

Model error estimation using the expectation maximization algorithm and a particle flow filter

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Lucini, M. M. ORCID: <https://orcid.org/0000-0001-6615-7711>,
Van Leeuwen, P. J. and Pulido, M. (2021) Model error
estimation using the expectation maximization algorithm and a
particle flow filter. SIAM/ASA Journal on Uncertainty
Quantification, 9 (2). pp. 681-707. ISSN 2166-2525 doi:
<https://doi.org/10.1137/19M1297300> Available at
<https://centaur.reading.ac.uk/99536/>

It is advisable to refer to the publisher's version if you intend to cite from the
work. See [Guidance on citing](#).

Published version at: <http://dx.doi.org/10.1137/19M1297300>

To link to this article DOI: <http://dx.doi.org/10.1137/19M1297300>

Publisher: Society for Industrial and Applied Mathematics

All outputs in CentAUR are protected by Intellectual Property Rights law,
including copyright law. Copyright and IPR is retained by the creators or other
copyright holders. Terms and conditions for use of this material are defined in
the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Model Error Estimation Using the Expectation Maximization Algorithm and a Particle Flow Filter*

María Magdalena Lucini[†], Peter Jan van Leeuwen[‡], and Manuel Pulido[†]

Abstract. Model error covariances play a central role in the performance of data assimilation methods applied to nonlinear state-space models. However, these covariances are largely unknown in most of the applications. A misspecification of the model error covariance has a strong impact on the computation of the posterior probability density function, leading to unreliable estimations and even to a total failure of the assimilation procedure. In this work, we propose the combination of the expectation maximization (EM) algorithm with an efficient particle filter to estimate the model error covariance using a batch of observations. Based on the EM algorithm principles, the proposed method encompasses two stages: the expectation stage, in which a particle filter is used with the present updated value of the model error covariance as given to find the probability density function that maximizes the likelihood, followed by a maximization stage, in which the expectation under the probability density function found in the expectation step is maximized as a function of the elements of the model error covariance. This novel algorithm here presented combines the EM algorithm with a fixed point algorithm and does not require a particle smoother to approximate the posterior densities. We demonstrate that the new method accurately and efficiently solves the linear model problem. Furthermore, for the chaotic nonlinear Lorenz-96 model the method is stable even for observation error covariance 10 times larger than the estimated model error covariance matrix and also is successful in moderately large dimensional situations where the dimension of the estimated matrix is 40×40 .

Key words. particle filters, state-space models, model error covariance, EM algorithm

AMS subject classifications. 62M05, 62M20, 60G25, 93E10, 93E11

DOI. 10.1137/19M1297300

1. Introduction. Several research areas in which the object of research is a complex system, such as the atmosphere, the ocean, and biological systems, require estimating the state of the system through partial observational information which is distributed in time. Surrogate models of the system, which represent approximately the time evolution of the variables, are included as a source of information in the state estimation. This has a twofold aim; it regularizes the state estimation at a time and propagates information of the system and its uncertainty to the subsequent times. This estimation problem was called data assimilation in geophysical sciences [8, 18], and this terminology has subsequently been popularized in other

*Received by the editors November 4, 2019; accepted for publication (in revised form) February 16, 2021; published electronically May 20, 2021.

<https://doi.org/10.1137/19M1297300>

Funding: This work was funded by European Research Council (ERC) CUNDA project 694509 under the European Union Horizon 2020 research and innovation programme.

[†]UNNE-CONICET, Corrientes 3400, Argentina, and Department of Meteorology, University of Reading, Reading, Berkshire, RG6 6AH UK (lucini@exa.unne.edu.ar, pulido@unne.edu.ar).

[‡]Department of Atmospheric Science, Colorado State University, Fort Collins, CO 80523 USA, and Department of Meteorology, University of Reading, Reading, Berkshire, RG6 6AH UK (p.j.vanleeuwen@reading.ac.uk).

areas. The so-called state-space model is composed by two stochastic equations; one represents the evolution of the hidden state variables and is henceforth referred to as the dynamical model, and the other one maps the state of the system into the observation space through an observational model.

The increase of observational data availability, particularly indirect observations, and the increased complexity in surrogate models introduce nonlinear dependencies into the dynamical and observational models which in turn lead to non-Gaussian statistics. Accounting for these non-Gaussian statistics in high-dimensional state-space models is essential for the inference and represents one of the major challenges in the area [13, 35, 41]. Gaussianity-based techniques such as all the variants of the Kalman filter and optimization techniques based on maximum a posteriori estimation, also known as variational data assimilation, cannot deal with strongly non-Gaussian statistics. Monte Carlo techniques are one of the most promising methodologies that can fully consider the non-Gaussian uncertainty in the sequential inference [12, 15, 27]. These techniques aim to represent the state distribution by a set of realizations (referred to as particles) of this distribution. Particles are evolved between estimation times by the use of the dynamical model. However, one point of contention in these techniques is the particle degeneracy along the time sequence. After a few cycles, most of the particles finish with negligible weights, i.e., their likelihood to the observation is low or null, and only one particle remains with weight 1. This limitation is particularly important in high-dimensional systems. One solution is resampling; however, this brings another limitation since only a few particles with large weight remain, and therefore the diversity is lost, particularly in experiments which can only afford a small number of particles.

Recently a new framework for particle filters (PFs), called particle flow filters, has been proposed [4, 9]. Within this framework, particles are moved from the initial proposed distribution to the posterior distribution using flows that are consistent with Bayes's rule. Particle flows potentially avoid the need for resampling since the particles are moved to regions of higher likelihood of the state space. The flow is not unique and requires the solution of an ordinary differential equation under a given regularization. Alternatively, if the Gaussian assumption is taken, the expression for the flow can be derived analytically, and the flow is unique and known as Gaussian particle flow. In a recent work, Pulido and van Leeuwen [25] showed that under the assumption that the flow lies in a reproducing kernel Hilbert space, it is uniquely determined through a Monte Carlo integration, an interacting particle system. This approach combines Monte Carlo sampling with variational inference. It is shown that this sequential Monte Carlo filter does not require resampling even for long time sequences and can work in relatively high-dimensional systems.

A crucial assumption in PFs is that the model error uncertainty is known. The model error uncertainty is included in the evolution of the particles and is essential to account for the prediction uncertainty. This is particularly the case in particle flow filters, in which the sequential prior density is assumed to be known. In practice, model error uncertainty is not known, and it is highly dependent on the surrogate model we are using. The structure of the model error uncertainty, in terms, for example, of correlation between variables, is expected to be different if different surrogate models are used. Therefore, there is a current need for the development of model error estimation techniques that may be applied for sequential Monte Carlo filters.

Likelihood-based methods for model error covariance estimation may be broadly classified into two groups: maximum likelihood estimation and Bayesian inference. Maximum likelihood estimation methods have been applied assuming that covariance parameters are deterministic [5, 17, 35, 36]. In this case, non-a priori information on the parameters is required. When the gradient of the likelihood can be obtained analytically, Newton–Raphson optimization methods may be applied to maximize the likelihood function. When this is not feasible, one of the most popular methods for maximum likelihood estimation is expectation maximization (EM) [10]. One of the reasons for its widespread use is that it is readily applicable. In particular, contrary to Newton–Raphson optimization methods, EM does not depend on tunable parameters. The second group is a Bayesian approach in which the model error covariance parameters are interpreted as stochastic and the prior distribution of the parameters needs to be given. In this case, some hypotheses on the correlations between state-space variables and parameters are required [28]. Some authors [22, 33] assume that the uncertainty in the parameters does not affect the state density, while the authors of [29] use the marginalization of the hidden state for the parameter estimation in a hierarchical Bayesian framework.

In this article, we propose a new method based on a maximum likelihood approach to estimate the model error covariance matrix in state-space models. The article is organized as follows. In section 2 we formulate the problem, giving a brief overview of PFs and some details of a recently proposed particle flow filter to be used in the experiments before describing the proposed method to estimate the model error covariance matrix in subsection 2.3. Further details about its derivation are given in Appendix A and Appendix B. The proposed method is tested on a simple autoregressive linear model and on the Lorenz-96 model with 8 and 40 dimensions. The experiments’ design, results, and comparison with existing methods are shown in section 3, while in section 4 we present some conclusions.

2. Methodology.

2.1. Problem formulation. Consider a state-space model consisting of a nonlinear system and a nonlinear observation model, described by

$$(2.1) \quad \mathbf{x}_k = \mathcal{M}(\mathbf{x}_{k-1}) + \boldsymbol{\beta}_k,$$

$$(2.2) \quad \mathbf{y}_k = \mathcal{H}(\mathbf{x}_k) + \boldsymbol{\epsilon}_k,$$

where $\mathbf{x}_k \in \mathbb{R}^{N_x}$, called the state vector, is a hidden or latent process, and $\{\mathbf{y}_k\}_{k=1}^K$ is a time series of observations measured at times $k = 1, \dots, K$ with $\mathbf{y}_k \in \mathbb{R}^M$. The maps $\mathcal{M}(\cdot)$ and $\mathcal{H}(\cdot)$ denote the nonlinear deterministic dynamical model and the (possibly nonlinear) observation operator, respectively. The state process is assumed to be a Markov process, whereas the conditional density of the observations \mathbf{y}_k depends only on the current state \mathbf{x}_k for $k = 1, \dots, K$. The stochastic term $\boldsymbol{\beta}_k$ accounts for the missing physics in the model and its numerical approximations, whereas $\boldsymbol{\epsilon}_k$ is the observation noise. In this work, we assume that observational and model uncertainty are additive and $\boldsymbol{\beta}_k \sim \mathcal{N}(0, \mathbf{Q}_k)$ and $\boldsymbol{\epsilon}_k \sim \mathcal{N}(0, \mathbf{R}_k)$, where \mathbf{Q}_k and \mathbf{R}_k belong to the subspace of positive definite matrices in $\mathbb{R}^{N_x \times N_x}$ and $\mathbb{R}^{M \times M}$, respectively, representing the model error and observation error covariance matrices at time k . We denote by $\boldsymbol{\theta}_k \in \boldsymbol{\Theta}$ the vector of parameters $(\mathbf{Q}_k, \mathbf{R}_k)$.

Since \mathcal{M} and \mathcal{H} are surrogate models which approximate the system evolution and the processes that relate the observations with the hidden state, the model error and observa-

tional error covariances are expected to be largely unconstrained physically. Observational errors may be partially constrained from the knowledge of measurement errors, whereas representation errors, arising from the fact that models and observations often represent reality differently, are hard to determine in practice. Therefore, the parameters θ_k of the state-space model are unknown and need to be inferred from the data as well as the hidden state \mathbf{x}_k . In principle, unknown parameters from the dynamical and observational models may also be included in θ_k [24]; in this work we consider these parameters to be provided. In this parametric inference problem we need to assume an a priori underlying error distribution; the assumption $\beta_k \sim N(0, Q_k)$ was then taken. The technique here proposed could be extended to β_k belonging to an exponential family, as done in [7].

We assume that the model error covariance \mathbf{Q}_k and the observation error covariance \mathbf{R}_k vary slowly within K -cycles (the temporal scale of covariance variations is longer than K -cycles), that is, $\mathbf{Q}_k = \mathbf{Q}$, $\mathbf{R}_k = \mathbf{R} \forall k = 1, \dots, K$, and propose a method based on a time-batch of observations to estimate the model error covariance \mathbf{Q} for particle filters (PFs). The information provided by the observations along the K times, $\mathbf{y}_{1:K}$, where the subindices $1 : K$ denote the set $\{\mathbf{y}_1, \dots, \mathbf{y}_K\}$, is considered essential to regularize/constrain the $N_x \times N_x$ unknowns from \mathbf{Q} . The coupling between observations at different times is produced through the dynamical model \mathcal{M} .

The method presented here is based on maximum likelihood estimation: given a set of independent observations $\{\mathbf{y}_k, k = 1, \dots, K\}$ from a probability density function (pdf) represented by $p(\mathbf{y}_{1:K}; \theta)$, a nonlinear dynamical model \mathcal{M} , and an observation operator \mathcal{H} , we seek to maximize the likelihood of the observations as a function of the statistical parameters \mathbf{Q} given an observation error covariance \mathbf{R} in the presence of a hidden state $\mathbf{x}_{0:K}$. The statistical parameters to be estimated are assumed to vary slowly. This is an essential and unavoidable assumption in the state-space framework because this framework usually considers a single realization of the observation at each time. Hence, a time range with a long set of observations, say $K \geq 100$, is required to constrain the statistical parameters in which the model error covariance is assumed to be slowly varying. Online maximum likelihood techniques (e.g., [1, 7]) and hierarchical Bayesian inference [29] also need to assume slow variations in the statistical parameters.

