

Machine learning-based conditional mean filter: a generalization of the ensemble Kalman filter for nonlinear data assimilation

Truong-Vinh Hoang¹, Sebastian Krumscheid¹, Hermann G. Matthies², and
Raúl Tempone³

¹Chair of Mathematics for Uncertainty Quantification, RWTH Aachen University,
Germany, hoang@uq.rwth-aachen.de

²Technische Universität Braunschweig, Germany

³Computer, Electrical and Mathematical Sciences and Engineering, KAUST, and
Alexander von Humboldt professor in Mathematics of Uncertainty Quantification,
RWTH Aachen University.

August 2, 2022

Abstract

This paper presents the machine learning-based ensemble conditional mean filter (ML-EnCMF) — a filtering method based on the conditional mean filter (CMF) previously introduced in the literature. The updated mean of the CMF matches that of the posterior, obtained by applying Bayes' rule on the filter's forecast distribution. Moreover, we show that the CMF's updated covariance coincides with the expected conditional covariance. Implementing the EnCMF requires computing the conditional mean (CM). A likelihood-based estimator is prone to significant errors for small ensemble sizes, causing the filter divergence. We develop a systematic methodology for integrating machine learning into the EnCMF based on the CM's orthogonal projection property. First, we use a combination of an artificial neural network (ANN) and a linear function, obtained based on the ensemble Kalman filter (EnKF), to approximate the CM, enabling the ML-EnCMF to inherit EnKF's advantages. Secondly, we apply a suitable variance reduction technique to reduce statistical errors when estimating loss function. Lastly, we propose a model selection procedure for element-wisely selecting the applied filter, *i.e.*, either the EnKF or ML-EnCMF, at each updating step. We demonstrate the ML-EnCMF performance using the Lorenz-63 and Lorenz-96 systems and show that the ML-EnCMF outperforms the EnKF and the likelihood-based EnCMF.

Keywords: nonlinear filter, inverse problem, conditional expectation, weather forecast, deep learning

1 Introduction

Data assimilation combines numerical models and observations of a dynamical system to infer its states [38, 5, 39, 15, 21, 45]. This approach of integrating data into dynamical models is essential in different applications, *e.g.*, numerical weather prediction, environmental forecasting, and digital twins [5, 24, 41, 46, 20, 6, 36, 35]. Filtering is a sequential data assimilation technique, usually comprising two steps: a) forecasting the states using the numerical models up to an observation event (prediction step) and b) updating the forecasts by conditioning on new observations (update step). Ideally, the update step should be performed using Bayes' rule, where the forecasting states are encoded as a prior distribution, and the updated states are obtained as the Bayesian posterior. However, it is computationally difficult to accurately represent the posterior distribution in the case of nonlinear systems, even prohibitive for high-dimensional state spaces [2, 27, 43].

The ensemble Kalman filter (EnKF) extends the KF for nonlinear settings [20, 15, 25, 44, 19]. The EnKF uses an ensemble to approximate the forecast distribution, which is then updated by applying the KF formulation to each ensemble member. Because the EnKF uses a linear updating formulation, the updated ensemble mean is generally biased with respect to the Bayesian posterior mean. Due to its linear update map, EnKF exhibits poor performance in tracking strongly nonlinear dynamical systems with sparse observations [27, 40, 43]. A promising technique for developing nonlinear filters is to use methods of optimal transportation and coupling of random variables [37, 40, 43], which is not in the scope of the paper.

The ensemble conditional mean filter (EnCMF), introduced by Lei and Bickel as *nonlinear ensemble adjustment filter* in [27], is a natural extension of the EnKF. The EnCMF harnesses the advantage of the conditional expectation in characterizing the underlying conditional distribution and ultimately improves the approximation performance of the updated ensemble for the conditional distribution. Compared with the *ideal* Bayesian filter, the updated random variable (RV) of the conditional mean filter (CMF) does not follow the posterior distribution obtained by applying Bayes' rule on the filter's forecast distribution [14]. However, unlike the EnKF, the CMF's updated mean is identical to the posterior one [33, 48]. Moreover, we show here that, though the covariance of the CMF updated RV does not match the posterior covariance, it coincides with the expected conditional covariance.

For the practical implementation of the EnCMF, knowledge regarding the conditional mean (CM) or an accurate numerical approximation is required. Computing the CM in nonlinear data assimilation problems with high-dimensional state vectors/fields is significantly challenging. In [27], the CM is estimated using a likelihood-based approach via Bayes' rule, which is prone to significant statistical errors for small ensemble sizes due to the intractability of the posterior. Here, we aim to develop a *likelihood-free* nonlinear approximation for the CM using orthogonal projection. Although the idea of using orthogonal projection to approximate CM was briefly introduced in [27], the orthogonal projection-based EnCMF is implemented using only the quadratic regression and shows poor performance compared with the likelihood-based EnCMF (LL-EnCMF) in tracking Lorenz 96 system. However, we show herein that by using a systematical methodology for approximating the CM, including variance reduction and model selection procedures, the proposed machine learning-based ensemble conditional mean filter (ML-EnCMF) outperforms the LL-EnCMF particularly for small ensemble sizes as well as the EnKF.

We use artificial neural networks (ANNs) to approximate the conditional expectation in this work. Notably, we approximate the CM by combining the linear approximation obtained based on the EnKF with an ANN to represent the CM nonlinearity. With this combination, the proposed ML-EnCMF naturally inherits the advantage of the EnKF for data assimilation problems with *closely* linear dynamical models using small ensemble sizes [15].

The forecast ensemble is used as the dataset for training the ANN. Approximating the CM using

small-sized ensembles can lead to overfitting, which causes divergence of the filter. We develop suitable variance reduction and model selection procedures to address this issue. The applied variance reduction method aims to reduce statistical errors when estimating the loss function. The model selection procedure compares the EnKF and ML-EnCMF predictions using an *a priori mean squared error* (MSE) metric to select element-wisely the better one at each updating step. Notably, the performance of the ML-EnCMF is generally better than the EnKF, owing to the model selection procedure. The effectiveness of the ML-EnCMF is demonstrated through detailed numerical studies of tracking the Lorenz-63 and Lorenz-96 systems in their chaotic regime. Especially for the Lorenz-96 system, whose state vector is 40 dimensional, we propose a *localized* ANN structure to avoid overestimated long-range correlation and reduce computational costs for ANN training.

The remainder of this paper is organized as follows. In Sec. 2, we summarize the KF, its ensemble version, and the Bayesian filter. In Sec. 3, the theoretical property of the EnCMF is discussed. In Sec. 4, the ML-EnCMF is presented. In Sec. 5, the performance of the ML-EnCMF implemented for tracking the Lorenz systems is analyzed. Finally, in Sec. 6, conclusions and future research directions are discussed.

2 Filtering method

This section presents the framework of the filtering approach for data assimilation. The detailed theoretical results on filtering techniques can be found in [38, 39, 5, 15, 47, 24].

Considering the discrete-time filtering problem, which presents the evolution of a physical system from time t_{k-1} to t_k by a dynamical state equation:

$$Q_k = \Psi_k(Q_{k-1}), \quad k \in \{1, 2, \dots, \kappa\}, \quad \kappa \in \mathbb{N}, \quad (1)$$

where Q_k is an \mathbb{R}^n -valued RV of the model's states at time t_k , and Ψ_k is the corresponding dynamical operator. In practice, operator Ψ_k can be the (numerical) solution operator of a deterministic dynamical system or a stochastic ordinary differential equation system [34]. The distribution of the RV Q_k is interpreted as an uncertainty model of the actual states, denoted as $q^{\text{tr}}(t_k)$. Observations are recorded at time steps t_k ($k \in \{1, 2, \dots, \kappa\}$) and modeled as

$$Y_k = h(Q_k) + \Xi_k, \quad (2)$$

where Y_k and Ξ_k are the \mathbb{R}^m -valued RVs of the observations and measurement errors, respectively, and $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a known observation map. In practice, the observations can be spatio-temporally sparse. We assume that the measurement error RVs Ξ_k , with $k \in \{0, 1, \dots, \kappa\}$, are statistically independent, the sequence of RVs $Q_0, Q_1, \dots, Q_\kappa$ satisfies the Markov property, and RVs Ξ_k and Q_k are also assumed to be statistically independent. Additionally, the distributions of the RVs Q_k , Ξ_k , and Y_k are assumed absolutely continuous, *i.e.*, their densities exist.

