrical Engineering and Automation
Aalto University, Finland

Email: *adrien.corenflos@aalto.fi, †hany.abdulsamad@aalto.fi

Abstract—We present a novel approach to approximate Gaussian and mixture-of-Gaussians filtering. Our method relies on a variational approximation via a gradient-flow representation. The gradient flow is derived from a Kullback–Leibler discrepancy minimization on the space of probability distributions equipped with the Wasserstein metric. We outline the general method and show its competitiveness in posterior representation and parameter estimation on two state-space models for which Gaussian approximations typically fail: systems with multiplicative noise and multi-modal state distributions.

Index Terms—Kalman filtering, variational inference, state-space models, Wasserstein gradient flow.

I. INTRODUCTION

State-space models (or hidden Markov models) are a class of models widely used to represent latent dynamics that are partially or indirectly observed. They typically arise in ecological, economical, and tracking applications [for an introduction, see, e.g. 1]. Formally, state-space models are given by a set of dynamics and noisy observations, often depending on a set of parameters θ

$$\begin{aligned} X_0 &\sim p_0(\cdot | \theta), \\ X_{k+1} &\sim p_k(\cdot | X_k, \theta), \\ Y_k &\sim h_k(\cdot | X_k, \theta). \end{aligned} \quad (1)$$

While the problem of inference in such models is generally intractable, computing the filtering distribution $p(x_k | y_{0:k})$ can typically be done exactly if the state-space is finite (x_k can only take a finite number of values) or when all the (conditional) densities in (1) are Gaussian using the celebrated Kalman filter [2]. When this is not the case, relying on approximations becomes necessary. Two important types of approximations are approximate Gaussian filters [see, e.g. 3], and Monte Carlo filters [see, e.g. 4].

Although the standard filtering problem is important, one may also be interested in system identification, which, in the parametric context, refers to learning θ from a sequence of observations $\{y_0, \dots, y_K\}$.

In this article, we pay particular attention to two classes of models, typically neglected in approximate Gaussian filtering. The first class is that of models with multiplicative noise, for which stochastic volatility models are an illustrative example,

Adrien Corenflos is funded by the Academy of Finland project 321891 (ADAFUME). Hany Abdulsamad is funded by the Finnish Center for Artificial Intelligence (FCAI).

often used in economics to model financial returns [5]. These are usually given as an auto-regressive latent state x_k , and observations y_k following

$$\begin{aligned} Y_k &= \exp(X_k/2) \eta_k, \\ X_{k+1} &= \mu + \alpha(X_k - \mu) + \sigma \epsilon_k, \end{aligned} \quad (2)$$

where the noise processes are correlated

$$\begin{pmatrix} \epsilon_k \\ \eta_k \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right). \quad (3)$$

The second class we consider are systems for which the state (filtering) posterior is a multi-modal distribution. A simple example of this form can be given by constructing a latent state x_k with random walk dynamics while the observations y_k are a function of the absolute value of x_k :

$$\begin{aligned} X_{k+1} &= X_k + \epsilon_k, & \epsilon_k &\sim \mathcal{N}(0, 1), \\ Y_k &= |X_k| + \eta_k, & \eta_k &\sim \mathcal{N}(0, 1). \end{aligned} \quad (4)$$

If, for example, $X_0 \sim \mathcal{N}(0, \delta^2)$, it is straightforward to see that the state filtering distribution will be bi-modal and fully symmetric with respect to the x-axis.

A. Contributions

Existing approximate Gaussian filtering methods suffer from several drawbacks. The linearization methods of [6, 7] for example require computing conditional expectations $m_k(x) = \mathbb{E}[Y_k | X_k = x]$. In the case of the stochastic volatility model (2), this quantity will unequivocally be null (at least for $\rho = 0$, see Section III-A for details). Consequently, applying these methods to (2) will lead to a filtering (or a smoothing) solution independent of the gathered observations, which is problematic.