2.2. Particle filters. State-space models are generally used in sequential data assimilation to estimate or reconstruct the hidden state \mathbf{x}_k given the observations $\mathbf{y}_{1:K}$. This can be done by computing the filter densities $\{p(\mathbf{x}_k | \mathbf{y}_{1:k})\}_{k=1:K}$ or smoother densities $\{p(\mathbf{x}_k | \mathbf{y}_{1:s})\}_{k=1:K}$ with $s \geq k$. Having prior knowledge of the initial state \mathbf{x}_0 , that is, given a prior pdf $p(\mathbf{x}_0)$, the posterior pdf of a filter is the probability of the model state at time k , given all the available information up to time k . In a Markovian system with observations \mathbf{y}_k that are conditionally independent given the state, the filter densities can be computed recursively using Bayes's rule to obtain the posterior pdf

$$(2.3) \quad p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta) = \frac{p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)}{p(\mathbf{y}_k | \mathbf{y}_{1:k-1}; \theta)},$$

where the prior pdf $p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)$ is the forecast or prediction pdf, $p(\mathbf{y}_k | \mathbf{x}_k; \theta)$ is the observation likelihood defined by the observation model \mathcal{H} and the distribution of the observation error

ϵ_k , and $p(\mathbf{y}_k|\mathbf{y}_{1:k-1};\boldsymbol{\theta})$ is a normalizing factor. Note that we consider here the marginalized posterior pdf, in which the only state variable that is estimated is the current one, considering all the past and the current observations.

When the dynamical and observational models are linear and their errors Gaussian, the filter densities are Gaussian and the state can be computed using Kalman recursive algorithms. However, for nonlinear models and/or nonlinear observational functions it is not possible to get a known distribution or a closed form for these filter pdfs, and they should be somehow approximated. Classical PFs [2, 12, 38] are based on sequential importance sampling and resampling algorithms and provide different methods to approximate these pdfs.

The basic idea behind a PF is to represent the posterior pdf $p(\mathbf{x}_k|\mathbf{y}_{1:k})$ by a set of N_p particles $\{x_k^{(j)}\}_{j=1:N_p}$ with corresponding weights $\{w_k^{(j)}\}_{j=1:N_p}$ such that $\sum_{j=1}^{N_p} w_k^{(j)} = 1$. That is, at time k , $p(\mathbf{x}_k|\mathbf{y}_{1:k};\boldsymbol{\theta})$ is approximated by

$$(2.4) \quad p(\mathbf{x}_k|\mathbf{y}_{1:k};\boldsymbol{\theta}) \doteq \sum_{j=1}^{N_p} w_k^{(j)} \delta(\mathbf{x}_k - \mathbf{x}_k^{(j)}),$$

where \doteq denotes approximation by an ensemble of particles and $\delta(\cdot)$ is the Dirac δ function.

Initially, a set of N_p particles $\{x_0^{(j)}\}_{j=1:N_p}$ with corresponding weights $\{w_0^{(j)} = \frac{1}{N_p}\}_{j=1:N_p}$ is drawn from the prior pdf $p(\mathbf{x}_0)$. These particles are sequentially evolved in time using a forecasting, weighting, and resampling scheme to obtain $\{x_k^{(j)}\}_{j=1:N_p}$ and $\{w_k^{(j)}\}_{j=1:N_p}$ at each time step k . Different PFs were proposed depending on the resampling, forecasting, or weighting approaches taken [2, 25, 43].

2.2.1. Variational mapping particle filter. The variational mapping PF (VMPF) is a particle flow filter which is based on optimization and Monte Carlo sampling. At each cycle, a particle flow filter [4, 9] drives the particle samples from the prior density towards a desired target density. In the VMPF, the particles are moved deterministically via a sequence of maps, based on the optimal transport principle. The maps seek to minimize the Kullback–Leibler divergence (KLD) between the target density, i.e., the posterior density $p(\mathbf{x}_k|\mathbf{y}_{1:k})$, and an intermediate density $q(\mathbf{x}_k)$, that is, the density represented through the sample, the set of particles, at a given cycle k . At the i th optimization iteration, the KLD is given by

$$(2.5) \quad \mathcal{D}_{KL}(q_i(\mathbf{x}_k)||p(\mathbf{x}_k|\mathbf{y}_{1:k})) = \int q_i(\mathbf{x}_k) \log \frac{q_i(\mathbf{x}_k)}{p(\mathbf{x}_k|\mathbf{y}_{1:k})} d\mathbf{x}_k,$$

where $q_i(\mathbf{x}_k)$ is represented via N_p sample points, i.e., particles, $\mathbf{x}_{k,i}^{(1:N_p)} \sim q_i(\mathbf{x}_k)$.

The density q_i is the result of a map, $T_i(\mathbf{x}_{k,i-1}) = \mathbf{x}_{k,i-1} + \epsilon \mathbf{v}_i(\mathbf{x}_{k,i-1})$, which is a small perturbation to the identity map (ϵ is assumed to be small). This means $q_i = T_i^\# q_{i-1} = q_{i-1}(T_i^{-1})|\det J(T_i^{-1})|$. The maps are assumed to be in a reproducing kernel Hilbert space (RKHS).

The optimal map that gives the steepest descent direction in the RKHS is shown to be given by

$$(2.6) \quad \mathbf{v}_i(\mathbf{x}_k) = -\nabla \mathcal{D}_{KL}(\mathbf{x}_k) = \mathcal{E}_{\mathbf{x}'_k \sim q_i} \left[\mathbf{K}(\mathbf{x}'_k, \mathbf{x}_k) \nabla_{\mathbf{x}'_k} \log p(\mathbf{x}'_k|\mathbf{y}_{1:k}) + \nabla_{\mathbf{x}'_k} \mathbf{K}(\mathbf{x}'_k, \mathbf{x}_k) \right],$$

where \mathbf{K} is a kernel (assumed here to be Gaussian) [25].

Then the optimization is a sequence of (sufficiently smooth) mappings in which each particle (j) is moved along the steepest descent direction

$$(2.7) \quad \mathbf{x}_{k,i}^{(j)} = \mathbf{x}_{k,i-1}^{(j)} - \epsilon \nabla \mathcal{D}_{KL}(\mathbf{x}_{k,i-1}^{(j)}).$$

If observational and model errors are assumed to be Gaussian, the gradient of the log-posterior density is

$$(2.8) \quad \nabla_{\mathbf{x}} \log p(\mathbf{x}_k) = \mathbf{H}^\top \mathbf{R}_k^{-1} (\mathbf{y}_k - \mathcal{H}(\mathbf{x}_k)) - \mathbf{Q}_k^{-1} \left[\mathbf{x}_k - \sum_{j=1}^{N_p} \beta_{k-1}^{(j)} \mathcal{M}(\mathbf{x}_{k-1}^{(j)}) \right],$$

where $\beta_{k-1}^{(j)} \triangleq \frac{w_{k-1}^{(j)} \psi_{\mathbf{Q}_k}^{(j)}}{\sum_{m=1}^{N_p} w_{k-1}^{(m)} \psi_{\mathbf{Q}_k}^{(m)}}$, and $\psi_{\mathbf{Q}_k}^{(j)} = \exp(-\|\mathbf{x}_k - \mathbf{x}_{k-1}^{(j)}\|_{\mathbf{Q}_k}^2)$. They could be interpreted as weights of the forecast states which consider the distance between the particles to the point under consideration. A more detailed description of the VMPF is found in [25]. Recently, a generalization to embed also the observational operator in the RKHS was proposed in [26].

One of the main advantages of the VMPF is that it not only efficiently samples high-dimensional state spaces with a limited number of particles but also does not suffer from sample impoverishment. Since (2.6) moves the particles toward regions of high observation likelihood, by construction the resulting sample has a large effective sample size close to the number of particles. Thus, this particle flow filter does not require resampling even for long recursions, as shown in [25].

2.3. Parameter estimation. Let us assume that $p(\cdot; \boldsymbol{\theta})$ is a parametric distribution, with $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ the parameter space. Given a set of observations $\mathbf{y}_{1:K} = \{\mathbf{y}_k, k = 1, \dots, K\}$ taken along a time interval of length K , a maximum likelihood estimation method aims at finding the value of $\boldsymbol{\theta}$ that maximizes the (incomplete) likelihood of the observations,

$$(2.9) \quad L(\boldsymbol{\theta}) = p(\mathbf{y}_{1:K}; \boldsymbol{\theta}) = \int p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \boldsymbol{\theta}) d\mathbf{x}_{0:K},$$

or equivalently, the log-likelihood function

$$(2.10) \quad l(\boldsymbol{\theta}) = \ln p(\mathbf{y}_{1:K}; \boldsymbol{\theta}) = \ln \left(\int p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \boldsymbol{\theta}) d\mathbf{x}_{0:K} \right).$$

The likelihood function (2.9) can be interpreted as how probable the set of observations $\mathbf{y}_{1:K}$ would be for different choices of $\boldsymbol{\theta}$. An analytic form for the log-likelihood function is not achievable in practice, and the numerical evaluation of (2.10) may involve high-dimensional integrations, which is intractable. In some situations the optimization task can be accomplished by using numerical optimization routines like Newton–Raphson techniques to solve the nonlinear equations obtained by differentiating the log-likelihood function (2.10) [5, 16, 24]. However, even in these particular situations, other methods are preferable due to the difficulty of implementing optimization methods and tuning their parameters. Gradient optimization methods may not be stable numerically for certain sets of parameters.

The expectation maximization (EM) algorithm [10] is a widely used numerical method that aims at maximizing the log-likelihood of the observations as a function of the statistical parameters θ in the presence of a hidden state $\mathbf{x}_{0:K}$ in successive iterations without the need to evaluate the complete log-likelihood function.

It basically consists in maximizing iteratively an intermediate function defined as

$$(2.11) \quad \mathcal{G}(\theta', \theta) \triangleq \mathbb{E}_{\theta'} [\ln p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \theta)]$$

$$(2.12) \quad = \int \ln p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \theta) p(\mathbf{x}_{0:K} | \mathbf{y}_{1:K}; \theta') d\mathbf{x}_{0:K},$$

where $\theta', \theta \in \Theta$. This intermediate function \mathcal{G} is, generally, much simpler to maximize than the incomplete log-likelihood defined in (2.10).

Starting from an initial parameter θ_0 , the two steps of the EM algorithm at iteration s can be summarized as follows:

- **Expectation Step** (E-Step): Calculate the required densities to compute the intermediate function $\mathcal{G}(\theta_{s-1}, \theta)$ as in (2.12).
- **Maximization Step** (M-Step): find $\theta_s = \max_{\theta \in \Theta} \mathcal{G}(\theta_{s-1}, \theta)$.

The assumptions of a hidden Markov model and mutually independent observations in (2.1) and (2.2) allow us to express the joint pdf $p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \theta)$ as

$$(2.13) \quad p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \theta) = p(\mathbf{x}_0; \theta) \prod_{k=1}^K p(\mathbf{x}_k | \mathbf{x}_{k-1}; \theta) \prod_{k=1}^K p(\mathbf{y}_k | \mathbf{x}_k; \theta).$$

To estimate the parameters of the state-space model using the EM algorithm, the expectation of this last pdf under the conditional (smoother) pdf $p(\mathbf{x}_{0:K} | \mathbf{y}_{1:K}; \theta')$ must be computed in the E-Step of the algorithm. In the case of a linear Gaussian model this can be accomplished by means of a Kalman smoother [31]. This was further extended to the ensemble Kalman filter in [14, 24].

When the dynamical or observational model is nonlinear and therefore the joint density non-Gaussian, the expectation in (2.11) may be intractable, and a different approach must be taken. A generally used approach is to approximate the expression in (2.11) by generating samples of the smoother pdf $p(\mathbf{x}_{0:K} | \mathbf{y}_{1:K}; \theta')$ using a particle smoother [3, 17, 21, 23]. However, the use of particle smoothers in data assimilation represents a computational challenge, since they not only tend to degenerate rapidly but also have a poor performance in moderate to high-dimensional spaces, particularly if the time sequence is long (large K).

The requirement of a particle smoother in the E-Step of the EM algorithm is due to the fact that the likelihood of the observations $p(\mathbf{y}_{1:K}; \theta)$ is usually obtained by marginalizing the joint pdf $p(\mathbf{x}_{0:K}, \mathbf{y}_{1:K}; \theta)$ over the whole state $\mathbf{x}_{0:K}$ (cf. (2.9)).