Let $y_k^{\text{obs}} \in \mathbb{R}^m$ denote the observation data observed at time t_k . We denote the set of observation data up to time t_k as $Y_k^{\text{obs}} = \{y_1^{\text{obs}}, y_2^{\text{obs}}, \dots, y_k^{\text{obs}}\}$. An accurate representation of the conditional probability density function (PDF) $\pi_{Q_k|Y_k^{\text{obs}}}$ is not feasible for high-dimensional state dynamical systems. Based on the assumptions that $Q_0, Q_1, \dots, Q_\kappa$ is a Markov process and the measurement error RVs are statistically independent, the conditional PDF $\pi_{Q_k|Y_k^{\text{obs}}}$ can be sequentially approximated [15, 39].

The *ideal* Bayesian filtering procedure from time t_{k-1} to t_k comprises two steps: a *prediction* step, where the conditional PDF $\pi_{Q_{k-1}|Y_{k-1}^{\text{obs}}}$ is transformed to the forecast PDF $\pi_{Q_k|Y_{k-1}^{\text{obs}}}$ using the operator Ψ_k in Eq. (1), and an *update* step, where Bayes' rule is used to map $\pi_{Q_k|Y_{k-1}^{\text{obs}}}$ to $\pi_{Q_k|Y_k^{\text{obs}}}$

as

$$\pi_{Q_k|Y_k^{\text{obs}}}(q) = \frac{\pi_{Q_k|Y_{k-1}^{\text{obs}}}(q) \pi_{\Xi_k}(y_k^{\text{obs}} - h(q))}{\int_{\mathbb{R}^n} \pi_{Q_k|Y_{k-1}^{\text{obs}}}(q') \pi_{\Xi_k}(y_k^{\text{obs}} - h(q')) dq'}, \quad (3)$$

where π_{Ξ_k} is the PDF of RV Ξ_k .

Ensemble-based filters represent the uncertainty of the states via ensembles. At the prediction step, the *forecast ensembles* of the states and observations at time t_k , denoted as $\{q_k^{f(i)}\}_{i=1}^N$ and $\{y_k^{f(i)}\}_{i=1}^N$, are computed respectively as

$$\begin{aligned} q_k^{f(i)} &= \Psi_k(q_{k-1}^{a(i)}), \quad i = 1, 2, \dots, N, \\ y_k^{f(i)} &= h(q_k^{f(i)}) + \xi_k^{(i)}, \quad i = 1, 2, \dots, N, \end{aligned} \quad (4)$$

where N is the ensemble size, $\{q_{k-1}^{a(i)}\}_{i=1}^N$ is the *updated ensemble* obtained at time t_{k-1} , and $\xi_k^{(i)}$ denotes the independent and identically distributed (i.i.d.) samples of Ξ_k . Here, common meteorological notations is used [40], *i.e.*, the superscript f (forecast) to denote the prior statistics, and the superscript a (analysis) to denote the posterior statistics. Generally, the updated ensemble $\{q_k^{a(i)}\}_{i=1}^N$ are computed using a suitable map \mathcal{T}_k as

$$q_k^{a(i)} = \mathcal{T}_k(q_k^{f(i)}, y_k^{f(i)}, y_k^{\text{obs}}), \quad i = 1, 2, \dots, N. \quad (5)$$

For example, the EnKF's updating map \mathcal{T} is linear. A summary of the EnKF is given in Appendix A.

The members of the ensembles $\{q_k^{f(i)}\}_{i=1}^N$, $\{y_k^{f(i)}\}_{i=1}^N$, and $\{q_k^{a(i)}\}_{i=1}^N$ can be seen as samples of RVs Q_k^f , Y_k^f , and Q_k^a , respectively, given as

$$Q_k^f = \Psi_k(Q_{k-1}^a), \quad k \in \{1, 2, \dots, \kappa\}, \quad (6a)$$

$$Y_k^f = h(Q_k^f) + \Xi_k, \quad k \in \{1, 2, \dots, \kappa\}, \quad (6b)$$

$$Q_k^a = \mathcal{T}_k(Q_k^f, Y_k^f, y_k^{\text{obs}}), \quad k \in \{1, 2, \dots, \kappa\}, \quad (6c)$$

where $Q_0^a \equiv Q_0$.

Ideally, map \mathcal{T}_k is identified such that $\pi_{Q_k^a}$ is an accurate approximation of the Bayesian posterior density $\pi_{Q_k^f|Y_k^{\text{obs}}}$, given as

$$\pi_{Q_k^f|Y_k^{\text{obs}}}(q) = \frac{\pi_{Q_k^f}(q) \pi_{\Xi_k}(y_k^{\text{obs}} - h(q))}{\int_{\mathbb{R}^n} \pi_{Q_k^f}(q') \pi_{\Xi_k}(y_k^{\text{obs}} - h(q')) dq'}, \quad (7)$$

which obtained from Eq. (3) by substituting PDF $\pi_{Q_k|Y_{k-1}^{\text{obs}}}$ with its approximation, *i.e.*, the forecast PDF $\pi_{Q_k^f}$. Indeed, when $\pi_{Q_k^a}$ is identical to $\pi_{Q_k^f|Y_k^{\text{obs}}}$ for every $k \in \{1, 2, \dots, \kappa\}$, we have $\pi_{Q_k^f} \equiv \pi_{Q_k|Y_{k-1}^{\text{obs}}}$ and $\pi_{Q_k^a} \equiv \pi_{Q_k|Y_k^{\text{obs}}} \forall k \in \{1, 2, \dots, \kappa\}$.

Because the EnKF's updating map \mathcal{T}_k is linear, for the nonlinear setting, *i.e.*, maps Ψ_k are nonlinear, there are biased errors between the PDFs $\pi_{Q_k^a}$ and $\pi_{Q_k^f|Y_k^{\text{obs}}}$. Consequently, the mean and higher statistical moments of PDF $\pi_{Q_k^a}$ do not match those of PDF $\pi_{Q_k^f|Y_k^{\text{obs}}}$ [25].

3 Ensemble conditional mean filter

In this section, we analyze the conditional mean filter (CMF) and its ensemble version (EnCMF). The orthogonal projection property of conditional expectation, which is the base of the ML-EnCMF, will be discussed in detail. First, in Sec. 3.1, we recall some essential properties of

the conditional expectation tailored for the update step of the CMF and discuss two particular cases, namely, the CM and the conditional variance. In Sec. 3.2, the CMF and its properties are analyzed. In Sec. 3.3, the ensemble version of CMF is presented. Finally, in Sec. 3.4, we exemplify the EnCMF using a simple static inverse problem that highlights its differences from the EnKF.

3.1 Conditional mean and variance

Because our focus is on the *update* step of the filtering setting, we ignore subscript k from here onward whenever possible to simplify the notations. We assume that the variances of the RVs Q^f and Y^f are finite. Let σ_{Y^f} be the σ -algebra generated by the observation RV Y^f and $r : \mathbb{R}^n \rightarrow \mathbb{R}$ be an arbitrary function such that $r \circ Q^f$ — a RV composed of the forecast state RV Q^f and the function r — has a finite variance. The conditional expectation $\mathbb{E}[r \circ Q^f | Y^f]$ is a σ_{Y^f} -measurable RV defined as follows:

$$\int_{\mathcal{B}} \mathbb{E}[r \circ Q^f | Y^f](\omega) \mathbb{P}(d\omega) = \int_{\mathcal{B}} r \circ Q^f(\omega) \mathbb{P}(d\omega), \quad \forall \mathcal{B} \in \sigma_{Y^f}, \quad (8)$$

where Ω and \mathbb{P} are an underlying probability space and measure, respectively. The general theoretical properties of the conditional expectation can be found in, *e.g.*, [7] and [13, Chapter 4]. According to the Doob-Dynkin lemma, the conditional expectation $\mathbb{E}[r \circ Q^f | Y^f]$, as a σ_{Y^f} -measurable function, takes the form $\phi_{r \circ Q^f}(Y^f)$ for some almost surely unique measurable function $\phi_{r \circ Q^f}$ [7]. The conditional expectation has a geometric interpretation as the L_2 projection of the RV $r \circ Q^f$ onto the σ -algebra generated by the observation RV Y^f . For developing our filter, we use two particular cases of the conditional expectation: the CM as the main ingredient of the filter and the conditional variance for analyzing the variance of the updated RVs.