Additionally, classical linearization methods do not extend directly to multi-modal distributions, and they do not, to the best of our knowledge, enable handling mixtures of Gaussians. In view of this, our contributions are the following:

- 1) We rephrase the filtering problem as an iterative distribution fitting problem.
- 2) We apply the variational inference method from [8] to propagate Gaussian approximations between time steps and formulate it as a fixed point iteration for efficient gradient calculation.
- 3) We use our method for parameter estimation in stochastic volatility models and filtering of multi-model target distributions.

II. METHODOLOGY

A. Variational Inference via Wasserstein Gradient-flows

Let $\pi(x) \propto \exp(-V(x))$ be an arbitrary target distribution known up to a normalizing constant. Given a variational family of distributions, $q_\phi(x)$, $\phi \in \Phi$, and a measure of discrepancy $L(\phi) = D(\pi, q_\phi)$ between π and q_ϕ , it is natural to try and approximate π with q_ϕ by minimizing L .

A typical discrepancy is given by the Kullback–Leibler (KL) divergence [9]

$$L(\phi) = \text{KL}(q_\phi | \pi) := \int q_\phi(x) \log \frac{q_\phi(x)}{\pi(x)} dx. \quad (5)$$

This divergence presents a number of attractive properties for statistical inference: (i) it is positive, (ii) it only requires evaluating $V(x)$ and does not necessitate knowing the normalizing constant of π , and (iii), it is exact in the sense that $\text{KL}(q_\phi | \pi) = 0$ if and only if $q_\phi = \pi$. For more details, we refer the reader to [10].

For example, if $q_\phi(x) = \text{N}(x | \mu, \Sigma)$ is in the family of well-defined Gaussians, parameterized by their mean and covariance, then the inference problem can be cast as a minimization of $L(\cdot)$ with respect to (μ, Σ) . Furthermore, when ignoring the positivity constraints on Σ , it is possible to define a gradient flow on $\phi = (\mu, \Sigma)$, akin to a gradient descent on $L(\cdot)$ in continuous time

$$\frac{d\phi_t}{dt} = -\nabla_\phi L(\phi_t), \quad (6)$$

which, under an assumption of convexity, will converge to the minimizer of the objective $L(\cdot)$ [11].

While this procedure is correct in essence, it targets the problem indirectly by first assuming an arbitrary parameterization of the model which may or may not respect convexity. A more direct approach is to fit q_ϕ to π in terms of a minimization problem over the space of probability distributions, where we want to minimize $L(q) = D(\pi, q)$. In this case, it is possible to define an analog to the gradient flow (6) by equipping the space of probability distributions with the Wasserstein distance [12, Chap. 6]. Under this metric, we can define a trajectory of probability distributions $q_t(x) \in \mathcal{P}(\mathbb{R}^d)$ via a partial differential equation

$$\frac{\partial q_t(x)}{\partial t} = \nabla \cdot \left[q_t(x) \nabla \log \frac{q_t(x)}{\pi(x)} \right], \quad (7)$$

where $\nabla \cdot$ is the divergence operator, expressed in Euclidean coordinates.

Interestingly, by restricting q_t to represent a Gaussian distribution, it was shown in [8], following [13], that (7) can be reformulated into coupled ordinary differential equations (ODEs) on the mean μ_t and covariance Σ_t of q_t

$$\begin{aligned} \frac{d\mu_t}{dt} &= -\mathbb{E}[\nabla V(Z_t)] \\ \frac{d\Sigma_t}{dt} &= 2I - \mathbb{E}[\nabla V(Z_t) \otimes (Z_t - \mu_t)] \\ &\quad - \mathbb{E}[(Z_t - \mu_t) \otimes \nabla V(Z_t)], \end{aligned} \quad (8)$$

where I is the identity matrix with dimensions $d \times d$ and $Z_t \sim \text{N}(\mu_t, \Sigma_t)$ is Gaussian. Now, provided that we can compute

or approximate the expectations arising in (8), we can find a minimizer $q^*(x) = \text{N}(x | m, P)$ of $L(\cdot)$ by integrating these coupled ODEs until convergence.