Instead of using this last expression for the likelihood of the observations, and following the notation of [6], the likelihood of the observations (model evidence) can be decomposed as $p(\mathbf{y}_{1:K}; \theta) = \prod_{k=1}^K p(\mathbf{y}_k | \mathbf{y}_{1:k-1}; \theta)$, with the convention $\mathbf{y}_{1:0} = \{\emptyset\}$. Marginalizing this last

expression, we obtain

$$\begin{aligned}
 p(\mathbf{y}_{1:K}; \boldsymbol{\theta}) &= \prod_{k=1}^K p(\mathbf{y}_k | \mathbf{y}_{1:k-1}; \boldsymbol{\theta}) \\
 (2.14) \qquad &= \prod_{k=1}^K \int p(\mathbf{y}_k | \mathbf{x}_k; \boldsymbol{\theta}) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \boldsymbol{\theta}) d\mathbf{x}_k,
 \end{aligned}$$

and therefore we can rewrite the logarithm of the incomplete likelihood (2.9) as

$$(2.15) \qquad l(\boldsymbol{\theta}) = \log p(\mathbf{y}_{1:K}; \boldsymbol{\theta}) = \log \prod_{k=1}^K \int p(\mathbf{y}_k | \mathbf{x}_k; \boldsymbol{\theta}) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \boldsymbol{\theta}) d\mathbf{x}_k.$$

Using this last expression instead of the commonly used expression (2.10), the intermediate function $\mathcal{G}(\boldsymbol{\theta}', \boldsymbol{\theta})$ of the EM algorithm can be written as

$$(2.16) \qquad \mathcal{G}(\boldsymbol{\theta}', \boldsymbol{\theta}) = \sum_{k=1}^K \int p(\mathbf{x}_k | \mathbf{y}_{1:k}; \boldsymbol{\theta}') \log \left(\frac{p(\mathbf{y}_k | \mathbf{x}_k; \boldsymbol{\theta}) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \boldsymbol{\theta})}{p(\mathbf{x}_k | \mathbf{y}_{1:k}; \boldsymbol{\theta}')} \right) d\mathbf{x}_k,$$

as shown in Appendix A. Note that this last expression for $\mathcal{G}(\boldsymbol{\theta}', \boldsymbol{\theta})$ is written in terms of filter and forecast pdfs, while smoother pdfs are no longer required. As the denominator of (2.16) does not depend on $\boldsymbol{\theta}$, maximizing $\mathcal{G}(\boldsymbol{\theta}', \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ is equivalent to maximizing

$$(2.17) \qquad \mathcal{G}(\boldsymbol{\theta}', \boldsymbol{\theta}) = \sum_{k=1}^K \int p(\mathbf{x}_k | \mathbf{y}_{1:k}; \boldsymbol{\theta}') \log (p(\mathbf{y}_k | \mathbf{x}_k; \boldsymbol{\theta}) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \boldsymbol{\theta})) d\mathbf{x}_k$$

$$(2.18) \qquad = \sum_{k=1}^K \mathcal{E}_{\boldsymbol{\theta}'} [\log (p(\mathbf{y}_k | \mathbf{x}_k; \boldsymbol{\theta}) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \boldsymbol{\theta}))].$$

So far $\boldsymbol{\theta}$ denotes the set comprised by (\mathbf{Q}, \mathbf{R}) , the model error and observation error covariance matrices, respectively. As there is a certain degree of knowledge about the instrument's noise and how observations are measured or obtained, \mathbf{R} is usually determined empirically in practice by estimating these noises and the errors between the state and observation space. However, this is far from being an accurate representation of \mathbf{R} , and during the last few years great effort has been dedicated to the study and estimation of the observation error covariance matrix, and many works have been published on these topics (see [20, 32, 36, 37, 39] and references therein). Some of these works focus on studying the most plausible structure for \mathbf{R} given the nature of how observations were measured, taking into account correlations between observation errors, while some others just propose a fixed structure for \mathbf{R} and a methodology to estimate it. What these different approaches have in common is that they assume that the model error covariance matrix is already given, or known. The model error covariance matrix \mathbf{Q} is, perhaps, the most difficult one to estimate or determine, since it accounts for the model inaccuracies and deficiencies in representing the missing underlying physics, the errors in parameterizations, the unresolved and smaller scales, and the numerical schemes

used. Some works have been devoted to the joint estimation of (\mathbf{Q}, \mathbf{R}) , and an up to date and detailed review of these techniques is presented in [34].

The main purpose of this work is to provide a method to estimate \mathbf{Q} , the covariance matrix of the model error β . Assuming that \mathbf{R} is known, replacing $\theta = \mathbf{Q}$ and having in mind that by hypothesis the density $p(\mathbf{y}_k|\mathbf{x}_k; \mathbf{Q})$ is assumed to be independent of \mathbf{Q} , starting from an initial guess \mathbf{Q}_0 , the two steps of the EM algorithm at iteration s can be summarized as follows:

- **Expectation Step** (E-Step): Calculate the required densities to compute the intermediate function

$$(2.19) \quad \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q}) = \sum_{k=1}^K \mathcal{E}_{\mathbf{Q}_{s-1}} [\log(p(\mathbf{x}_k|\mathbf{y}_{1:k-1}; \mathbf{Q}))].$$

- **Maximization Step** (M-Step): find $\mathbf{Q}_s = \max_{\mathbf{Q} \in \Theta} \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q})$, where Θ is the space of positive definite matrices of order N_x .

Using a PF, the posterior pdf $p(\mathbf{x}_k|\mathbf{y}_{1:k}; \mathbf{Q}_{s-1})$ can be approximated by a set of particles and their corresponding weights as in (2.4) with $\theta = \mathbf{Q}_{s-1}$.

As we assume a Gaussian model noise $\beta \sim \mathcal{N}(0, \mathbf{Q})$ (see subsection 2.1), the transition density is given by $p(\mathbf{x}_k|\mathbf{x}_{k-1}; \mathbf{Q}) = \phi(x_k, \mathcal{M}(\mathbf{x}_{k-1}), \mathbf{Q})$, where $\phi(x_k, \mathcal{M}(\mathbf{x}_{k-1}), \mathbf{Q}) \triangleq \frac{1}{(2\pi)^{n/2}|\mathbf{Q}|^{1/2}} \exp\{-\frac{1}{2}(\mathbf{x}_k - \mathcal{M}(\mathbf{x}_{k-1}))^\top \mathbf{Q}^{-1}(\mathbf{x}_k - \mathcal{M}(\mathbf{x}_{k-1}))\}$.

Therefore, the prediction density $p(\mathbf{x}_k|\mathbf{y}_{1:k-1}; \mathbf{Q})$ can be approximated by

$$(2.20) \quad p(\mathbf{x}_k|\mathbf{y}_{1:k-1}; \mathbf{Q}) \doteq \sum_{i=1}^{N_p} w_{k-1}^{(i)} \phi(\mathbf{x}_k, \mathcal{M}(\mathbf{x}_{k-1}^{(i)}), \mathbf{Q}),$$

where $w_{k-1}^{(i)}, \mathbf{x}_{k-1}^{(i)}$ are the i th weight and particle, respectively, obtained by a PF at time step $k-1$.

Combining these approximated pdfs and computing the required expectations, the intermediate function $\mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q})$ given in (2.19) can now be written in terms of N_p particles as

$$(2.21) \quad \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q}) \doteq \sum_{k=1}^K \sum_{j=1}^{N_p} w_{k, \mathbf{Q}_{s-1}}^{(j)} \log \left(\sum_{i=1}^{N_p} w_{k-1, \mathbf{Q}}^{(i)} \phi(x_k^{(j)}, \mathcal{M}(\mathbf{x}_{k-1}^{(i)}), \mathbf{Q}) \right).$$

Differentiating (2.21) with respect to \mathbf{Q} , we can determine the root of $\frac{\partial \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q})}{\partial \mathbf{Q}} = 0$ to obtain the maximum of the intermediate function $\mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q})$ at iteration s . By doing this, we obtain (Appendix B)

$$(2.22) \quad \mathbf{Q} = f_{\mathbf{Q}_{s-1}}(\mathbf{Q}) = \frac{1}{K} \sum_{k=1}^K \left[\sum_{j=1}^{N_p} w_{k, \mathbf{Q}_{s-1}}^{(j)} \left(\frac{1}{S_{j,i}} \sum_{i=1}^{N_p} w_{k-1, \mathbf{Q}}^{(i)} \psi_k^{j,i} \mathbf{B}_k^{j,i} \right) \right],$$

where $\psi_k^{j,i} \triangleq \exp\{-\frac{1}{2}(\mathbf{x}_k^{(j)} - \mathcal{M}(\mathbf{x}_{k-1}^{(i)}))^\top \mathbf{Q}^{-1}(\mathbf{x}_k^{(j)} - \mathcal{M}(\mathbf{x}_{k-1}^{(i)}))\}$, $\mathbf{B}_k^{j,i} \triangleq (\mathbf{x}_k^{(j)} - \mathcal{M}(\mathbf{x}_{k-1}^{(i)}))(\mathbf{x}_k^{(j)} - \mathcal{M}(\mathbf{x}_{k-1}^{(i)}))^T$, and $S_{j,i} \triangleq \sum_{i=1}^{N_p} w_{k-1, \mathbf{Q}}^{(i)} \psi_k^{j,i}$.

We have to take into account that particles and weights indexed by j , that is, $\mathbf{x}_k^{(j)}$ and $w_{k, \mathbf{Q}_{s-1}}^j$ in (2.22), are computed in the expectation step using \mathbf{Q}_{s-1} as the model error covariance matrix. On the other hand, particles and weights indexed by i , that is, $\mathbf{x}_{k-1}^{(i)}$, $w_{k-1, \mathbf{Q}}^i$, should be computed with a filter that assumes \mathbf{Q} as the model error covariance matrix. However, this \mathbf{Q} is the matrix to be found in the maximization step.

Summarizing, the two steps of this EM algorithm using a PF at iteration s are as follows:

- **E-Step:** Given \mathbf{Q}_{s-1} , use a PF with a model error covariance \mathbf{Q}_{s-1} to calculate the weights and particles which are needed for the intermediate function $\mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q})$.
- **M-Step:** Solve (2.22).

The M-Step at iteration s requires solving an implicit equation for the covariance matrix \mathbf{Q} , where the inverse of \mathbf{Q} is involved. In this work we propose to solve this equation using a fixed point algorithm in \mathbf{Q} . This fixed point algorithm requires extra iterations at each iteration of the EM algorithm; however the conducted experiments showed that fewer than six iterations of the fixed point algorithm are enough to satisfy the required stopping criteria. The Banach fixed point theorem, or contractive mapping theorem, guarantees the existence (and uniqueness) of a fixed point of certain mappings (functions) defined on a complete metric space, as long as these mappings are contractive. It is not possible to show analytically that the function involved in our fixed point algorithm is contractive due to the nonlinearity in the dynamical model, and since we are not under the hypothesis of the Banach fixed point theorem we cannot guarantee that our algorithm converges to a fixed point. What we can guarantee, based on empirical evidence, is that the algorithm stops after a few iterations, satisfying a stopping criterion based on the Frobenius norm defined below.

The pseudocode of the proposed algorithm is presented in Algorithm 2.1. Within this pseudocode, PF (\mathbf{Q}) indicates that the particles and corresponding weights are obtained by using a PF with \mathbf{Q} as model error covariance matrix. The algorithm evaluates the fixed point function, (2.22), to obtain the new updated value of the parameters.

The stopping criterion for the fixed point algorithm (FP) is defined as either $\text{stop}_C(\mathbf{Q}_{\text{FP}}, \mathbf{Q}_{\text{FP}0}) = \frac{\|\mathbf{Q}_{\text{FP}} - \mathbf{Q}_{\text{FP}0}\|_F}{\|\mathbf{Q}_{\text{FP}}\|_F}$, where $\|\cdot\|_F$ is the Frobenius norm, is smaller than a previously set threshold, or when the maximum number of fixed point iterations is reached.

The algorithm in its current form does not assume any a priori knowledge of the structure of \mathbf{Q} , and therefore the full matrix is estimated. That is, at each iteration of the EM algorithm the full covariance matrix \mathbf{Q} is computed (M-Step) by means of a fixed point algorithm. The updated value of \mathbf{Q} , at the m th iteration of the fixed point algorithm, is given by $\mathbf{Q}_m = f_{\mathbf{Q}_{EM}}(\mathbf{Q}_{m-1})$, where \mathbf{Q}_{m-1} is the matrix computed in the previous iteration of the fixed point algorithm and $f_{\mathbf{Q}_{EM}}(\mathbf{Q}_{m-1})$ is calculated using the formula for $f_{\mathbf{Q}_{EM}}(\mathbf{Q}_{m-1})$ given in (2.22). Hence, the resulting \mathbf{Q}_m is a linear combination of the positive semidefinite matrices $\mathbf{B}_k^{j,i}$ multiplied by nonnegative real numbers ($0 \leq w_{k,*}^{(j)} < 1$, $0 < \psi_k^{j,i}$, and $S_{j,i} > 0$), and therefore the updated value \mathbf{Q}_m is also a positive semidefinite matrix.