Conditional mean

The CM $\mathbb{E}[Q^f | Y^f]$ is the vector-valued RV $\mathbb{E}[Q^f | Y^f] := \phi_{Q^f} \circ Y^f$ for some *unique* function $\phi_{Q^f} : \mathbb{R}^m \rightarrow \mathbb{R}^n$. The value of map ϕ_{Q^f} at y is identical to the mean of the conditional PDF $\pi_{Q^f | Y^f}(\cdot | y)$:

$$\phi_{Q^f}(y) = \int q \pi_{Q^f | Y^f}(q | y) dq. \quad (9)$$

Particularly, $\phi_{Q^f}(y^{\text{obs}})$, which is obtained by evaluating map ϕ_{Q^f} with the observation data y^{obs} , is the mean of the posterior PDF $\pi_{Q^f | y^{\text{obs}}}$. Usually, the CM does not have an analytical solution and is approximated using the likelihood-based (see Appendix B) or the orthogonal projection-based approach in the filtering context. Our paper focuses on the orthogonal projection-based approach.

Using the orthogonal projection property, the CM can be identified as follows

$$\phi_{Q^f} = \arg \min_{g \in \mathcal{S}(\mathbb{R}^m, \mathbb{R}^n)} \mathbb{E}[\|Q^f - g \circ Y^f\|^2], \quad (10)$$

where $\mathcal{S}(\mathbb{R}^m, \mathbb{R}^n)$ is the set of all functions $g : \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that the variance of the RV $g(Y^f)$ is finite, and $\|\cdot\|$ denotes the usual Euclidean norm. When limiting the function g in Eq. (10) to be linear, the suboptimal approximation of the CM has a closed-form, which is given in the following lemma.

Lemma 1 (Linear approximation of the CM). *The linear approximation g_1 of the map ϕ_{Q^f} is defined as the orthogonal projection of RV Q^f onto the sub- σ -algebra $\sigma_{Y^f}^* = \{g_1^* \circ Y^f\}$ for all linear functions $g_1^* : \mathbb{R}^m \rightarrow \mathbb{R}^n$. The map g_1 can be analytically obtained as*

$$g_1(y) = \mathbf{K}y + b, \quad (11)$$

where \mathbf{K} is the generalized Kalman gain,

$$\mathbf{K} = \text{Cov} [Q^f, Y^f] \text{Cov} [Y^f]^{-1}, \quad (12)$$

and $b = \mathbb{E} [Q^f - \mathbf{K}Y^f]$.

A proof of Lemma 1 can be found in [33]. Notice that, the generalized Kalman gain \mathbf{K} defined in Eq. (12) is identical to its basic version \mathbf{K}^l (Eq. (48b)) for linear observation maps, see Eq. (2).

Conditional variance

The conditional covariance matrix $\text{Cov} [Q^f | Y^f]$ is an $n \times n$ -matrix valued RV defined as follows

$$\text{Cov} [Q^f | Y^f] \equiv \mathbb{E} [Q^f Q^{f\top} | Y^f] - \mathbb{E} [Q^f | Y^f] \mathbb{E} [Q^f | Y^f]^\top. \quad (13)$$

3.2 Conditional mean filter

The mean of the posterior PDF $\pi_{Q^f | y^{\text{obs}}}$ coincides with the CM evaluated at y^{obs} , $\phi_{Q^f}(y^{\text{obs}})$; hence, it is desirable to design a filter such that its updated RV also has the mean value $\phi_{Q^f}(y^{\text{obs}})$. This filter should also agree with the KF in the case of the linear-Gaussian setting. The CMF proposed by Lei and Bickel in [27] fulfills these requirements. The CMF formulates the updated RV as follows

$$Q^a = Q^f + \phi_{Q^f}(y^{\text{obs}}) - \phi_{Q^f}(Y^f). \quad (14)$$

The transformation in Eq. (14) is a nonlinear version of the general form \mathcal{T} in Eq. (6c). Moreover, we show here that, though the covariance of the CMF updated RV does not match the posterior covariance, it coincides with the expected conditional covariance. These properties of the CMF are formulated in the following theorem.

Theorem 1. *The updated RV obtained using the CMF in Eq. (14) satisfies the following properties:*

A) *For the linear-Gaussian setting, i.e., Q^f and Ξ are Gaussian RVs, and maps Ψ and h are linear, the CMF coincides with the KF, such that*

$$Q^f + \phi_{Q^f}(y^{\text{obs}}) - \phi_{Q^f}(Y^f) \equiv Q^f + \mathbf{K}^l(y^{\text{obs}} - Y^f). \quad (15)$$

B) *The mean of the RV Q^a expressed in Eq. (14) is identical to that of the posterior PDF $\pi_{Q^f | y^{\text{obs}}}$ (Eq. (7)):*

$$\mathbb{E} [Q^a] = \phi_{Q^f}(y^{\text{obs}}), \quad (16)$$

C) *The covariance of the RV Q^a expressed in Eq. (14) is equal to the expected conditional covariance:*

$$\text{Cov} [Q^a] = \mathbb{E} [\text{Cov} [Q^f | Y^f]]. \quad (17)$$

Proof. A) For the linear-Gaussian setting, we achieve that $\phi_{Q^f}(Y^f) \equiv \mathbf{K}^l Y^f + b$ using Lemma 1. The updated RV Q^a obtained using Eq. (14) becomes $Q^a = Q^f + \mathbf{K}^l(y^{\text{obs}} - Y^f)$, which is identical to the KF (Eq. (48a)).

B) Let Q^I be defined as $Q^I := Q^f - \mathbb{E} [Q^f | Y^f]$. Based on the law of total expectation, we obtain the mean and variance of Q^I as follows:

$$\mathbb{E} [Q^I] = 0_n, \quad (18)$$

where 0_n is the n -dimensional zeros vector, and

$$\text{Cov} [Q^I] = \text{Cov} [Q^f] - \text{Cov} [\mathbb{E} [Q^f | Y^f]], \quad (19)$$

respectively. Thus, the mean of the updated RV Q^a is identical to $\phi_{Q^f}(y^{\text{obs}})$:

$$\mathbb{E}[Q^a] = \mathbb{E}[Q^f] + \phi_{Q^f}(y^{\text{obs}}) - \phi_{Q^f}(y^{\text{obs}}), \quad (20)$$

which is the mean of the posterior PDF as expressed in Eq. (9).

C) Using the law of total variance

$$\text{Cov}[Q^f] = \mathbb{E}[\text{Cov}[Q^f|Y^f]] + \text{Cov}[\mathbb{E}[Q^f|Y^f]], \quad (21)$$

we obtain

$$\text{Cov}[Q^a] = \text{Cov}[Q^f] - \text{Cov}[\mathbb{E}[Q^f|Y^f]] = \mathbb{E}[\text{Cov}[Q^f|Y^f]]. \quad (22)$$

□

Property A of Theorem 1 indicates that the CMF is a natural extension of the KF. Compared with the posterior PDF, $\pi_{Q^f|y^{\text{obs}}}$, the updated RV Q^a of the CMF expressed in Eq. (14) shows an identical mean vector (Property B). However, the distribution of the centered updated RV, $Q^a - \phi_{Q^f}(Y^f)$, and the covariance matrix $\text{Cov}[Q^a]$, in particular, are independent of the observational data y^{obs} and not identical to those of the posterior PDF, besides special cases such as the linear-Gaussian setting. An approach to improve the CMF is to match the covariance and possibly higher moments of the updated RV with the empirical posterior moments, as proposed in [27, 33], which requires significant additional computational resources. In an extreme scenario where the conditional mean does not encode any information, *i.e.*, the conditional mean has zero variance, the CMF exhibits no information gain, *i.e.*, $Q^a \equiv Q^f$. This limitation also applies to the EnKF.