B. Filtering as Variational Inference

The filtering problem involves inferring the posterior distribution $p(x_k | y_{0:k})$ for each time step k . To do so, it is often possible to rely on the familiar *innovation-prediction* decomposition [3, Chap. 4]

$$\begin{aligned} p(x_k | y_{0:k}) &\propto p(y_k | x_k) p(x_k | y_{0:k-1}), \\ p(x_k | y_{0:k-1}) &= \int p(x_k | x_{k-1}) p(x_{k-1} | y_{0:k-1}) dx_{k-1}. \end{aligned} \quad (9)$$

For simplicity, we will focus on state-space models with affine Gaussian transition models

$$p(x_k | x_{k-1}) = \text{N}(x_k | A_{k-1} x_{k-1} + b_{k-1}, Q_{k-1}), \quad (10)$$

covering those previously presented in (2) and (4).

In that case, assuming a Gaussian approximation is available for the previous time step $k-1$,

$$p(x_{k-1} | y_{0:k-1}) \approx \text{N}(x_{k-1} | m_{k-1}, P_{k-1}), \quad (11)$$

then the *prediction* step of the filter leads to another Gaussian distribution $p(x_k | y_{0:k-1}) \approx \text{N}(x_k | \bar{m}_k, \bar{P}_k)$ with mean and covariance

$$\begin{aligned} \bar{m}_k &= A_{k-1} m_{k-1} + b_{k-1}, \\ \bar{P}_k &= A_{k-1} P_{k-1} A_{k-1}^\top + Q_{k-1}. \end{aligned} \quad (12)$$

Given this predictive distribution, the *innovation* step computes the approximate filtering distribution

$$p(x_k | y_{0:k}) \propto \exp(-V(x_k)), \quad (13)$$

where the potential function $V(x_k)$ is given as

$$V(x_k) = -\log p(y_k | x_k) - \log \text{N}(x_k | \bar{m}_k, \bar{P}_k). \quad (14)$$

Consequently, in order to find the parameters of a variational Gaussian $\text{N}(x_k | m_k, P_k)$ that approximates the posterior $p(x_k | y_{0:k})$, it suffices to follow the recipe from Section II-A and integrate (8) up to stationarity, starting from the predictive parameters (\bar{m}_k, \bar{P}_k) , or any other approximation of (m_k, P_k) .

When the transition $p(x_k | x_{k-1})$ is not Gaussian, it is possible to use a similar variational approach to propagate the Gaussian approximation $p(x_{k-1} | y_{0:k-1}) \approx \text{N}(x_{k-1} | m_{k-1}, P_{k-1})$ through the non-Gaussian dynamics and approximate the marginal $p(x_k | y_{0:k-1})$ with another Gaussian $\text{N}(x_k | \bar{m}_k, \bar{P}_k)$. This method was used in [13] to propagate the Gaussian approximation through dynamics defined by a stochastic differential equation. Combining this with our approach is, therefore, *de facto* possible. However, due to the limited scope, we leave this aspect for future work and only consider Gaussian dynamics in the remainder of this article.

Finally, because the likelihood of the observations is given by $p(y_{0:k}) = p(y_{0:k-1}) \int p(y_k | x_k) p(x_k | y_{0:k-1}) dx_k$, it is easy to derive an approximation of the marginal log-likelihood of the model by recursion. That is because the quantity $\ell_k = \int p(y_k | x_k) p(x_k | y_{0:k-1}) dx_k$ is evaluated as part of (8),

and can therefore be used to provide an online estimate of the log-likelihood increments, which, in turn, can be used, for example, in a system identification scenario. We return to this point in Sections II-D and III-A.