In practice, these numerical calculations appear to be robust to roundoff errors so that it is not required to impose any further restriction to guarantee the positive semidefiniteness of \mathbf{Q} . For higher-dimensional problems, symmetry may be assumed a priori, and the upper triangular matrix of the model error covariance may be obtained via sparse matrix multiplication algorithms.

Algorithm 2.1. EM algorithm to estimate \mathbf{Q} using a PF and a fixed point algorithm.

Given $\mathbf{Q}_{EM}, \mathbf{x}_0^{1:N_p}, w_0^{1:N_p}, \epsilon_{EM}, \epsilon_{FP}$
 $\text{iter}_{EM} = 1$
while ($\text{iter}_{EM} \leq \text{Maxiter}_{EM}$) **and** ($\text{stop}_{EM} > \epsilon_{EM}$) **do**
 E-Step:
 Compute $\{\mathbf{x}_k^{(j)}, w_{k, \mathbf{Q}_{EM}}^{(j)}\}$ using $\text{PF}(\mathbf{Q}_{EM})$
 M-Step:
 $\text{iter}_{FP} = 1$; $\text{stop}_{FP} = 2\epsilon_{FP}$
 $\mathbf{Q}_{FP0} = \mathbf{Q}_{EM}$
 while ($\text{iter}_{FP} \leq \text{Maxiter}_{FP}$) **and** ($\text{stop}_{FP} > \epsilon_{FP}$) **do**
 $\mathbf{Q}_{FP} = f_{\mathbf{Q}_{EM}}(\mathbf{Q}_{FP0})$ from (2.22)
 Compute $\text{stop}_{FP} = \text{stop}_C(\mathbf{Q}_{FP}, \mathbf{Q}_{FP0})$
 Compute $\{\mathbf{x}_{k-1}^{(i)}, w_{k-1, \mathbf{Q}_{FP}}^{(i)}\}$ using $\text{PF}(\mathbf{Q}_{FP})$
 $\text{iter}_{FP} = \text{iter}_{FP} + 1$; $\mathbf{Q}_{FP0} = \mathbf{Q}_{FP}$
 end while
 Compute $\text{stop}_{EM} = \text{stop}_C(\mathbf{Q}_{EM}, \mathbf{Q}_{FP})$
 $\mathbf{Q}_{EM} = \mathbf{Q}_{FP}$
 $\text{iter}_{EM} = \text{iter}_{EM} + 1$
end while

3. Numerical experiments' design. In order to evaluate the capabilities and performance of the proposed methodology, numerical experiments were designed using two different dynamical models \mathcal{M} , a univariate linear Gaussian model and the Lorenz-96 model [19]. For each of these models we conducted twin experiments with different settings. We first generated a set of noisy observations using the dynamical model with known parameters. Then, using these synthetic observations and the same stochastic dynamical model, we estimated the model error covariance \mathbf{Q} with the proposed algorithm and compared the results with those obtained with some classical methodologies. These experiments are useful to assess the convergence and performance of the proposed methodology.

3.1. Linear model. A one-dimensional linear Gaussian state-space model is defined as

$$(3.1) \quad \begin{aligned} x_k &= \nu x_{k-1} + \beta_k, \\ y_k &= x_k + \epsilon_k, \end{aligned}$$

where $\{x_k\}_{k=0:K}, \{y_k\}_{k=1:K} \in \mathbb{R}$, $\beta_k \sim \mathcal{N}(0, \mathbf{Q})$, $\epsilon_k \sim \mathcal{N}(0, \mathbf{R})$, ν is the autoregressive coefficient, and \mathbf{Q}, \mathbf{R} are the error variances. The implementation of the EM algorithm to estimate the parameters of this linear Gaussian model using the Kalman filter and smoother was first discussed by Shumway and Stoffer in [31], whereas in [30] the same authors provide a more detailed derivation of this implementation. Discussions of the convergence of the EM algorithm for this model can be found in [11, 40, 42]. A set of one-dimensional experiments was conducted by generating $K = 100$ noisy observations $\mathbf{y}_{1:K}$ using the linear model (3.1) with known parameter values $\nu = 0.8$, $\mathbf{Q}_{true} = \mathbf{I}_1$, $\mathbf{R} = \mathbf{I}_1$. Then, using the same model, these synthetic observations, and an initial guess \mathbf{Q}_0 sampled from a uniform distribution $\mathcal{U}[0.5, 1.5]$,

the model error variance \mathbf{Q} was estimated by four different algorithms:

- EM+VMPF: The algorithm here proposed, which is a version of the EM algorithm for a PF without the need of a particle smoother. The PF used is the variational mapping particle filter (VMPF) described in [subsection 2.2.1](#).
- EM+SIR: The algorithm here proposed, coupled with a classical sampling importance resampling (SIR) filter [\[2\]](#).
- EM+KF+KS: The classical EM algorithm based on the Gaussian assumption that requires a Kalman filter and smoother as presented in [\[30\]](#).
- EM+EnKF+EnKS: A version of the EM algorithm in conjunction with the ensemble Kalman filter and an ensemble Kalman smoother as presented in [\[24\]](#) with $N_p = 50$ ensemble members.

The procedure was repeated independently 50 times for each algorithm in order to have an empirical distribution of the estimators. That means the “true” state and a set of observations $\{\mathbf{y}_{1:K}\}$ were generated independently for each experiment realization, and an independent initial guess \mathbf{Q}_0 was sampled from a uniform distribution $\mathcal{U}[0.5, 1.5]$. Then the model error covariance \mathbf{Q} was estimated. This procedure was repeated 50 times for each algorithm. With these results we can have an approximation of the empirical distribution of the estimators.

As the algorithm here proposed is suitable to be used with any PF, we tested the EM+PF algorithm with two different PFs, namely the classical SIR filter with $N_p = 1000$ particles (for a detailed explanation see [\[2\]](#)) and the recently proposed PF based on optimal transportation and referred to as the variational mapping particle filter (VMPF) ([subsection 2.2.1](#) and [\[25\]](#)) with $N_p = 20$ particles.

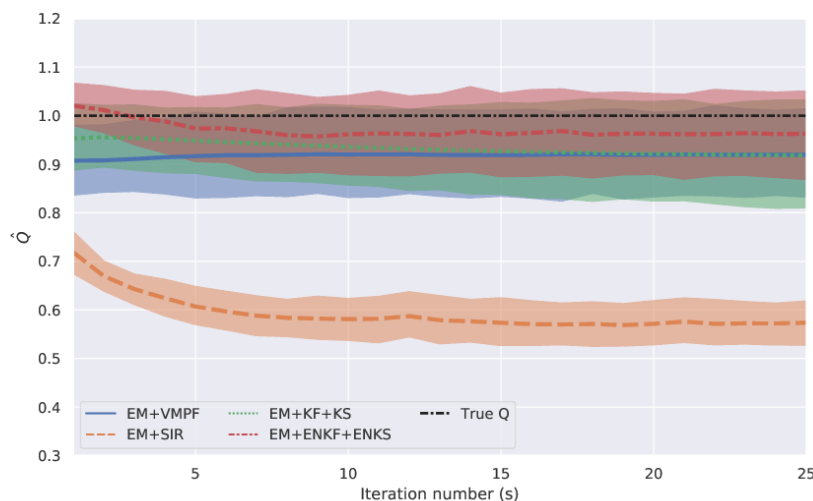


Figure 1. $\hat{\mathbf{Q}}$ as function of the iteration number using the EM+VMPF (blue), EM+SIR (orange), EM+KF+KS (green), and EM+EnKF+EnKS (red) algorithms, for the linear model ([subsection 3.1](#)) with true parameter value $\mathbf{Q}_{true} = \mathbf{I}_1$, $\mathbf{R} = \mathbf{I}_1$, $\nu = 0.8$, and $K = 100$. The 95% confidence intervals were generated by running 50 independent repetitions of these estimation experiments and 25 EM iterations. For the EM+VMPF only $N_p = 20$ particles were used, whereas for the EM+SIR and the EM+EnKF+EnKS $N_p = 1000$ particles and $N_p = 50$ ensemble members, respectively, were used.

Figure 1 shows the mean and 95% confidence intervals obtained by the 50 independent repetitions of the experiments for each algorithm. The EM+VMPF, EM+KF+KS, and EM+EnKF+EnKS algorithms show a similar behavior, stabilizing after 5 to 6 iterations of their respective EM versions; the EM+EnKF+EnKS one is the less biased (in terms of the mean value of the confidence interval). The 95% confidence intervals for these three methods include the true value $\mathbf{Q}_{true} = \mathbf{I}_1$. With only $N_p = 20$ particles the EM+VMPF algorithm results are very similar to the EM+KF+KS ones. On the other hand, for this simple linear model and using $N_p = 1000$ particles, the results obtained by using the SIR filter are biased, and the 95% confidence interval never reaches the true value of $\mathbf{Q}_{true} = \mathbf{I}_1$. In the experiments that follow, we only use EM+VMPF for the PF implementation of the proposed algorithm. Experiments coupling the EM algorithm with the SIR filter for high-dimensional state and observational spaces, including the Lorenz-96 system, are not feasible due to the large number of particles required [25].

3.2. Lorenz-96 model. In this section we show results of twin experiments using the chaotic and nonlinear Lorenz-96 system as the dynamical model \mathcal{M} in (2.1). The Lorenz-96 is one of the most used toy models in data assimilation within the geoscience community due to its ability to mimic certain properties of the atmospheric predictability at a low computational cost.

It is defined by the ordinary differential equations

$$(3.2) \quad \frac{dX_n}{dt} = -X_{n-2}X_{n-1} + X_{n-1}X_{n+1} - X_n + F,$$

where X_n is the state variable of the model at variable (position) n , $n = 1, \dots, N_x$, and the domain is assumed to be periodic, that is, $X_{-1} \equiv X_{N-1}$, $X_0 \equiv X_N$, $X_{N+1} \equiv X_1$. F is the forcing constant, and, as usual, here it is set as $F = 8$ to represent chaotic dynamics.

We used a fourth-order Runge–Kutta scheme with a model time step of $\delta_t = 0.005$ to integrate the Lorenz-96 equations (3.2). In the first set of experiments with the Lorenz-96 system, the number of variables is set to $N_x = 8$, meaning that an 8×8 model error covariance \mathbf{Q} has to be estimated. The observation error covariance \mathbf{R} was chosen as a diagonal matrix $\mathbf{R} = \sigma_R^2 \mathbf{I}$. We observed every grid point; in other words, the observation model \mathcal{H} is assumed to be the identity transformation. The observations are taken every $\Delta_t = 0.05$, which represents 10 model time steps in all the experiments performed.

For these twin experiments, we simulated $K = 500$ noisy observations $\mathbf{y}_{1:K}$ using the Lorenz-96 model with known model error and observation error covariances \mathbf{Q}_{true} , \mathbf{R} . Using the same model and these synthetic observations, the full model error covariance matrix \mathbf{Q} was estimated by three different algorithms:

- EM+VMPF: The algorithm here proposed, which is a version of the EM algorithm for a PF without the need of a particle smoother. The PF used is the VMPF, as in the previous section.
- EM+VMPF+EnKS: A “hybrid” method, consisting of the “classical” batch EM, where the VMPF is used in the forward pass of the E-Step and a Gaussian smoother is used in the backwards pass.
- EM+EnKF+EnKS, as presented in [24].

In all these algorithms the number of particles/ensemble members was $N_p = 20$, except when otherwise stated. The method here proposed is designed to be used with a PF, avoiding the need of a particle smoother due to the computational disadvantages they present. A fair competitor to analyze the performance of our methodology would be the classical EM using a PF and a particle smoother. As already mentioned, particle smoothers tend to degenerate rapidly and have a poor performance in moderate-dimensional spaces, and so we propose to use the EnKS smoother combined with the VMPF instead for comparison purposes.

The performance of the VMPF in the Lorenz-96 dynamical system has been thoroughly evaluated in [25]. In particular, that work showed that the effective sample size was close to the number of particles, and no resampling was required for the 40-variable Lorenz-96 system. The EM+EnKF+EnKS uses the perturbed observations EnKF, while the EnKS is based on the linear RTS smoother. Furthermore, since this work attempts to estimate the full model error covariance matrix, we did not use either localization or inflation in the experiments. However, we note that model error may be interpreted as additive inflation so that in principle model error estimation should account for underdispersion of the ensemble.