Remark 1. *Usually, the conditional variance and its range are significantly diminished compared with the variance of the forecast RV Q^f in practice because uncertainty is reduced thanks to observational data. Hence, the error between the expected conditional covariance and the posterior covariance ($\text{Cov}[Q^f|Y^f = y^{\text{obs}}]$) is much smaller than the forecast one ($\text{Cov}[Q^f]$). In other words, compared with the posterior PDF $\pi_{Q^f|y^{\text{obs}}}$, the updated RV Q^a has an approximate covariance matrix (Theorem 1, Property C). We illustrate this argument in Sec. 3.4. Moreover, Property C is advantageous when combining EnCMF with the method of A-optimal design of experiments [3, 12], which seeks for the experiment setup to minimize the trace of the expected conditional covariance.*

3.3 Ensemble approximation of conditional mean filter

Similar to the EnKF, the EnCMF uses the ensemble technique to represent the forecast and updated RVs. In the forecast step, the distributions of the RVs Q^f and Y^f are approximated using the N -sized ensembles, $\{q^{f(i)}\}_{i=1}^N$ and $\{y^{f(i)}\}_{i=1}^N$, respectively, as expressed in Eq. (6). In the update step, the updated ensemble $\{q^{a(i)}\}_{i=1}^N$ of the RV Q^a is evaluated using Eq. (14):

$$q^{a(i)} = q^{f(i)} + \phi_{Q^f}(y^{\text{obs}}) - \phi_{Q^f}(q^{f(i)}), \quad i = 1, 2, \dots, N. \quad (23)$$

Implementing the EnCMF requires the evaluation of map ϕ_{Q^f} of the CM. This task can be performed using two approaches: a) by evaluating the conditional PDF in Eq. (9) via the likelihood function or b) using the orthogonal projection property expressed in Eq. (10). The likelihood-based approach, which uses the sampling technique to estimate values of the conditional mean for different samples of measurement RV Y^f (summarized in Appendix B), has limited applicability for data assimilation problems with high-dimensional state-spaces because the conditional distribution becomes intractable. The orthogonal projection-based approach is favorable in such a situation. In Sec. 3.4, we implement the EnCMF with a simple one-dimensional example, in which the likelihood-based approach is applied. In Sec. 4, we discuss the approximation of the CM using the

orthogonal projection-based approach combined with machine learning (ML), which will lead to the ML-EnCMF method.

Algorithm 1 presents a pseudocode of the ML-EnCMF. The forecast step is similar to those of the EnKF. The update step using ML-based approximation of the CM will be presented in Sec. 4. In terms of execution schedule, the prediction step and the approximation of the CM do not require observational data; hence they can be performed prior to the arrival of the data for accelerating the filtering process.

Algorithm 1 ML-EnCMF algorithm for performing assimilation at time t_k .

Require: Updated ensembles obtained from the assimilation step performed at t_{k-1} $\{q_{k-1}^{a(i)}\}_{i=1}^N$ and observational data y_k^{obs}

Prediction:

1: Evaluate forecast ensembles $\{q_k^{f(i)}\}_{i=1}^N$ ▷ Eq. (4)

Update:

2: Approximate the map $\phi_{Q_k^f}$ (explained in Sec. 4) ▷ Algorithm 2

3: Evaluate the updated ensemble $\{q_k^{a(i)}\}_{i=1}^N$ ▷ Eq. (27b)

return Updated ensemble $\{q_k^{a(i)}\}_{i=1}^N$

3.4 Illustration of the EnCMF for a simple static inverse problem

We consider the following inverse problem to compare the EnCMF with the EnKF and the ideal Bayesian filter. Given a one-dimensional RV $Q^f \sim \mathcal{N}(0, 2^2)$ and the following nonlinear observation map,

$$Y^f = h(Q^f) + \Xi \quad \text{where} \quad h(q) = \begin{cases} q & \text{if } q \leq 0 \\ q^2 & \text{if } q > 0 \end{cases}, \quad \Xi \sim \mathcal{N}(0, 0.5^2), \quad (24)$$

the task is to evaluate the Bayesian posterior for different values of the observation y^{obs} . We solve the problem using the EnCMF and EnKF and compare their updated ensembles with the Bayesian posterior in terms of three statistical characteristics: mean, variance, and PDF. The observation map considered here is nonlinear; therefore, we employ Eq. (12) for computing Kalman gain. For this simple problem, we use the likelihood-based approach, explained in Appendix B, for approximating the conditional mean (Eqs. (51, 53)) and implementing the EnCMF (Eq. (54)). A large ensemble size of 10000 is selected to minimize statistical errors.

We refer to Eq. (12) and apply the MC method with the same ensemble size to approximate the Kalman gain. For comparison, Fig. 1 depicts the CM map ϕ_{Q^f} with its linear approximation, obtained using Eq. (11). We observe a significant error in the linear approximation. Consequently, a significant bias between the means of the updated ensemble of the EnKF and the Bayesian posterior is predicted.

We approximate the conditional expectation of the second moment $\mathbb{E}[(Q^f)^2 | Y^f]$ and the

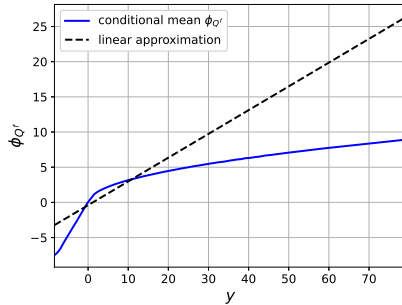


Figure 1: CM computed using Eq. (9) and its linear approximation (Eq. (11)).

conditional variance $\text{Var} [Q^f | Y^f]$ using the conditional PDF as follows

$$\mathbb{E} \left[(Q^f)^2 | Y^f = y \right] = \int q^2 \pi_{Q^f | Y^f}(q|y) dq, \quad (25a)$$

$$\begin{aligned} \text{Var} [Q^f | Y^f] &= \mathbb{E} \left[(Q^f)^2 | Y^f \right] - (\mathbb{E} [Q^f | Y^f])^2 \\ &= \mathbb{E} \left[(Q^f)^2 | Y^f \right] - (\phi_{Q^f}(Y^f))^2, \end{aligned} \quad (25b)$$

where the likelihood-based approach is applied to compute $\mathbb{E} \left[(Q^f)^2 | Y^f = y \right]$, see Eq. (52). Fig. (2) presents the empirical PDF of the conditional variance.

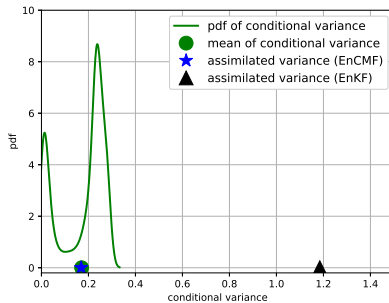


Figure 2: Empirical PDF of the conditional variance $\text{Var} [Q^f | Y^f]$ expressed in Eq. (25b) compared with the variances of the updated ensembles obtained using the EnKF and EnCMF. The expected conditional variance, $\mathbb{E} [\text{Var} [Q^f | Y^f]]$, and the updated ensemble variance of the EnCMF show closely identical estimated values of approximately 0.17.

We implement the EnKF, using Eqs. (4, 12, 50), and the LL-EnCMF, using Eqs. (4, 53, 54). The updated ensemble variances in both filters are invariant when varying measurement data. For the EnCMF, the updated ensemble variance is a non-biased estimator of the mean of the conditional variance; thus, the EnCMF shows significant improvement in estimating the conditional variance than the EnKF. Notably, the absolute error between the EnCMF's updated variance and the conditional variance (within $[0, 0.18)$) is much smaller than the prior variance (equal to 4). This observation is an illustration for the Remark 1.

We consider different observation scenarios using various q^{tr} values to evaluate the *synthesized* observation data as $y^{\text{obs}} = h(q^{\text{tr}}) + \xi$, where ξ is an i.i.d. sample of the error RV Ξ . Fig. 3 presents the empirical densities of updated ensembles and the posterior PDF. Although the empirical PDF of the updated ensemble obtained using the EnCMF does not coincide with the Bayesian posterior, it still fits the posterior significantly better than the EnKF. In particular, the updated ensemble mean obtained using the EnCMF is closely identical to the Bayesian posterior mean and in good agreement with the truth value q^{tr} . In contrast, the updated ensemble of the EnKF exhibits significant biased errors in terms of the mean value compared with the Bayesian posterior. This numerical experiment reconfirms the theoretical statement in Theorem 1.