C. The Multi-modal Filtering Case

To generalize the variational technique from the previous section, let us suppose that the filtering distribution at time $k - 1$ is instead given by a mixture density

$$p(x_{k-1} | y_{0:k-1}) = \sum_{i=1}^{N_i} w^{(i)} N(x_{k-1} | m_{k-1}^{(i)}, P_{k-1}^{(i)}), \quad (15)$$

with $w^{(i)} = 1/N_i, \forall i \in [1, N_i]$. In this case, when the transition dynamics are Gaussian, it is easy to show that

$$p(x_k | y_{0:k-1}) = \sum_{i=1}^{N_i} w^{(i)} N(x_k | \bar{m}_k^{(i)}, \bar{P}_k^{(i)}), \quad (16)$$

where $\forall i \in [1, N_i]$

$$\begin{aligned} \bar{m}_k^{(i)} &= A_{k-1} m_{k-1}^{(i)} + b_{k-1}, \\ \bar{P}_k^{(i)} &= A_{k-1} P_{k-1}^{(i)} A_{k-1}^\top + Q_{k-1}. \end{aligned} \quad (17)$$

These updates correspond to a tractable *prediction* step. As a consequence, we only need to understand how to perform the *innovation* step to arrive at the posterior $p(x_k | y_{0:k})$.

Conveniently, [8] also shows that the duality between the gradient flow (7) and the coupled ODEs (8) extends to the case of the finite variational mixture of Gaussians

$$q_t(x) = \sum_{i=1}^{N_i} \omega^{(i)} q_t^{(i)}(x) = \sum_{i=1}^{N_i} \omega^{(i)} N(x | \mu_t^{(i)}, \Sigma_t^{(i)}), \quad (18)$$

with $\omega^{(i)} = 1/N_i, \forall i \in [1, N_i]$. In this case, rather than a pair of ODEs, we obtain a system of such ODEs

$$\begin{aligned} \frac{d\mu_t^{(i)}}{dt} &= -\mathbb{E} \left[\nabla \log \frac{q_t}{\pi} (Z_t^{(i)}) \right] \\ \frac{d\Sigma_t^{(i)}}{dt} &= -\mathbb{E} \left[\nabla^2 \log \frac{q_t}{\pi} (Z_t^{(i)}) \right] \Sigma_t^{(i)} \\ &\quad - \Sigma_t^{(i)} \mathbb{E} \left[\nabla^2 \log \frac{q_t}{\pi} (Z_t^{(i)}) \right], \end{aligned} \quad (19)$$

where, for all i , $Z_t^{(i)} \sim N(\mu_t^{(i)}, \Sigma_t^{(i)})$ is Gaussian, and ∇^2 denotes the Hessian operator.

This result means that, given a Gaussian mixture approximation of the predictive $p(x_k | y_{1:k-1})$, we can obtain a Gaussian mixture approximation of $p(x_k | y_{0:k})$ by integrating (19) and following a similar approach to Section II-B.

D. Numerical Considerations and Implementation

In practice, the integrals arising in (8) and (19) are not available in closed form. Therefore, we need to resort to numerical integration. This can be done by using any form of deterministic or stochastic Gaussian quadrature, for example, Monte Carlo [see, e.g., 14] or sigma-points [3, see, e.g.,] methods. The two approaches have their pros and cons: Monte Carlo will give the correct solution on average, while

Algorithm 1 Uni-modal Wasserstein Gradient-flow Filter

- 1: **input:** Measurements $y_{0:K}$ and prior (m_0, P_0)
 - 2: **output:** Filtering posterior distributions $(m_{1:K}, P_{1:K})$ and marginal log-likelihood $\ell = \log p(y_{0:K})$
 - 3: Set $\ell \leftarrow 0$
 - 4: **for** $k \leftarrow 1$ **to** K **do**
 - 5: Set $\ell_k \leftarrow \log \mathbb{E}_{x \sim N(\bar{m}_k, \bar{P}_k)} p(y_k | x_k)$
 - 6: Set $\ell \leftarrow \ell + \ell_k$ ▷ Log-likelihood
 - 7: $(m_k, P_k) \leftarrow (\bar{m}_k, \bar{P}_k)$
 - 8: **while** Not converged **do** ▷ Innovation
 - 9: $(m_k, P_k) \leftarrow I(m_k, P_k)$
 - 10: **end while**
 - 11: Set $\bar{m}_{k+1} \leftarrow A_k m_k + b_k$ ▷ Prediction
 - 12: Set $\bar{P}_{k+1} \leftarrow A_k P_k A_k^\top + Q_k$
 - 13: **end for**
 - 14: **return** $(m_{1:K}, P_{1:K}), \ell$
-