To assess the proposed methodology to estimate the model error covariance matrix, experiments with two different structures, usually assumed in practice, were proposed for the true model error covariance matrix \mathbf{Q}_{true} : (a) an isotropic noncorrelated covariance matrix, where $\mathbf{Q}_{true} = \sigma^2 \mathbf{I}$, with \mathbf{I} the identity matrix of order N_x , and (b) an isotropic tridiagonal covariance matrix with diagonal values σ_d^2 and both sub-/superdiagonal values σ_{sd}^2 . In the first case we assume that model errors for different model variables are uncorrelated and have the same variance σ_d^2 , whereas for the second case we assume an a priori spatial covariance structure with correlations between the first neighbors. We remark that this is the structure of the “true” model error covariance used to generate the synthetic observations via (2.1) and (2.2). On the other hand, we do not assume any constraint on the structure of the covariance matrix in the estimation algorithm, with the exception of the positive semidefinite property (and hence symmetry) of the estimated matrix, which is satisfied by construction.

Following a similar procedure as for the linear model case (subsection 3.1), 50 independent realizations of this experiment were performed in order to show the estimator empirical distribution. The nonzero values of the initial guesses \mathbf{Q}_0 (a $N_x \times N_x$ matrix with the same structure of \mathbf{Q}_{true}) and σ_d^2 and σ_{sd} were sampled from $\mathcal{U}[0.05, 0.5]$ and $\mathcal{U}[0.01, 0.15]$, respectively.

In Figure 2 we show the empirical distribution of the estimator of \mathbf{Q} (top panel) for true parameter value $\mathbf{Q}_{true} = \sigma^2 \mathbf{I}$ with $\sigma^2 = 0.2$, $\mathbf{R} = 0.5 \mathbf{I}$, and $K = 500$ observations and the Frobenius norm $\|\mathbf{Q}_{true} - \hat{\mathbf{Q}}_{(s)}\|_F$ (bottom panel), for the algorithm proposed in this work, EM+VMPF (blue), the EM+VMPF+EnKS algorithm (yellow), and the EM+EnKF+EnKS algorithm (red) proposed by [24], which requires an ensemble Kalman filter and an ensemble Kalman smoother. Since the full covariance matrix is estimated, the Frobenius norm indicates both deviations from nonzero and zero values in the estimation. Thus, it is a good diagnostic of the overall structural errors in the estimation.

For ease of visualization, and just for plotting purposes, in this case of an isotropic uncorrelated covariance assumption $\mathbf{Q} = \sigma^2 \mathbf{I}$, at the s th iteration of the EM algorithm we kept $[\hat{\sigma}_d^2]_s$ as the average of the diagonal values of $\hat{\mathbf{Q}}_{(s)}$. As we repeated each experiment independently 50 times, we had a series of 50 values of $[\hat{\sigma}_d^2]_s$ at the s iteration of the EM algorithm to con-

struct a violin object that describes the empirical distribution of the corresponding estimator at each iteration.

With only $N_p = 20$ particles, the EM+VMPF and the EM+EnKF+EnKS methods provide good estimates of \mathbf{Q} , stabilizing in about 10 iterations with the median value (white circle) around the true diagonal value (top panel of Figure 2). The EM+VMPF (blue) algorithm produces estimates less biased than the EM+EnKF+EnKS (red) algorithm despite the fact that it does not require a particle smoother (and therefore uses less observational information in the state estimates). Moreover, the violin objects show that the empirical distribution for the EM+VMPF estimates is symmetric and highly concentrated around the true value, whereas the EM+EnKF+EnKS estimator empirical distributions have a greater dispersion, are not symmetric, and have tails towards higher values of σ_d^2 , meaning that in the 50 repetitions performed this method overestimates the value of σ_d^2 . On the other hand, the EM+VMPF+EnKS estimator empirical distribution has a large dispersion and does not reach the true value σ_d^2 .

The bottom panel of Figure 2 shows the empirical distribution of the Frobenius norm $\|\mathbf{Q}_{true} - \hat{\mathbf{Q}}_{(s)}\|_F$ for the three methods. The EM+EnKF+EnKS algorithm (red), despite having a greater dispersion, has a slightly better performance than the other two.

To examine the sensitivity of the proposed algorithm to different values of the observation error covariance \mathbf{R} a second set of experiments was performed. We assumed $\mathbf{R} = \sigma_R^2 \mathbf{I}$ with $\sigma_R^2 \in \{0.1, 0.5, 1.0, 2.0\}$, and $\mathbf{Q} = 0.2\mathbf{I}$. For each setting we performed 30 independent realizations of the same experiment. The mean and 95% confidence intervals for $[\hat{\sigma}_d^2]_s$ for each value of σ_R^2 are shown in Figure 3. The larger the \mathbf{R} values are, the harder it is to estimate model error, likely because of sampling noise. The results for $\sigma_R^2 = 1.0$ and 2.0 are noisier and demonstrate small biases, even after 25 iterations of the EM scheme. As clearly stated in [34] and references therein, the quality of reconstructed state vectors and estimation procedures when using variational or ensemble-based methods largely depends on the relative amplitudes between observation and model errors. In [44] the authors also mention that a small increment in the magnitude of the observation error \mathbf{R} significantly affects the accuracy and behavior of the method they propose to estimate the model error \mathbf{Q} iteratively in the observation space using the implicit equally weighted PF (IEWPF) [43]. Their method provides reasonable estimates of the diagonal values of \mathbf{Q}_{true} as long as the diagonal values of the observation error matrix \mathbf{R} are relatively small (they assume a diagonal \mathbf{R}), with $\sigma_R^2 = 0.2$ being the largest one. The larger σ_R^2 , the less accurate the estimation procedure. As shown in Figure 3, the estimation with EM+VMPF gives results with a good performance for observational variances ranging from 0.1 to 2.0. The estimation error increases with the increase of the observational variance; however, the estimation error is lower than 20% in all the cases shown.

Experiments designed for the estimation of a tridiagonal isotropic model error covariance are shown in Figure 4. The diagonal value of \mathbf{Q}_{true} was set to $\sigma_d^2 = 0.2$ and sub-/superdiagonal values were $\sigma_{sd}^2 = 0.05$, as defined in [44], $\mathbf{R} = 0.5\mathbf{I}$, $K = 500$, and $N_p = 20$. For this isotropic tridiagonal model error covariance assumption, at the s th iteration of the EM algorithm we computed $\hat{\sigma}_d^2$ as the average of the diagonal values of $\hat{\mathbf{Q}}_s$, and $\hat{\sigma}_{sd}^2$ as the average of the sub-/superdiagonal values of $\hat{\mathbf{Q}}_s$. The violin plots at the s th iteration of the EM algorithm were generated with $\hat{\sigma}_d^2$, $\hat{\sigma}_{sd}^2$ obtained for each realization of the experiment.

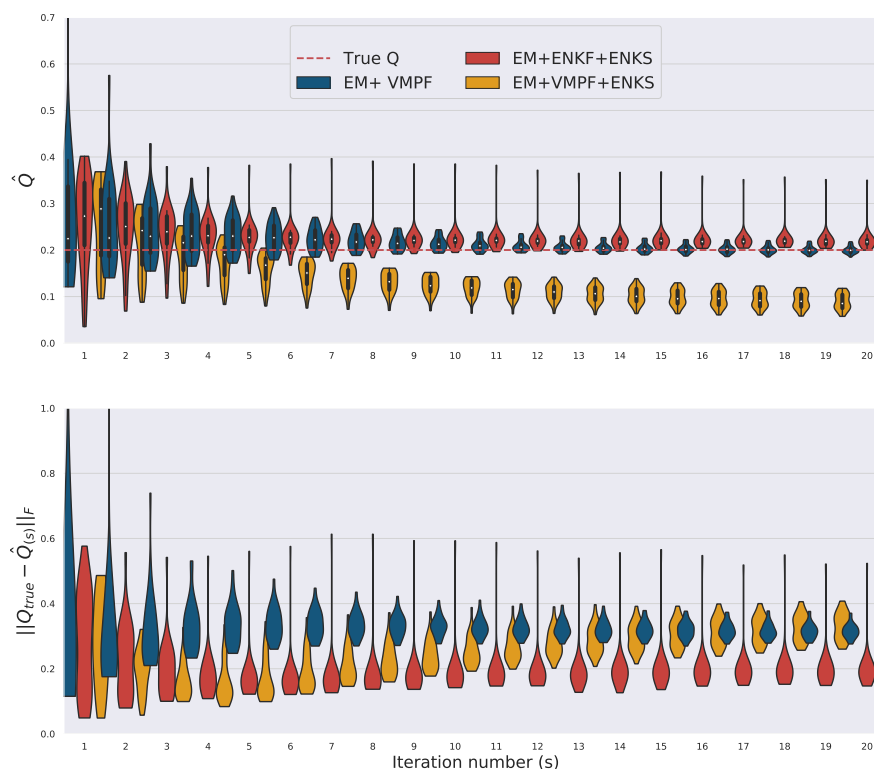


Figure 2. *EM+VMPF (blue), EM+VMPF+EnKS (yellow), and EM+EnKF+EnKS (red) performances when estimating \mathbf{Q} . Mean diagonal elements of $\hat{\mathbf{Q}}_{(s)}$ (top) and Frobenius norm $\|\mathbf{Q}_{true} - \hat{\mathbf{Q}}_{(s)}\|_F$ (bottom) for the 8 variables Lorenz-96 model (3.2) with true model error covariance $\mathbf{Q}_{true} = \sigma^2 \mathbf{I}$, $\sigma^2 = 0.2$, and $K = 500$ (subsection 3.2). The violin plots were generated by running 50 independent repetitions of the algorithm with $N_p = 20$ particles.*

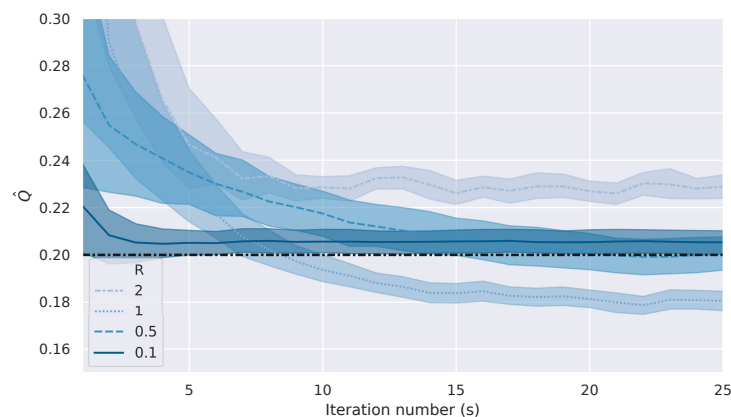


Figure 3. *Sensitivity of EM+VMPF to different values of observation error \mathbf{R} in the 8-variable Lorenz-96 model (subsection 3.2) with true model error covariance $\mathbf{Q}_{true} = \sigma^2 \mathbf{I}$, with $\sigma^2 = 0.2$, $N_x = 8$, and $K = 500$.*

As shown in Figure 4, the EM+VMPF (blue) and EM+EnKF+EnKS (red) methods converge rapidly to the true value σ_d^2 (top panel), whereas the hybrid EM+VMPF+EnKS method does not capture the variances. The empirical distribution of the EM+VMPF estimates shows a median value (white circle) closer to the true value of σ_d^2 but has a greater dispersion. However, as in the diagonal \mathbf{Q} case, the EM+EnKF+EnKS proposal tends to overestimate σ_d^2 . Both the EM+VMPF and EM+EnKF+EnKS methods show a similar performance when estimating the sub-/superdiagonal elements of \mathbf{Q} , with EM+VMPF being more biased in terms of the median value. However, after 17 iterations the violin object obtained by the EM+VMPF proposal is completely contained within the violin object obtained by the EM+EnKF+EnKS algorithm. The EM+VMPF+EnKS method never reaches the covariances between neighboring variables, underestimating them. A different behavior is observed for the empirical distribution of the Frobenius norm. In this case, EM+EnKF+EnKS outperforms the EM+VMPF estimation procedure, which in turn outperforms the EM+VMPF+EnKS performance. Thus, EM+VMPF has a better performance when estimating variance and covariances between neighbor variables, but it does not perform so well when estimating the null long correlations due to its nonparametric statistics.

Regarding computational costs and times, EM+VMPF and EM+VMPF+EnKS needed 3.1 minutes and 1.2 minutes, respectively, to complete 20 iterations of their EM versions. The extra cost in the former one is because it needs to do the iterations of the fixed point algorithm at each EM iteration. On average 4 fixed point iterations are required even when we set the maximum number of iterations to 6; i.e., the convergence criterion is met in 4 iterations on average (see Algorithm 2.1). Therefore, the overall extra cost of the smoother-free EM in this experiment is about 3 times the classical batch EM using the same number of particles.