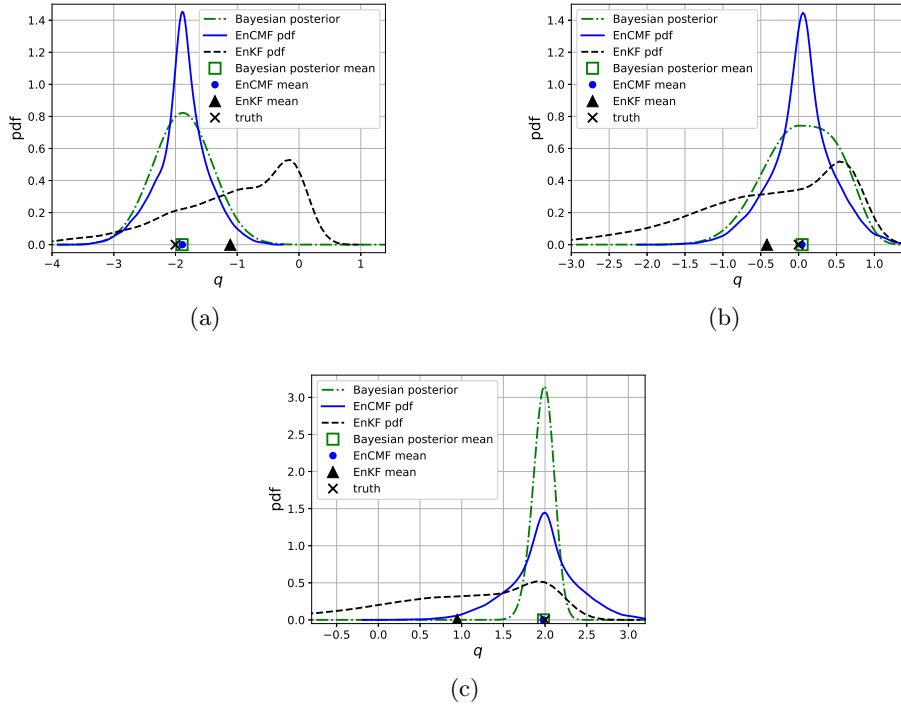


Figure 3: Comparison of the empirical densities of the updated ensembles and the Bayesian posterior: (a) actual value $q^{\text{tr}} = -2$ (minus standard deviation of the prior distribution), (b) actual value $q^{\text{tr}} = 0$ (mean of the prior distribution), (c) actual value $q^{\text{tr}} = +2$ (plus standard deviation of the prior distribution).

A counter-example to illustrate the limitation of the CMF, discussed in Sec. 3.2, is to replace the observation map in Eq. (24) by $h = q^2$ [14]. In this case, the CMF fails to incorporate measurement data because the conditional mean is deterministic $\phi_{Q^f}(Y^f) = 0$. Indeed the updated RV of the CMF is identical to the forecast.

4 ML-based EnCMF

This section presents the ML-EnCMF. While our paper focuses on the update step, improving the efficiency of the forecast step using ML techniques is also actively investigated, *e.g.*, in [1, 8, 9, 16].

Based on the orthogonal projection property (Eq. (10)), we develop an approximation of the CM using ANN. The updated ensemble is obtained by inserting the CM approximation in Eq. (23).

The section is organized as follows. In Sec. 4.1, we discuss the approximation of the CM using a combination of the KF linear updating map and an ANN. In Sec. 4.2 and Sec. 4.3, the variance reduction and the model selection methods are respectively reported. Finally, Sec. 4.4 reports a pseudo-algorithm for the training and model selection procedures.

4.1 ANN-based approximation of the conditional mean

Let $g_{\text{NN}}(\cdot; \boldsymbol{\theta}) : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be an ANN map, where $\boldsymbol{\theta}$ denotes its hyper-parameters, *i.e.*, the network's weights and biases. The ANN structure is not fixed and adapted following the natural representation of the states. For example, convolution neural networks can be used for spatial/-field states to represent the spatial correlation, as illustrated in Sec. 5.3. Here, we approximate the CM, *i.e.*, map ϕ_{Q^f} , by combining the linear approximation g_1 (Eq. (11)) and the ANN $g_{\text{NN}}(\cdot; \boldsymbol{\theta})$ as

$$\phi_{Q^f}(\cdot) \approx g_1(\cdot) + a \odot g_{\text{NN}}(\cdot; \boldsymbol{\theta}), \quad (26)$$

where $a = [a_1, \dots, a_n]$ is a n -dimensional vector with $a_\alpha \in \{0, 1\}$ with $\alpha = 1, \dots, n$, and \odot denotes the element-wise multiplication. In other words, the α th-component of the state vector is updated using either the combined nonlinear map ($a_\alpha = 1$) or the Kalman linear map ($a_\alpha = 0$). The value of vector a is identified based on the model selection described in Sec. 4.3 below.

Instead of relying solely on the ANN to approximate the map ϕ_{Q^f} , the combination in Eq. (26) is more robust. Indeed, the proposed combination inherits from the acceptable performance of the EnKF in closely linear-Gaussian settings with small ensemble sizes. Performance-enhancing techniques successfully applied for the EnKF with small-sized ensembles, *e.g.*, covariance tapering [17], are also applicable for the linear part of the ML-EnCMF. Moreover, the ML-EnCMF generally performs better than the EnKF because the filter is only activated when it approximates the CM better than the EnKF, owing to the model selection procedure.

After the ANN g_{NN} is trained, given the observation data y^{obs} , the updated RV Q^a of the CMF (Eq. (14)) and its ensemble (Eq. (23)) are approximated by the following relations

$$Q^a = Q^f + \mathbf{K}(y^{\text{obs}} - Y^f) + a \odot (g_{\text{NN}}(y^{\text{obs}}; \boldsymbol{\theta}) - g_{\text{NN}}(Y^f; \boldsymbol{\theta})), \quad (27a)$$

$$q^{a(i)} = q^{f(i)} + \mathbf{K}(y^{\text{obs}} - y^{f(i)}) + a \odot (g_{\text{NN}}(y^{\text{obs}}; \boldsymbol{\theta}) - g_{\text{NN}}(y^{f(i)}; \boldsymbol{\theta})), \quad (27b)$$

where $i = 1, \dots, N$, respectively. The mean and covariance estimators of the updated ensemble are given as

$$\mathbb{E}[Q^a] \approx g_1(y^{\text{obs}}) + a \odot g_{\text{NN}}(y^{\text{obs}}; \boldsymbol{\theta}), \quad (28a)$$

$$\begin{aligned} \text{Cov}[Q^a] \approx & \frac{1}{N} \sum_{i=1}^N [q^{f(i)} - g_1(y^{f(i)}) - a \odot g_{\text{NN}}(y^{f(i)}; \boldsymbol{\theta})] \\ & [q^{f(i)} - g_1(y^{f(i)}) - a \odot g_{\text{NN}}(y^{f(i)}; \boldsymbol{\theta})]^\top, \end{aligned} \quad (28b)$$

using Theorem 1.

We train the map g_{NN} by using the orthogonal projection property of the CM stated in Eq. (10). For a given tensor of hyper-parameters $\boldsymbol{\theta}$, we use the MSE $J(\boldsymbol{\theta})$ defined as

$$J(\boldsymbol{\theta}) = \mathbb{E} \left[\|Q^f - g_1(Y^f) - g_{\text{NN}}(Y^f; \boldsymbol{\theta})\|^2 \right] \quad (29)$$

as a metric for the approximation in Eq. (26). The hyper-parameters of the ANN are obtained by solving the following optimization problem

$$\boldsymbol{\theta} = \arg \min_{\boldsymbol{\theta}_*} J(\boldsymbol{\theta}_*) + \mathcal{R}(\boldsymbol{\theta}_*), \quad (30)$$

where $\mathcal{R}(\boldsymbol{\theta}_*)$ is a regularization term applied to reduce overfitting, usually composed of L_1 or L_2 norms of the weights [10].