deterministic quadrature is bound to be biased. Nonetheless, deterministic rules are sometimes more practical when no stochasticity should be allowed in the system. Whichever is chosen, we will obtain an approximation

$$\frac{d\mu_t}{dt} \approx F_m(\mu_t, \Sigma_t), \quad \frac{d\Sigma_t}{dt} \approx F_P(\mu_t, \Sigma_t), \quad (20)$$

of (8) for the choice of $V(x_k)$ given by (14). Moreover, the same approximation scheme can be used to compute log-likelihood increments $\log \mathbb{E} [p(y_k | x_k)]$ under the predictive $N(x_k | \bar{m}_k, \bar{P}_k)$. In Section III, we use Gauss–Hermite [see, e.g., 3, Chap. 5] quadrature integration rules with order five.

We now assume that we have chosen an integration method I for which the stationary solution of (20) is a fixed point: $(m_k, P_k) = I(m_k, P_k)$. This can, for example, be an Euler integration step with a small step size. This fixed-point perspective allows us to leverage the implicit function theorem and bypass the loop when computing gradients for system identification. For the sake of brevity, we omit the details here and refer to [15] for the method and to our code¹ for a Python implementation.

Therefore, the final algorithm for variational uni-modal filtering from Section II-B is given by Algorithm 1. Due to space constraints, we do not reproduce the algorithm for the multi-modal case from Section II-C: the procedure follows the same steps, albeit for a larger system of ODEs.

III. EMPIRICAL EVALUATION

As discussed in Section I, introducing our method is motivated by the problems posed by multiplicative noise and multimodality in approximate Gaussian filtering. Consequently, we aim to demonstrate its effectiveness on these two problems. We compare the variational Wasserstein filter (VWF) to a bootstrap particle filter (PF, [16]) using 500 particles and the *continuous* resampling scheme from [17]. This resampling allows using the particle filter in a standard parameter estimation

¹Implementation of the fixed-point iteration is available under https://github.com/hanyas/wasserstein-flow-filter/blob/master/wasserstein_filter/utlis.py

scenario. The code to reproduce these experiments can be found at <https://github.com/hanyas/wasserstein-flow-filter>.

A. Stochastic Volatility with Leverage

First, we consider the stochastic volatility model as given by (2) and (3). Because the noises ϵ_k and η_k are correlated, we need to construct the joint distribution over X_k and its generating noise ϵ_k and form the augmented two-dimensional state $\zeta_k = [X_k \ \epsilon_k]^\top$. The dynamics of ζ_k are then given by

$$\begin{pmatrix} X_{k+1} \\ \epsilon_{k+1} \end{pmatrix} = \begin{pmatrix} \alpha & \sigma \\ 0 & 0 \end{pmatrix} \begin{pmatrix} X_k \\ \epsilon_k \end{pmatrix} + \begin{pmatrix} \mu(1-\alpha) \\ 0 \end{pmatrix} + q_k \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad (21)$$

where q_k is a standard Gaussian random variable. With these dynamics, we have the following observation model

$$Y_k = \exp(X_k/2)(\rho\epsilon_k + \sqrt{1-\rho^2}r_k), \quad (22)$$

with r_k being a standard Gaussian random variable. Contrary to the original form of the model, the noise processes are now de-correlated so that we can apply our method to the 2-dimensional system defined by ζ_k .

To perform an empirical comparison, we simulate the model and record three trajectories with $K = 1000$, $K = 1500$, and $K = 2000$ observations. We use parameters $\alpha_* = 0.975$, $\mu_* = 0.5$, $\sigma_*^2 = 0.02$, and $\rho_* = -0.8$. These values correspond to those in [17], typical of a standard stock market.