We also evaluated the performance of the proposed algorithm using less frequent observations. This means that longer model integrations are required, which in turn enhances nonlinearities, resulting in non-Gaussian forecast distributions. We used the 8-variable Lorenz-96 model assuming observations are available every $\Delta_t = 0.25$. In all the previous experiments with the Lorenz-96 model, we assumed observations were every $\Delta_t = 0.05$. The synthetic observations were generated with a true tridiagonal model error covariance matrix \mathbf{Q} , with diagonal values $\sigma_d^2 = 0.5$ and sub/superdiagonal values $\sigma_{sd} = 0.1$, $\mathbf{R} = 0.5\mathbf{I}_8$, $K = 500$. Following a procedure similar to that for previous experiments and settings, independent realizations of the current experiment were performed, and \mathbf{Q}_{true} was estimated using the EM+VMPF, EM+EnKF+EnKS, and EM+VMPF+EnKS algorithms with $N_p = 20$ particles. The mean values of $\hat{\sigma}_d^2$ and $\hat{\sigma}_{sd}$ at the s th iteration of the algorithm are shown in Figure 5. The EM+VMPF (blue solid line) algorithm recovers the structure of \mathbf{Q} better than the EM+EnKF+EnKS (red dashed lines) and the EM+VMPF+EnKS (yellow dot-dashed lines) algorithms. The EM+VMPF algorithm outperforms the other two algorithms when estimating the variances σ_d^2 , being EM+EnKF+EnKS and EM+VMPF+EnKS highly biased. The three methods show similar performances when estimating the sub/superdiagonal values of \mathbf{Q} , σ_{sd} .

The low performance of EM+EnKF+EnKS in this estimation problem appears to be associated to the Gaussianity assumption. The non-Gaussianity resulting from long assimilation cycles has a negative impact on the Gaussian-based algorithms, with the EM+EnKF+EnKS performance being the most affected. On the other hand, the smoother-free algorithm based

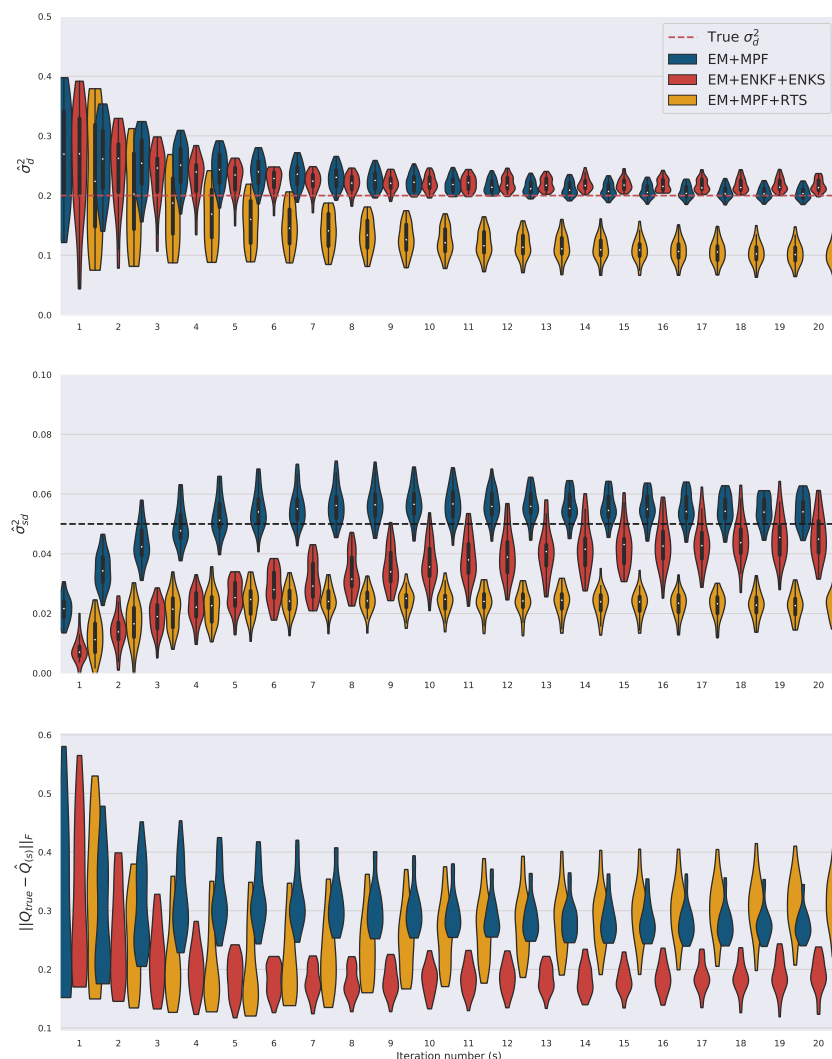


Figure 4. Estimation of \mathbf{Q} as a function of the iteration for the EM+VMPF (blue), EM+EnKF+EnKS (red), EM+VMPF+EnKS (yellow) for the Lorenz-96 model with tridiagonal true model error covariance \mathbf{Q}_{true} , with $\sigma_d^2 = 0.2$, $\sigma_{sd}^2 = 0.05$, $N_x = 8$, and $K = 500$ (subsection 3.2). The violin plots were generated by running 50 independent repetitions of these algorithms with $N_p = 20$ particles. Upper panels show the average of the estimated diagonal values of $\hat{\mathbf{Q}}_s$ (top panel) and the average of the sub/superdiagonal values of $\hat{\mathbf{Q}}_s$ (middle panel). The Frobenius norm $\|\mathbf{Q}_{true} - \hat{\mathbf{Q}}_{(s)}\|_F$ is shown in the bottom panel as function of the algorithm iteration.

on a PF, EM+VMPF, can capture the non-Gaussian densities and gives a relatively good estimation. A similar result was found in [7] for an online EM algorithm. They show that importance-sampling online EM with a PF has a much better performance than a Gaussian smoother EM for long assimilation cycles.

A higher-dimensional experiment was also performed for the chaotic Lorenz-96 model with 40 dimensions and $F = 8$, where a 40×40 model error covariance matrix has to be estimated.

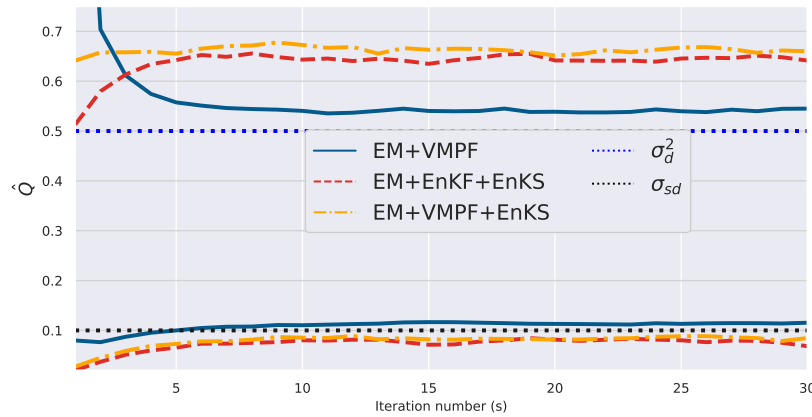


Figure 5. Estimation of \mathbf{Q} as a function of the iteration, for the 8-variable Lorenz-96 model using $\Delta_t = 0.25$ with \mathbf{Q}_{true} a tridiagonal matrix, $\sigma_d^2 = 0.5$, $\sigma_{sd} = 0.1$, $\mathbf{R} = 0.5 \mathbf{I}$, $K = 500$ (subsection 3.2). Each line shows the mean of the diagonal elements of $\hat{\mathbf{Q}}_s$ and the mean of the subdiagonal elements $\hat{\sigma}_{(sd)}$ for the EM+VMPF (blue solid line), the EM+EnKF+EnKS (red dashed line), and EM+VMPF+EnKS (yellow dot-dashed line) algorithms with $N_p = 20$ particles. True values of $\sigma_d^2 = 0.5$ (blue) and $\sigma_{sd} = 0.1$ (black) are shown as horizontal dotted lines.

We prescribed an isotropic tridiagonal model error covariance matrix \mathbf{Q}_{true} , with diagonal values $\sigma_d^2 = 0.2$ and both sub-/superdiagonal values equal to $\sigma_{sd}^2 = 0.05$ to generate the synthetic observations. Following a procedure similar to that used in the previous experiments, we performed independent repetitions of this experiment and estimated \mathbf{Q} using the EM+VMPF, the EM+VMPF+EnKS, and the EM+EnKF+EnKS algorithms. We used $N_p = 20$ particles for the EM+VMPF smoother-free algorithm, whereas the EM+EnKF+EnKS and the EM+VMPF+EnKS algorithms required $N_p = 50$ ensemble members/particles in order to avoid smoother divergence. We note that early experiments we conducted with the 40-variables Lorenz-96 model using $N_p = 20$ particles with both EM+EnKF+EnKS and EM+VMPF+EnKS showed divergence.

The observation error covariance was set to $\mathbf{R} = 0.5\mathbf{I}$, and the batch of observations is $K = 500$. At the s th iteration of the EM algorithm we computed $\hat{\sigma}_d^2$ as the average of the diagonal values of $\hat{\mathbf{Q}}_s$, $\hat{\sigma}_{sd}^2$ as the average of the sub-/superdiagonal values of $\hat{\mathbf{Q}}_s$ and the Frobenius norm $\|\mathbf{Q}_{true} - \hat{\mathbf{Q}}_{(s)}\|_F$. The results of these estimates at the s th iteration of the algorithm are shown in the top, middle, and bottom panels, respectively, of Figure 6. The performances of the EM+VMPF (blue) and the EM+EnKF+EnKS (red) algorithms are similar in terms of the estimation of the diagonal values of \mathbf{Q} and Frobenius norm. When estimating the diagonal elements of \mathbf{Q} both algorithms tend to stabilize after 10 iterations overestimating them, with the EM+VMPF algorithm being the less biased. EM+VMPF shows an excellent performance when estimating the sub/superdiagonal elements σ_{sd} , while EM+EnKF+EnKS slightly underestimates it.

The full covariance matrix is estimated at each iteration, and the Frobenius norm indicates both deviations from nonzero and zero values in the estimation. As shown in the bottom panel of Figure 6, the overall structural errors in the estimation of \mathbf{Q} by the EM+VMPF and the

EM+EnKF+EnKS algorithms are similar. On the other hand, EM+VMPF+EnKS (yellow) has a poor performance, underestimating σ_d^2 and σ_{sd} . We remark here that EM+VMPF uses $N_p=20$ particles, while EM+EnKF+EnKS and EM+VMPF+EnKS require $N_p=50$ particles.

Regarding computational times, the experiments took an average of 53.2 minutes and 93.9 minutes to complete 20 iterations for EM+VMPF and EM+VMPF+EnKS, respectively. All the experiments were run in Python on a laptop with an i5 processor. In this larger-dimensional experiment the number of iterations in the fixed point algorithm is again about 4; however, due to the larger number of particles required for the smoother (50 particles), EM+VMPF presents a lower computational cost. The required number of particles is an important factor in practice for medium and large systems.

Despite the increase in the dimensionality, the experiments' results show that the smoother-free version of the EM algorithm for a PF provides unbiased results when estimating the diagonal and sub/superdiagonal values of \mathbf{Q} , showing a better performance than the hybrid EM version for a PF combined with a Gaussian Rauch–Tung–Striebel (RTS) smoother (EM+VMPF+EnKS).

4. Conclusions. In this work a novel method to estimate the model error uncertainty in dynamical systems is introduced and evaluated. It assumes that both the model and observation errors are additive and Gaussian with zero mean and covariance matrices \mathbf{Q} and \mathbf{R} , respectively. The methodology here presented is based on the maximization of a likelihood criterion using the principles of the EM algorithm and a particle filter. We aim at maximizing the complete likelihood of the observations by marginalizing this likelihood function. The resulting likelihood is expressed sequentially. By taking this approach, in the E-Step of the EM algorithm we only have to compute filtering densities, avoiding the need to compute smoothing densities, which are known to be one of the main drawbacks when using particle filters in data assimilation. The trade-off of avoiding the need of a particle smoother in the E-Step of the EM algorithm is the need to solve an implicit equation for \mathbf{Q} in the M-Step of the EM algorithm. This problem was tackled by means of a fixed point algorithm, and despite the fact that an analytical proof of its convergence is not straightforward to obtain due to the nonlinearity of the model dynamics, empirical results show that it converges to a solution of this implicit equation.

The EM algorithm coupled with the VMPF presents, in general, an overall excellent performance. It gives very promising results in the experiments performed with a simple linear first order autoregressive system and a chaotic Lorenz-96 system with 8 and 40 variables. In the first case, results were compared with those obtained by different methods already proposed in the literature, showing a good performance in terms of bias and root mean squared error (RMSE). The new method is suitable for non-Gaussian posterior densities from nonlinear dynamical and observational models, unlike the Kalman filter/smoothing and its ensemble variants.