Remark 2. *An important metric to assess the performance of a filter is the L_2 norm of the error between the updated ensemble mean and the ground-truth state q^{tr} . For example, that metric for the ML-EnCMF is evaluated as $\|q^{\text{tr}} - g_1(y_{\text{obs}}) - g_{\text{NN}}(y_{\text{obs}}; \boldsymbol{\theta})\|^2$ supposing vector a is a unit one. Therefore, the value of $J(\boldsymbol{\theta})$ is the apriori MSE of the updated ensemble mean. Furthermore, $J(\boldsymbol{\theta})$ is equal to the total variance of the updated ensemble (see Eq. (28b)), which is a measure of the prediction uncertainty. In other words, minimizing the function J is equivalent to minimizing the apriori MSE of the updated ensemble mean and simultaneously minimizing the prediction uncertainty.*

4.2 Reduced variance estimator of the MSE metric

In our approach, the metric $J(\boldsymbol{\theta})$ is estimated using the forecast ensembles. Let $\mathcal{D} = \{(y^{f(i)}, q^{f(i)})\}_{i=1}^N$ be the dataset comprising pairs $(y^{f(i)}, q^{f(i)})$ collected from the forecast ensembles, evaluated using Eq. (4) in the prediction step. The dataset \mathcal{D} is divided into two sets: an N_T -sized training dataset \mathcal{D}_T for tuning the hyper-parameters and an N_S -sized test dataset \mathcal{D}_S for testing the ANN; here, $N_T + N_S = N$. In both training and testing processes, the metric $J(\boldsymbol{\theta})$ is estimated based on the corresponding datasets, \mathcal{D}_T and \mathcal{D}_S , respectively. Statistical errors are inherent in such an estimation; hence, we present a variance reduction technique to reduce these errors.

The crude MC estimator of the metric J using the training dataset is obtained as follows

$$\widehat{J}(\boldsymbol{\theta}|\mathcal{D}_T) = \frac{1}{N_T} \sum_{(y^{f(i)}, q^{f(i)}) \in \mathcal{D}_T} \|q^{f(i)} - g_1(y^{f(i)}) - g_{\text{NN}}(y^{f(i)}; \boldsymbol{\theta})\|^2, \quad (31)$$

where $y^{f(i)} = h(q^{f(i)}) + \xi^{(i)}$ as stated in Eq. (4). We use the following reduced-variance estimator obtained by suitably augmenting the data to reduce the statistical errors in evaluating the MSE metric

$$\widehat{J}^{\text{vr}}(\boldsymbol{\theta}|\mathcal{D}_T) = \frac{1}{N_T} \sum_{(y^{f(i)}, q^{f(i)}) \in \mathcal{D}_T} \frac{1}{M} \sum_{j=1}^M \|q^{f(i)} - g_1(h(q^{f(i)}) + \xi^{(i,j)}) - g_{\text{NN}}(h(q^{f(i)}) + \xi^{(i,j)}; \boldsymbol{\theta})\|^2, \quad (32)$$

where M is the data augmentation multiplier, and $\xi^{(i,j)}$ with $i = 1, 2, \dots, N$ and $j = 1, 2, \dots, M$ are i.i.d. samples of RV Ξ . When M is sufficiently large, the estimator $\widehat{J}^{\text{vr}}(\boldsymbol{\theta}|\mathcal{D}_T)$ shows minor statistical errors compared with $\widehat{J}(\boldsymbol{\theta}|\mathcal{D}_T)$ (see discussion in Appendix C).

Intuitively, this variance reduction technique decreases the sensitivity of the trained ANN to observation noise because additional noisy data are used for training. From the implementation perspective, using the reduced-variance estimator $\widehat{J}^{\text{vr}}(\boldsymbol{\theta}|\mathcal{D}_T)$ is equivalent to training the ANN on an $N \times M$ augmented dataset:

$$\mathcal{D}^a = \{(y^{f(i,j)}, q^{f(i)}), \quad i = 1, \dots, N, \quad j = 1, \dots, M\}, \quad (33)$$

where $y^{f(i,j)} = h(q^{f(i)}) + \xi^{(i,j)}$, which does not require any modification of the implemented training algorithm.

4.3 Model selection

Training the ANN with possibly many hyper-parameters on small datasets can lead to overfitting. Although the CM may be nonlinear, the linear approximation (Eq. (11)) may still be preferable over an ANN-based approximation that is significantly overfitted owing to robustness. Our proposed model selection compares the ML-EnCMF with the EnKF using the apriori MSE of the updated ensemble mean.

Let L_α and J_α , $\alpha = 1, \dots, n$, be the apriori MSE of the EnKF and ML-EnCMF ensemble means, respectively, of the α -th state component, defined as

$$L_\alpha = \mathbb{E} \left[\left([Q^f]_\alpha - [g_1(Y^f)]_\alpha \right)^2 \right], \quad (34a)$$

$$J_\alpha = \mathbb{E} \left[\left([Q^f]_\alpha - [g_1(Y^f)]_\alpha - [g_{\text{NN}}(Y^f; \boldsymbol{\theta})]_\alpha \right)^2 \right]. \quad (34b)$$

We identify the vector a in Eq. (26) as

$$a_\alpha = \mathbf{1}(\widehat{L}_\alpha > \widehat{J}_\alpha), \quad \alpha = 1, \dots, n \quad (35)$$

where $\mathbf{1}(\cdot)$ is a logical operator yielding one if the condition (\cdot) is valid and zero otherwise. In Eq. (35), \widehat{L}_α and \widehat{J}_α are estimated, respectively, as:

$$\widehat{L}_\alpha = \frac{1}{N_S} \sum_{(y^{f(i)}, q^{f(i)}) \in \mathcal{D}_S} \frac{1}{M} \sum_{j=1}^M \left([q^{f(i)}]_\alpha - [g_1(h(q^{f(i)} + \xi^{(i,j)}))]_\alpha \right)^2, \quad (36a)$$

$$\widehat{J}_\alpha = \frac{1}{N_S} \sum_{(y^{f(i)}, q^{f(i)}) \in \mathcal{D}_S} \frac{1}{M} \sum_{j=1}^M \left([q^{f(i)}]_\alpha - [g_1(h(q^{f(i)} + \xi^{(i,j)}))]_\alpha - [g_{\text{NN}}(h(q^{f(i)} + \xi^{(i,j)}); \boldsymbol{\theta})]_\alpha \right)^2, \quad (36b)$$

where we apply the variance reduction technique, similar to Eq. (32), and use the test dataset to avoid bias errors.

4.4 Algorithm

Given the forecast ensembles $\{q_k^{f(i)}\}_{i=1}^N$, $\{y_k^{f(i)}\}_{i=1}^N$, we estimate the Kalman gain, using Eq. (48b) (Eq. (12)) for linear (nonlinear) observational maps, accordingly. We then approximate the map $\phi_{Q_k^f}$ as stated in Eq. (26). The training pseudocode is presented in Algorithm 2. Using the forecast ensemble, $\mathcal{D} = \{(y_k^{f(i)}, q_k^{f(i)})\}_{i=1}^N$, we first generate the augmented dataset (Eq. (33)), which is used to compute the reduced-variance estimators of the MSE metric as expressed in Eqs. (32, 36). Then, to solve the MSE minimizing problem described in Eq. (30), this algorithm uses the mini-batch version of the stochastic gradient descent method, a widely accepted approach for ANN training [23]. We include a call-back procedure into the training process to select the hyper-parameters with the best performance [18]. Finally, the model selection procedure (explained in Sec. 4.3) is performed. Algorithms 1 and 2, together, summarize our ML-EnCMF implementation.

Algorithm 2 Training algorithm of the ANNs used for approximating the map ϕ_{Qf} .