To illustrate the problem of using methods based on linearizing the conditional observation mean $\mathbb{E}[Y_k | X_k, \epsilon_k]$, we evaluate the marginal log-likelihood of the data approximated by an extended Kalman filter (EKF) targeting (21) with parameters $(\alpha_*, \mu_*, \sigma_*)$ while varying levels of correlation ρ and number of observations K . The resulting curves for the particle filter of [17], our method, and the extended Kalman filter are shown in Figure 1. As the (model) correlation ρ between the two noise-generating processes decreases, the predictive value of the observations becomes negligible. As a consequence, the marginal likelihood, as approximated by the EKF, will struggle to capture the right level of correlation. This means that calibrating the model using an extended (or similar) Kalman filter will result in an inconsistent estimate, at least for the parameter ρ .

To further confirm this hypothesis, we follow [17] and perform joint maximum likelihood parameter estimation given $K = 1000$ observations. We compare our method to theirs and to an extended Kalman filter. The likelihood is optimized in all cases using the gradient of the log-likelihood, delivered by automatic reverse differentiation. In the particular case of the variational Wasserstein filter, differentiation is made efficient by our fixed point formulation of the variational calibration.

For a statistical comparison, we repeat this experiment over ten trials associated with different random seeds and, therefore, ten distinct sets of observations. We report the mean plus or minus one standard deviation of the parameter estimates for each algorithm. However, because the particle filter is inherently a stochastic method, even for a fixed set of observations, we perform 25 *sub-trials* per set of observations and use the median as a parameter estimate per trial. The




Fig. 1. Comparing the (normalized) marginal log-likelihood as a function of the leverage parameter ρ in a stochastic volatility model. We plot the estimates as returned by an extended Kalman filter (red), a bootstrap particle filter (blue), and a variational Wasserstein filter (green). The dotted, dashed, and solid lines correspond to different numbers of observations $K = 1000$, $K = 1500$, and $K = 2000$. The true value is $\rho = -0.8$ (vertical line). The variational Wasserstein and particle filters deliver consistent approximations independent of the data size, while the extended Kalman filter does not.

TABLE I
STATISTICS FOR PARAMETER ESTIMATION OF THE STOCHASTIC VOLATILITY MODEL. RESULTS ARE AVERAGED OVER TEN TRIALS.

	μ	α	σ	ρ
True	0.50	0.975	0.14	-0.80
PF [17]	0.56 (± 0.08)	0.972 (± 0.007)	0.15 (± 0.02)	-0.85 (± 0.06)
VWF	0.56 (± 0.07)	0.972 (± 0.009)	0.15 (± 0.02)	-0.80 (± 0.04)
EKF	0.69 (± 0.33)	0.780 (± 0.590)	0.21 (± 0.25)	-0.58 (± 0.54)

results are given in Table I and confirm our intuition. The extended Kalman filter provides inconsistent solutions, while our method delivers estimates comparable to the estimates of the particle filter.

B. Multi-Modal Example

We now turn to the multi-modal motivating example (4). We simulate $K = 500$ observations from the model, where we have taken the variance at origin to be $\delta^2 = 1$. We then perform variational filtering using the mixture version of Algorithm 1 with $N_i = 2$ mixture components.

Our goal is a correct representation of the filtering posterior. Thus we will not assess the performance in terms of root mean square error, which is inappropriate for such problems. Instead, we qualitatively report the resulting filtering distributions in Figure 2. There is visually hardly any difference between the particle filtering estimate and our method.

One caveat to this positive result is that the bi-modal variational Wasserstein filter tends to collapse when the two modes are too close to each other. It is still unclear whether this is a feature of the general method or the ODE solver, and further investigations are warranted.