The proposed smoother-free variant of the EM algorithm with the particle filter is compared with a classical EM algorithm that uses the particle filter and a Gaussian smoother. The experiments also present a comparison with the EM algorithm coupled with the ensemble Kalman filter. The Gaussian assumption in the Kalman filter leads to an analytical solution of the maximization step. In this sense, the combination of the EM with the Kalman filter

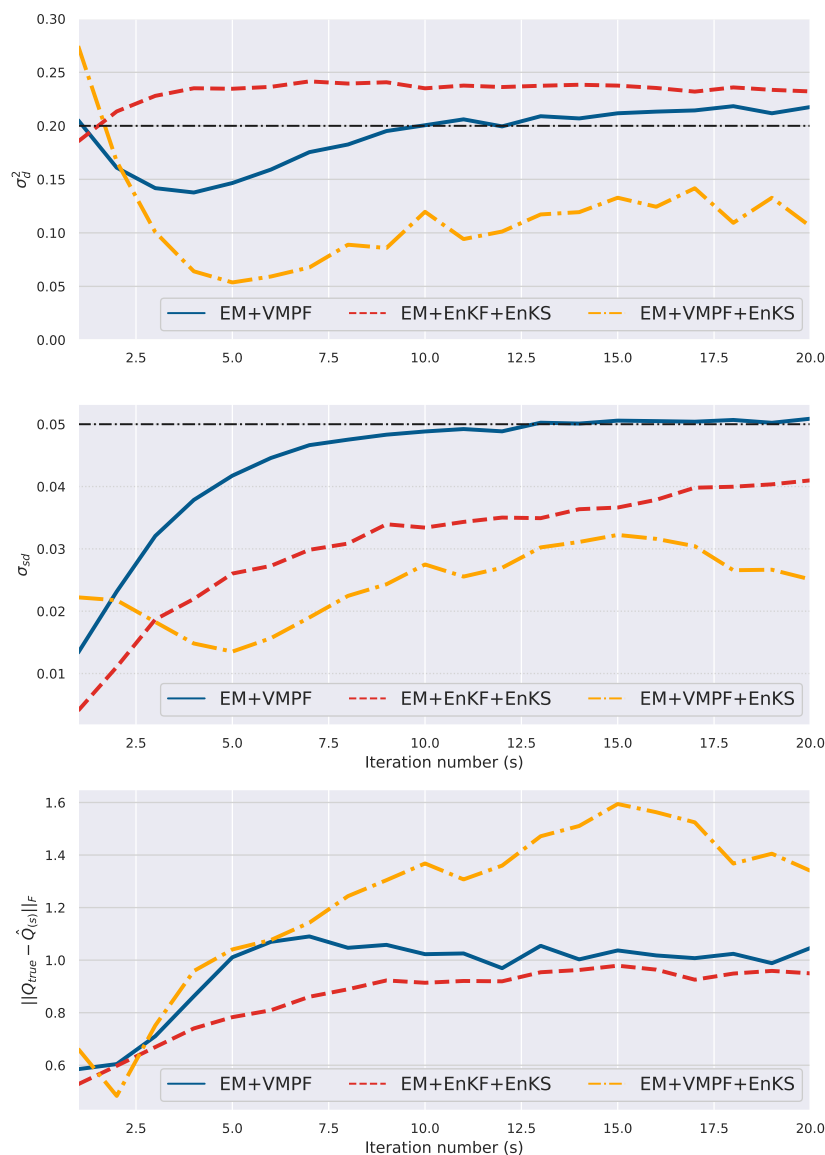


Figure 6. Estimation of \mathbf{Q} as a function of the iteration, for the 40-variable Lorenz-96 model with $\mathbf{Q} = 0.2 \mathbf{I}$, $\mathbf{R} = 0.5 \mathbf{I}$, $K = 500$ (subsection 3.2) for the EM+VMPF (blue solid line) algorithm with $N_p = 20$ particles, and the EM+EnKF+EnKS (red dashed line), and EM+VMPF+EnKS (yellow dot-dashed line) algorithms with $N_p = 50$ particles each. Each line shows the mean of the diagonal elements of $\hat{\mathbf{Q}}_s$ (top panel), the mean of the subdiagonal elements $\hat{\sigma}_{(sd)}$ (middle panel), and Frobenius norm $\|\mathbf{Q}_{true} - \hat{\mathbf{Q}}_{(s)}\|_F$ (bottom panel). True values of $\sigma_d^2 = 0.2$ and $\sigma_{sd} = 0.05$ are shown as horizontal dot-dashed lines in black.

could be considered as a gold standard for the proposed algorithm to achieve. The results of the experiments showed that the smoother-free algorithm coupled with the particle filter performance may be as good as the performance of the EM with the Kalman filter. In a more equally based comparison, the smoother-free variant of the EM algorithm gives a performance

that is superior to that of the classical EM algorithm coupled with the particle filter and a Gaussian smoother.

In the case of the Lorenz-96 system, its performance was tested for different scenarios showing good convergence properties, even for less frequent observations. It is stable even for $\mathbf{R} = 10\mathbf{Q}$, although a small bias appears in the estimate. The new method outperforms the traditional EM algorithm with the EnKS [14] for a diagonal \mathbf{Q} and for the diagonal of a tridiagonal \mathbf{Q} . However, off-diagonal elements' estimates were always noisier than those using an EnKS. It also works in moderately large parameter estimation problems of dimension 40×40 and of state space 40. In the case of less frequently observed data, longer model integrations are required, which result in non-Gaussian forecast distributions. The smoother-free EM algorithm coupled with the VMPF showed better performance in this scenario than methods based on the ensemble Kalman filter and smoothers.

In general we found that the extra computational cost of the smoother-free EM for the same number of particles is about 3–4 times the classical batch EM due to the fixed point algorithm iterations. However, in practice a smaller number of particles may be required if the smoother step is avoided, substantially diminishing the extra cost. At the current form, the algorithm here proposed would be expensive (in computational time and also in terms of memory requirements) for state dimensions larger than $N_x > 100$ and prohibitive for $N_x > 1000$, since the algorithm uses the explicit form of \mathbf{Q} and its inverse is needed at the M-Step.

The computational cost of the algorithm here proposed is directly related to the number of iterations needed for convergence. All the experiments performed achieved convergence to a narrow neighborhood of the true value of \mathbf{Q} in as few as 10–15 iterations of the EM algorithm. Each EM iteration requires the computation of K filtering densities computed by using a particle filter with N_p particles, whereas the M-Step requires solving a fixed point algorithm. In our experiments, we set the maximum number of iterations for the fixed point algorithm to 6. In practice, an average of 4 fixed point iterations were enough to satisfy the stopping criterion. In turn, each of these fixed point iterations also require computing K filtering densities computed by using a particle filter with N_p particles.

The numerical results obtained in the conducted experiments and these last considerations about computational costs suggest that the method works well for moderately large models, and in the examined situations even better and faster than classical schemes, but as previously mentioned, it would be too expensive in its current implementation for large-dimensional models.

Model error covariances are essential in particle flow filters. The conducted experiments show that these particle flow filters, in particular VMPF, can work with an adaptive model error without a priori information on this covariance, whereas in previous studies a fixed known model error covariance was used [25]. Model error covariances impact on the prior density and also on the kernel covariance in the VMPF. The overall excellent performance of the estimates may also be attributed to the strong sensitivity of VMPF performance to model error covariance. In this sense, there is a positive feedback between the model error covariance estimates of the EM algorithm and the state estimates of the filter.

Appendix A. Derivation of $\mathcal{G}(\theta', \theta)$. As explained in subsection 2.3, the likelihood of the observations can be decomposed as $p(\mathbf{y}_{1:K}; \theta) = \prod_{k=1}^K p(\mathbf{y}_k | \mathbf{y}_{1:k-1}; \theta)$, with the convention $\mathbf{y}_{1:0} = \{\emptyset\}$. Marginalizing this last expression, we obtain

$$\begin{aligned} p(\mathbf{y}_{1:K}; \theta) &= \prod_{k=1}^K p(\mathbf{y}_k | \mathbf{y}_{1:k-1}; \theta) \\ (A.1) \quad &= \prod_{k=1}^K \int p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta) d\mathbf{x}_k, \end{aligned}$$

and taking the logarithm we can rewrite this last expression as

$$l(\theta) = \log p(\mathbf{y}_{1:K}; \theta) = \log \prod_{k=1}^K \int q_{\theta'}(\mathbf{x}_k) \frac{p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)}{q_{\theta'}(\mathbf{x}_k)} d\mathbf{x}_k,$$

where $q_{\theta'}(\mathbf{x}_k)$ is a pdf whose support includes the support of the likelihood of the observations. In principle, θ' is not necessarily equal to θ . Using Jensen's inequality,

$$(A.2) \quad l(\theta) \geq \sum_{k=1}^K \int q_{\theta'}(\mathbf{x}_k) \log \left(\frac{p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)}{q_{\theta'}(\mathbf{x}_k)} \right) d\mathbf{x}_k.$$

Let $\mathcal{G}(q, \theta)$ be an intermediate function defined as

$$(A.3) \quad \mathcal{G}(q, \theta) = \sum_{k=1}^K \int q(\mathbf{x}_k) \log \left(\frac{p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)}{q(\mathbf{x}_k)} \right) d\mathbf{x}_k.$$

Using Bayes's rule, the recursive posterior density at time k is

$$(A.4) \quad p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta) = \frac{p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)}{p(\mathbf{y}_k | \mathbf{y}_{1:k-1}; \theta)}.$$

If we choose $q(\mathbf{x}_k) = p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta)$, and so replace (A.4) in (A.3), it can be shown that $\mathcal{G}(q, \theta) = l(\theta)$. That means that for fixed $\theta = \theta'$, the function q that maximizes the intermediate function, \mathcal{G} , is the recursive posterior density $p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta)$. This density can be inferred by a filter method, and this corresponds to the Expectation Step.

On the other hand, maximizing $\mathcal{G}(q, \theta)$ with respect to θ gives a lower bound of $l(\theta)$.

The Maximization Step consists in maximizing $\mathcal{G}(q', \theta)$ as a function of θ , where q' is the density obtained in the Expectation Step. If we now write $q_{\theta'}(\mathbf{x}_k) = p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta')$ and make an abuse of notation in the expression $\mathcal{G}(q_{\theta'}, \theta)$ by replacing $q_{\theta'}$ by the parameter that identifies

it, then

$$(A.5) \quad \mathcal{G}(\theta', \theta) = \sum_{k=1}^K \int p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta') \log \left(\frac{p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)}{p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta')} \right) d\mathbf{x}_k$$

$$(A.6) \quad = \sum_{k=1}^K \int p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta') \log (p(\mathbf{y}_k | \mathbf{x}_k; \theta) p(\mathbf{x}_k | \mathbf{y}_{1:k-1}; \theta)) d\mathbf{x}_k$$

$$(A.7) \quad - \sum_{k=1}^K \int p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta') \log (p(\mathbf{x}_k | \mathbf{y}_{1:k}; \theta')) d\mathbf{x}_k.$$

Appendix B. Equation for \mathbf{Q} . We want to find the root of $\frac{\partial}{\partial \mathbf{Q}} \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q}) = 0$, where $\mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q})$:

$$\mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q}) \doteq \sum_{k=1}^K \sum_{j=1}^{N_p} w_{k, \mathbf{Q}_{s-1}}^{(j)} \log \left(\sum_{i=1}^{N_p} w_{k-1, \mathbf{Q}}^{(i)} \phi \left(x_k^{(j)}, \mathcal{M}(\mathbf{x}_{k-1}^{(i)}), \mathbf{Q} \right) \right)$$

and $\phi(\mathbf{x}_k, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{N_x/2} |\boldsymbol{\Sigma}|^{1/2}} \exp \left\{ -\frac{1}{2} (\mathbf{x}_k - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_k - \boldsymbol{\mu}) \right\}$.