Require: Ensembles $\{q^{f(i)}\}_{i=1}^N$ and $\{y^{f(i)}\}_{i=1}^N$, number of epochs n_e , ANN g_{NN} , Kalman gain

Initialization:

- 1: Set the initial hyper-parameters θ^0
- 2: Generate $N \times M$ -size augmented dataset ▷ Eq. (33)
- 3: Compute the testing metric $m_0 = \widehat{J}^{\text{tr}}(\theta^0 | \mathcal{D}_S)$
▷ similar to Eq. (32) with test dataset

Training:

- 4: **for** $\nu = 1, 2, \dots, n_e$ **do**
 - 5: Train ANN using mini-batch stochastic gradient descent method [23, 4] to update θ^ν
 - 6: Compute the testing metric: $m_\nu = \widehat{J}^{\text{tr}}(\theta^\nu | \mathcal{D}_S)$
▷ similar to Eq. (32) with test dataset
 - 7: **if** $m_\nu < \min\{m_0, m_1, \dots, m_{\nu-1}\}$ **then**
 - 8: $\theta = \theta^\nu$ ▷ Call-back procedure
 - 9: **end for**
 - 10: **Model selection** ▷ Eq. (35)
- return** Hyper-parameters θ and vector a
-

5 Numerical experiments: results and discussion

In this section, we demonstrate the performance of ML-EnCMF for the data assimilation of Lorenz 63 (L63) and Lorenz 96 (L96) systems in comparison with the EnKF and LL-EnCMF. Moreover the challenge in applying the ML-EnCMF for large-scale data assimilation applications is also discussed. This section is organized as follows. In Sec. 5.1, we first describe the general setup of the experiments. Then, in Sec. 5.2 and 5.3, we discuss the numerical results for tracking L63 and L96 systems, respectively. Finally, in the Sec. 5.4 we discuss the computational cost of the ML-EnCMF and the challenges for large-scale data assimilation applications.

5.1 Setup

This subsection describes the setup of the investigated data assimilation problems, particularly dynamical models, performance metrics, and a numerical experiment for assessing errors in approximating the CM.

5.1.1 Dynamical models

Lorenz-63 system

The L63 model is a simplified model of atmospheric convection [29] comprising three ODEs:

$$\frac{dq_1}{dt} = \sigma(q_2 - q_1), \quad \frac{dq_2}{dt} = q_1(\rho - q_3) - q_2, \quad \frac{dq_3}{dt} = q_1q_2 - \beta q_3, \quad (37)$$

where q_1 , q_2 , and q_3 are proportional to the rate of convection, horizontal temperature variation, and vertical temperature variation, respectively, and σ , β , and ρ are the system parameters. We use $\sigma = 10$, $\beta = 8/3$, and $\rho = 28$ as a conventional setting of the L63 system, resulting in a chaotic behavior.

Lorenz-96 system

The L96 model is an idealized model of a one-dimensional latitude band of the Earth's atmosphere [30]. The system is defined using a set of ODEs over the periodic domain of a 40-dimensional state vector, $\mathbf{q} = [q_1, q_2, \dots, q_n]^\top$ with $n = 40$:

$$\frac{dq_\alpha}{dt} = (q_{\alpha+1} - q_{\alpha-2})q_{\alpha-1} - q_\alpha + F, \quad \alpha = 1, \dots, 40; \quad q_{n+1} = q_1, \quad (38)$$

where $q_0 = q_n$, $q_{-1} = q_{n-1}$, and F is the forcing constant. We select $F = 8$, which is known to cause a chaotic behavior.

For both models, the *synthesized* ground-truth state vector at $t = 0$, $\mathbf{q}^{\text{tr}}(t = 0)$, is given as a sample of the normal distribution $\mathcal{N}(0_n, \mathbf{I}_n)$. The RV Q_0 is also assumed to follow this distribution. We simulate the L63 and L96 models using a fourth-order explicit Runge-Kutta algorithm with a time step $\Delta t = 0.01$.

5.1.2 Performance metrics

At each assimilation step k , the mean vector \bar{q}_k^a of the updated ensemble $\{q_k^{a(i)}\}_{i=1}^N$ is compared with the ground-truth state $\mathbf{q}^{\text{tr}}(t_k)$ using the component-average root MSE (RMSE) criterion, $rmse_k$, which is defined as

$$rmse_k = \frac{\|\bar{q}_k^a - \mathbf{q}^{\text{tr}}(t_k)\|}{\sqrt{n}}, \quad k = 1, 2, \dots, \kappa, \quad (39)$$

where $n = 3$ for L63 and $n = 40$ for L96. We use the *average* of these values $\{rmse_1, rmse_2, \dots, rmse_\kappa\}$ as a performance metric. We also monitor the average ensemble spread, \overline{ens} , defined as

$$\overline{ens} = \frac{1}{\kappa} \sum_{k=1}^{\kappa} \left[\frac{\text{tr}(\text{Cov}[Q_k^a])}{n} \right]^{1/2}, \quad (40)$$

to measure the contraction of the ensemble. In Eq. (40), $\text{tr}(\text{Cov}[Q_k^a])$ is estimated from ensemble $\{q_k^{a(1)}, \dots, q_k^{a(N)}\}$. The final metric that we consider herein is the average coverage probability f_{cv} of the 95%-confidence interval bounded between 2.5% and 97.5% quantiles of each marginal distribution, which is estimated as follows

$$f_{\text{cv}} = \frac{1}{n \times \kappa} \sum_{k=1, 2, \dots, \kappa; \alpha=1, 2, \dots, n} \mathbf{1}(q_\alpha^{\text{tr}}(t_k) \in I_{\alpha, k}), \quad (41)$$

where $I_{\alpha, k}$ is the estimated 95%-confidence interval of the α -th component of the updated ensemble at t_k . For small ensemble sizes, *e.g.*, $N = 20, 30$, and 60 , we compute the 90%, 93.3%, and 93.3%-coverage probability, respectively. The aforementioned performance metrics are commonly used to evaluate data assimilation algorithms, see *e.g.* [27, 43, 26].

5.1.3 Assessing errors of conditional mean approximations

Here we discuss a numerical experiment for assessing errors in approximating the CM using the ML-EnCMF and LL-EnCMF caused by small ensemble sizes N . First, a filtering procedure with κ assimilation steps using LL-EnCMF with a large ensemble size N_{ref} , *i.e.*, $N_{\text{ref}} \gg N$, is implemented as the reference EnCMF. The reference EnCMF provides an approximation of the CM at each assimilation step with negligible errors. Secondly, the forecast ensemble of the state and measurement at each assimilation step $k = 1, \dots, \kappa$, $\{(q_k^{f(i)}, y_k^{f(i)})\}_1^{N_{\text{ref}}}$, is randomly divided into two sets D_1 and D_2 with sizes N and $N_{\text{ref}} - N$, respectively. We then apply the ML-EnCMF and LL-EnCMF to approximate the CM using set D_1 as the input forecast ensemble and compute their component-average RMSE $rmse_\phi$ based on set D_2 as

$$rmse_{\phi,k} = \left[\frac{1}{N_{\text{ref}} - N} \sum_{(x^f, y^f) \in D_2} \frac{\|\hat{\phi}_{Q^f,k}(y^f|D_1) - \phi_{Q^f,k}^{\text{ref}}(y^f)\|^2}{n} \right]^{1/2}, \quad k = 1, \dots, \kappa, \quad (42)$$

where $\hat{\phi}_{Q^f,k}(\cdot|D_1)$ is the approximation of the CM map $\phi_{Q^f,k}$ via the ML-EnCMF (Eq. (26)) or the LL-EnCMF (Eq. 53) using the same N -sized forecast ensemble D_1 , and $\phi_{Q^f,k}^{\text{ref}}$ is the reference solution. Finally, we evaluate the average value of $rmse_{\phi,k}$ as

$$\overline{rmse}_\phi = \frac{1}{\kappa} \sum_{k=1}^{\kappa} rmse_{\phi,k}. \quad (43)$$

We use the average value \overline{rmse}_ϕ as an indicator for the misfit between the reference CM and its approximations. Running this numerical experiment requires a significant computational cost due to the condition $N_{\text{ref}} \gg N$. Therefore, we implement the experiment only for L63 cases.

5.2 Lorenz 63 system

This subsection reports the numerical results in tracking the L63 system. The data assimilation setting is described as follows. The initial ensembles are obtained by running the EnKF for 2000 assimilation steps. Then, we apply different techniques, *i.e.*, EnKF, ML-EnCMF, LL-EnCMF, for the next 2000 steps and evaluate their performance. We assume that every state is observed using the complete observation model:

$$y_k^{\text{obs}} = \mathbf{q}^{\text{tr}}(t_k) + \xi_k, \quad \Xi_k \sim \mathcal{N}(0_3, 2^2 \mathbf{I}_3). \quad (44)$$

Here we consider two scenarios for the observational period $\Delta T_{\text{obs}} := t_{k+1} - t_k$, $\Delta T_{\text{obs}} = 0.5$ and $\Delta T_{\text{obs}} = 1$. The selected observational periods are more extended than the usually reported ones ($\Delta T_{\text{obs}} \leq 0.1$), which amplifies the nonlinear characteristic.