IV. CONCLUSION

In this article, we have presented a novel approach for approximate Gaussian filtering, which avoids the reliance on




Fig. 2. Performing filtering on a multi-modal dynamical system. We compare the filtering result of a particle filter (top) with that of a variational Wasserstein filter equipped with a mixture-of-Gaussians posterior representation (bottom). Both filters capture the bi-modal posterior distribution (red and blue) inferred from the observations (green) induced by the true states (black).

enabling assumptions [6], which are usually not amenable to multiplicative noise and not extendable to multi-modal distributions. Several questions remain open:

a) Assumptions: We have assumed, for simplicity, that the transition dynamics are governed by affine Gaussian densities. Overcoming this limitation is an interesting research question. In fact, in [13], the coupled ODEs (8) were originally introduced to propagate the Gaussian through a nonlinear stochastic differential equation.

b) Numerics: The numerical scheme chosen here is, primarily for ease of exposition, different from that of [8], which uses an iterative method called JKO [11] after [18], instead of integrating the differential equation (8) directly. It is unclear which approach is the best fit for filtering applications.

c) Mixture weights: The weights of the mixture in the ODE (20) are not allowed to vary. This modeling restriction needs to be relaxed. However, that may lead to identifiability issues (for example, a single Gaussian can be represented with a mixture of two Gaussians with different weights in infinite ways). Furthermore, introducing time-varying weights should be done carefully.

V. INDIVIDUAL CONTRIBUTIONS

The original idea and redaction of the article are due to Adrien Corenflos. Both authors contributed to the design of the methodology. The implementation and experiments are due to Hany Abulsamad.

REFERENCES

[1] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*. Courier Corporation, 2007.

[2] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Journal of Basic Engineering*, 1960.

[3] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.

[4] N. Chopin and O. Papaspiliopoulos, *An Introduction to Sequential Monte Carlo*. Springer, 2020.

[5] S. L. Heston, “A closed-form solution for options with stochastic volatility with applications to bond and currency options,” *The Review of Financial Studies*, 1993.

[6] A. F. García-Fernández, L. Svensson, M. R. Morelande, and S. Särkkä, “Posterior linearization filter: Principles and implementation using sigma points,” *IEEE Transactions on Signal Processing*, 2015.

[7] F. Tronarp, Á. F. García-Fernández, and S. Särkkä, “Iterative filtering and smoothing in nonlinear and non-Gaussian systems using conditional moments,” *IEEE Signal Processing Letters*, 2018.

[8] M. Lambert, S. Chewi, F. Bach, S. Bonnabel, and P. Rigollet, “Variational inference via Wasserstein gradient flows,” *arXiv preprint arXiv:2205.15902*, 2022.

[9] J. M. Joyce, “Kullback–Leibler divergence,” in *International Encyclopedia of Statistical Science*. Springer, 2011.

[10] K. P. Murphy, *Probabilistic Machine Learning: An Introduction*. MIT Press, 2022.

[11] F. Santambrogio, “Euclidean, metric, and Wasserstein gradient flows: An overview,” *Bulletin of Mathematical Sciences*, 2017.

[12] C. Villani, *Optimal Transport: Old and New*. Springer Berlin Heidelberg, 2009.

[13] S. Särkkä, “On unscented Kalman filtering for state estimation of continuous-time nonlinear systems,” *IEEE Transactions on Automatic Control*, 2007.

[14] G. Pagès, *Numerical Probability*. Springer, 2018.

[15] B. Christianson, “Reverse accumulation and attractive fixed points,” *Optimization Methods & Software*, 1994.

[16] N. J. Gordon, D. J. Salmond, and A. F. Smith, “Novel approach to nonlinear/non-Gaussian Bayesian state estimation,” in *IEEE Proceedings of Radar and Signal Processing*, 1993.

[17] S. Malik and M. K. Pitt, “Particle filters for continuous likelihood evaluation and maximisation,” *Journal of Econometrics*, 2011.

[18] R. Jordan, D. Kinderlehrer, and F. Otto, “The variational formulation of the Fokker–Planck equation,” *SIAM Journal on Mathematical Analysis*, 1998.