Denoting $\boldsymbol{\beta}_k^{(i,j)} = \mathbf{x}_k^{(j)} - \mathcal{M}(\mathbf{x}_{k-1}^{(i)})$, we have

$$\begin{aligned} \frac{\partial}{\partial \mathbf{Q}} \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q}) &= \sum_{k=1}^K \sum_{j=1}^{N_p} w_{k, \mathbf{Q}_{s-1}}^{(j)} \frac{\partial}{\partial \mathbf{Q}} \log \left[\sum_{i=1}^{N_p} \frac{w_{k-1}^{(i)}}{(2\pi)^{N_x/2} |\mathbf{Q}|^{1/2}} \right. \\ &\quad \times \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta}_k^{(i,j)})^T \mathbf{Q}^{-1} \boldsymbol{\beta}_k^{(i,j)} \right\} \left. \right] \\ &= \sum_{k=1}^K \sum_{j=1}^{N_p} w_{k, \mathbf{Q}_{s-1}}^{(j)} \frac{\partial}{\partial \mathbf{Q}} \left[-\frac{N_x}{2} \log(2\pi) - \frac{1}{2} \log |\mathbf{Q}| \right. \\ &\quad \left. + \log \left(\sum_{i=1}^{N_p} w_{k-1}^{(i)} \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta}_k^{(i,j)})^T \mathbf{Q}^{-1} \boldsymbol{\beta}_k^{(i,j)} \right\} \right) \right] \\ &= -\frac{K}{2} \mathbf{Q}^{-1} + \frac{1}{2} \mathbf{Q}^{-1} \left[\sum_{k=1}^K \sum_{j=1}^{N_p} w_{k, \mathbf{Q}_{s-1}}^{(j)} \frac{1}{S_i} \sum_{i=1}^{N_p} w_{k-1}^{(i)} \right. \\ &\quad \times \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta}_k^{(i,j)})^T \mathbf{Q}^{-1} \boldsymbol{\beta}_k^{(i,j)} \right\} \boldsymbol{\beta}_k^{(i,j)} \boldsymbol{\beta}_k^{(i,j)T} \left. \right] \mathbf{Q}^{-1}, \end{aligned}$$

where $S_i = \sum_{i=1}^{N_p} w_{k-1}^{(i)} \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta}_k^{(i,j)})^T \mathbf{Q}^{-1} \boldsymbol{\beta}_k^{(i,j)} \right\}$.

Thus, \mathbf{Q} that satisfies $\frac{\partial}{\partial \mathbf{Q}} \mathcal{G}(\mathbf{Q}_{s-1}, \mathbf{Q}) = 0$ is given by

$$(B.1) \quad \mathbf{Q} = \frac{1}{K} \sum_{k=1}^K \left[\sum_{j=1}^{Np} w_{k, \mathbf{Q}_{s-1}}^{(j)} \left(\frac{1}{S_i} \sum_{i=1}^{Np} w_{k-1, Q}^{(i)} \psi_k^{j,i} \mathbf{B}_k^{j,i} \right) \right],$$

where $\psi_k^{j,i} = \exp\{-\frac{1}{2}(\beta_k^{(i,j)})^T \mathbf{Q}^{-1} \beta_k^{(i,j)}\}$ and $\mathbf{B}_k^{j,i} = \beta_k^{(i,j)} \beta_k^{(i,j)^T}$.

Acknowledgments. We would like to acknowledge the editor and the two reviewers for a thorough evaluation of the work and for their constructive and insightful comments. They have certainly contributed to improving the quality of this article.

REFERENCES

- [1] C. ANDRIEU AND A. DOUCET, *Online expectation-maximization type algorithms for parameter estimation in general state space models*, in Proceedings of the 2003 IEEE Conferences on Acoustics, Speech and Signal Processing (ICASSP '03) (Hong Kong, China), Vol. VI, IEEE, Washington, DC, 2003, pp. VI–69.
- [2] M. S. ARULAMPALAM, S. MASKELL, N. GORDON, AND T. CLAPP, *A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking*, IEEE Trans. Signal Process., 50 (2002), pp. 174–188.
- [3] M. BRIERS, A. DOUCET, AND S. MASKELL, *Smoothing algorithms for state-space models*, Ann. Inst. Statist. Math., 62 (2010), pp. 61–89.
- [4] P. BUNCH AND S. GODSILL, *Approximations of the optimal importance density using Gaussian particle flow importance sampling*, J. Amer. Statist. Assoc., 111 (2016), pp. 748–762.
- [5] O. CAPPÉ, E. MOULINES, AND T. RYDEN, *Inference in Hidden Markov Models*, Springer, New York, 2005.
- [6] A. CARRASSI, M. BOCQUET, A. HANNART, AND M. GHIL, *Estimating model evidence using data assimilation*, Quart. J. Roy. Meteorol. Soc., 143 (2017), pp. 866–880.
- [7] T. COCUCCI, M. PULIDO, M. LUCINI, AND P. TANDEO, *Model error covariance estimation in particle and ensemble Kalman filters using an online expectation-maximization algorithm*, Quart. J. Roy. Meteorol. Soc., 147 (2021), pp. 526–543, <https://doi.org/10.1002/qj.3931>.
- [8] R. DALEY, *Atmospheric Data Analysis*, Cambridge Atmospheric Space Sci. Ser. 2, Cambridge University Press, Cambridge, UK, 1993.
- [9] F. DAUM, J. HUANG, AND A. NOUSHIN, *Exact particle flow for nonlinear filters*, in Signal Processing, Sensor Fusion, and Target Recognition XIX, Proc. Vol. 7697, SPIE, Bellingham, WA, 2010, 769704.
- [10] A. P. DEMPSTER, N. M. LAIRD, AND D. B. RUBIN, *Maximum likelihood from incomplete data via the EM algorithm*, J. Roy. Statist. Soc. Ser. B, 39 (1977), pp. 1–38.
- [11] R. DOUC, E. MOULINES, AND D. STOFFER, *Nonlinear Time Series: Theory, Methods and Applications with R Examples*, Chapman & Hall/CRC, Boca Raton, FL, 2014.
- [12] A. DOUCET AND A. M. JOHANSEN, *A tutorial on particle filtering and smoothing: Fifteen years later*, in The Oxford Handbook of Nonlinear Filtering, Oxford University Press, Oxford, UK, 2011, pp. 656–704.
- [13] M. DOWD, E. JONES, AND J. PARSLow, *A statistical overview and perspectives on data assimilation for marine biogeochemical models*, Environmetrics, 25 (2014), pp. 203–213.
- [14] D. DREANO, P. TANDEO, M. PULIDO, B. AIT-EL-FQUIH, T. CHONAVEL, AND I. HOTEIT, *Estimating model-error covariances in nonlinear state-space models using Kalman smoothing and the expectation-maximization algorithm*, Quart. J. Roy. Meteorol. Soc., 143 (2017), pp. 1877–1885.
- [15] N. J. GORDON, D. J. SALMOND, AND A. F. SMITH, *Novel approach to nonlinear/non-Gaussian Bayesian state estimation*, IEE Proc. F, 140 (1993), pp. 107–113.
- [16] N. GUPTA AND R. MEHRA, *Computational aspects of maximum likelihood estimation and reduction in sensitivity function calculations*, IEEE Trans. Automat. Control, 19 (1974), pp. 774–783.

- [17] N. KANTAS, A. DOUCET, S. S. SINGH, J. MACIEJOWSKI, AND N. CHOPIN, *On particle methods for parameter estimation in state-space models*, *Statist. Sci.*, 30 (2015), pp. 328–351.
- [18] K. LAW, A. STUART, AND K. ZYGALAKIS, *Data Assimilation*, Springer, New York, 2015.
- [19] E. N. LORENZ, *Predictability: A problem partly solved*, in *Proceedings of the Seminar on Predictability*, Vol. 1, ECMWF, Shinfield Park, Reading, UK, 1995.
- [20] X. LUO AND T. BHAKTA, *Estimating observation error covariance matrix of seismic data from a perspective of image denoising*, *Comput. Geosci.*, 21 (2017), pp. 205–222.
- [21] X. LUO, R. J. LORENTZEN, A. S. STORDAL, AND G. NÆVDAL, *Toward an enhanced Bayesian estimation framework for multiphase flow soft-sensing*, *Inverse Problems*, 30 (2014), 114012, <https://doi.org/10.1088/0266-5611/30/11/114012>.
- [22] P. MATISKO AND V. HAVLENA, *Noise covariance estimation for Kalman filter tuning using Bayesian approach and Monte Carlo*, *Internat. J. Adapt. Control Signal Process.*, 27 (2013), pp. 957–973.
- [23] J. OLSSON, O. CAPPÉ, R. DOUC, AND E. MOULINES, *Sequential Monte Carlo smoothing with application to parameter estimation in nonlinear state space models*, *Bernoulli*, 14 (2008), pp. 155–179.
- [24] M. PULIDO, P. TANDEO, M. BOCQUET, A. CARRASSI, AND M. LUCINI, *Stochastic parameterization identification using ensemble Kalman filtering combined with maximum likelihood methods*, *Tellus A: Dynamic Meteorology and Oceanography*, 70 (2018), pp. 1–17.
- [25] M. PULIDO AND P. J. VAN LEEUWEN, *Kernel embedding of maps for sequential bayesian inference: The variational mapping particle filter*, *J. Comput. Phys.*, 396 (2019), pp. 400–415.
- [26] M. PULIDO, P. J. VAN LEEUWEN, AND D. POSSELT, *Kernel embedded nonlinear observational mappings in the variational mapping particle filter*, in *Computational Science—ICCS 2019, Part IV, Lecture Notes in Comput. Sci.* 11539, Springer, Cham, 2019, pp. 141–155.
- [27] C. ROBERT AND G. CASELLA, *Monte Carlo Statistical Methods*, 2nd ed., Springer, New York, 2004.
- [28] N. SANTITISSADEEKORN AND C. JONES, *Two-stage filtering for joint state-parameter estimation*, *Monthly Weather Rev.*, 143 (2015), pp. 2028–2042.
- [29] G. SCHEFFLER, J. RUIZ, AND M. PULIDO, *Inference of stochastic parameterizations for model error treatment using nested ensemble Kalman filters*, *Quart. J. Roy. Meteorolog. Soc.*, 145 (2019), pp. 2028–2045.
- [30] R. SHUMWAY AND D. STOFFER, *Time Series Analysis and Its Applications: With R Examples*, Springer, New York, 2010.
- [31] R. H. SHUMWAY AND D. S. STOFFER, *An approach to time series smoothing and forecasting using the EM algorithm*, *J. Time Ser. Anal.*, 3 (1982), pp. 253–264.
- [32] L. M. STEWART, S. L. DANCE, AND N. K. NICHOLS, *Data assimilation with correlated observation errors: Experiments with a 1-D shallow water model*, *Tellus A: Dynamic Meteorology and Oceanography*, 65 (2013), 19546.
- [33] J. STROUD AND T. BENGTTSSON, *Sequential state and variance estimation within the ensemble Kalman filter*, *Monthly Weather Rev.*, 135 (2007), pp. 3194–3208.
- [34] P. TANDEO, P. AILLIOT, M. BOCQUET, A. CARRASSI, T. MIYOSHI, M. PULIDO, AND Y. ZHEN, *Joint estimation of model and observation error covariance matrices in data assimilation: a review*, *Monthly Weather Rev.*, 148 (2020), pp. 3973–3994.
- [35] G. UENO, T. HIGUCHI, T. KAGIMOTO, AND N. HIROSE, *Maximum likelihood estimation of error covariances in ensemble-based filters and its application to a coupled atmosphere–ocean model*, *Quart. J. Roy. Meteorolog. Soc.*, 136 (2010), pp. 1316–1343.
- [36] G. UENO AND N. NAKAMURA, *Iterative algorithm for maximum-likelihood estimation of the observation-error covariance matrix for ensemble-based filters*, *Quart. J. Roy. Meteorolog. Soc.*, 140 (2014), pp. 295–315.
- [37] G. UENO AND N. NAKAMURA, *Bayesian estimation of the observation-error covariance matrix in ensemble-based filters*, *Quart. J. Roy. Meteorolog. Soc.*, 142 (2016), pp. 2055–2080.
- [38] P. J. VAN LEEUWEN, *Particle filtering in geophysical systems*, *Monthly Weather Rev.*, 137 (2009), pp. 4089–4114.
- [39] J. WALLER, S. DANCE, AND N. NICHOLS, *On diagnosing observation-error statistics with local ensemble data assimilation*, *Quart. J. Roy. Meteorolog. Soc.*, 143 (2017), pp. 2677–2686.
- [40] G. C. WEI AND M. A. TANNER, *A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms*, *J. Amer. Statist. Assoc.*, 85 (1990), pp. 699–704.

- [41] C. K. WIKLE AND L. M. BERLINER, *A Bayesian tutorial for data assimilation*, Phys. D, 230 (2007), pp. 1–16.
- [42] C.-F. J. WU, *On the convergence properties of the EM algorithm*, Ann. Statist., 11 (1983), pp. 95–103.
- [43] M. ZHU, P. J. VAN LEEUWEN, AND J. AMEZCUA, *Implicit equal-weights particle filter*, Quart. J. Roy. Meteorolog. Soc., 142 (2016), pp. 1904–1919.
- [44] M. ZHU, P. J. VAN LEEUWEN, AND W. ZHANG, *Estimating model error covariances using particle filters*, Quart. J. Roy. Meteorolog. Soc., 144 (2018), pp. 1310–1320.