We implement the ML-EnCMF using a two-hidden-layer dense ANN structure with 20 nodes at each hidden layer and the rectified linear unit as the activation function. The variance reduction technique, see Eq. (32), is applied such that $N \times M = 6000$. The forecast ensembles are divided into training and testing datasets with a size ratio of $N_T : N_S = 8 : 2$. At each assimilation step, we train the ANN for a maximum of 100 epochs, a learning rate of 0.001, and a batch size of 128 using Algorithm 2.

For the LL-EnCMF, we use inflation to avoid degeneracy. The inflation modifies the updated ensemble as

$$q^{a(i)} \leftarrow \bar{q}^a + \delta(q^{a(i)} - \bar{q}^a), \quad i = 1, \dots, N, \quad (45)$$

where \bar{q}^a is the mean of the ensemble, and δ (≥ 1) is the inflation coefficient. For a sufficiently large ensemble size, $\delta = 1$, in other words, the inflation is not active. We test different inflation coefficients in range $[1, 1.3]$ and select those yielding the best results. We note that the inflation technique is not applied for the ML-EnCMF or EnKF as an additional challenge for assessing these methods' robustness.

We run four data assimilation procedures, which perform $\kappa = 2000$ updating steps each following the above-described setting. The estimated performance metrics vary insignificantly when adding more assimilation procedures. The average performance metrics of the ML-EnCMF, EnKF, and LL-EnCMF for $\Delta T_{\text{obs}} = 0.5$ and $\Delta T_{\text{obs}} = 1$ are reported on Tab. 1 and 2, respectively. Here, one can compare the average RMSE of the updated mean with the standard deviation of the measurement error to obtain the relative magnitude.

Table 1: L63 system: performance metrics (average RMSE - average spread (average coverage probability f_{cv})) of the ML-EnCMF compared with the LL-EnCMF and EnKF with $\Delta T_{\text{obs}} = 0.5$. The LL-EnCMF uses inflation coefficients 1.25, 1.2, and 1.05 for ensemble sizes 20, 30, and 60, respectively. For $N = 20, 30$, and 60 we compute the coverage probability for 90%, 93.3%, and 93.3%-confident intervals, respectively.

N	20	30	60	100	200
ML-EnCMF	1.27-0.96 (0.83)	1.11 - 0.97 (0.90)	0.94 - 0.97 (0.92)	0.86 - 0.97 (0.94)	0.81 - 0.96 (0.95)
EnKF	1.37-1.17 (0.86)	1.29 - 1.21 (0.90)	1.24 - 1.27 (0.93)	1.23 - 1.29 (0.93)	1.22 - 1.30 (0.93)
LL-EnCMF	1.43 - 1.02 (0.85)	1.21-1.04 (0.92)	0.99-0.91 (0.93)	0.90-0.89 (0.93)	0.85 - 0.90 (0.95)

Table 2: L63 system: performance metrics (average RME - median RMSE (average coverage probability f_{cv})) of the ML-EnCMF compared with the LL-EnCMF and EnKF with $\Delta T_{\text{obs}} = 1$. The LL-EnCMF uses inflation coefficients 1.25, 1.2, and 1.1 for ensemble sizes 20, 30, and 60, respectively. For $N = 20, 30$, and 60 we compute the coverage probability for 90%, 93.3%, and 93.3%-confident intervals, respectively.

N	20	30	60	100	200
ML-EnCMF	1.66 - 1.19 (0.82)	1.50 - 1.20 (0.90)	1.22 - 1.21 (0.90)	1.14-1.18 (0.93)	1.06 - 1.17 (0.93)
EnKF	1.67 -1.46 (0.87)	1.61 - 1.55 (0.92)	1.53 - 1.57 (0.91)	1.51 - 1.60 (0.93)	1.51 - 1.61 (0.93)
LL-EnCMF	2.52 - 1.11 (0.72)	1.78-1.15 (0.86)	1.35 - 1.12 (0.90)	1.18 - 1.06 (0.93)	1.05 - 1.06 (0.93)

In comparison with the EnKF, the ML-EnCMF consistently exhibits better performance. The improvement increases with the ensemble size because the nonlinearity is better approximated. For instance, with $N = 200$, the average RMSE is reduced by 34% and 30%, for $T_{\text{obs}} = 0.5$ and $T_{\text{obs}} = 1$, respectively. A similar trend is also observed for the average ensemble spread. While the average ensemble spread is significantly reduced using the ML-EnCMF, the average coverage probability is still comparable with the EnKF. It is noteworthy that, the proposed filter shows better performance not only with large ensemble sizes but also small ones, *e.g.*, $N \leq 60$. For $N = 30$ and $T_{\text{obs}} = 1$, the EnKF and ML-EnCMF show similar performance because the ANN

cannot capture the CM’s nonlinearity with a small ensemble size and under significant uncertainty (compared with the case of $N = 30$ and $T_{\text{obs}} = 0.5$)

Compared with the LL-EnCMF, the ML-EnCMF exhibits significantly better performance in terms of average RMSE for small ensemble sizes $N \leq 60$, which is a crucial advantage of the ML-EnCMF because the ensemble size is usually small in practice. Notably, the LL-EnCMF exhibits worse performance compared with the EnKF for $N = 20$ and $N = 30$ ($T_{\text{obs}} = 1$). Another disadvantage of the LL-EnCMF compared to the ML-EnCMF is the strong dependence on the inflation coefficient. For example, the average RMSEs obtained by the LL-EnCMF without using the inflation technique in case $\Delta T_{\text{obs}} = 0.5$ are 2.33 and 1.65, for ensemble sizes of 20 and 30, respectively. For $N \geq 100$, both methods exhibit similar performance metrics.

We implement the numerical experiment discussed in Sec. 5.1.3 using $N_{\text{ref}} = 10000$. The average RMSEs of the CM approximations using the ML-EnCMF and LL-EnCMF \overline{rmse}_ϕ are plotted in Fig. 4. The CM approximating errors of ML-EnCMF are smaller compared with the LL-EnCMF for $N \leq 60$. Furthermore, those errors of ML-EnCMF and LL-EnCMF are getting closer as the ensemble size increases. For $N \geq 100$, the approximating errors yielded by the ML-EnCMF and LL-EnCMF are negligible compared with the standard deviation of the measurement error. This trend of the approximating CM errors explains the previous numerical observation that the ML-EnCMF outperforms the LL-EnCMF for small ensemble sizes ($N \leq 60$), while for $N \geq 100$ both methods exhibit similar performance metrics.

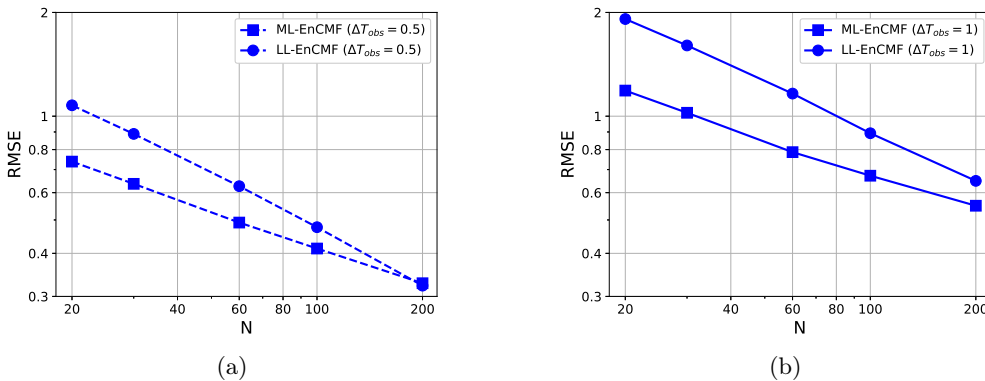


Figure 4: Lorenz 63: average RMSE of the CM approximations using LL-EnCMF and ML-EnCMF \overline{rmse}_ϕ , a) $\Delta T_{\text{obs}} = 0.5$, b) $\Delta T_{\text{obs}} = 1$.

5.3 Lorenz 96 system

This subsection reports the numerical results in tracking the L96 system. The initial ensembles are obtained by running the EnKF for 2000 assimilation steps using a complete observation model with standard Gaussian noise. Then, we apply different techniques, *i.e.*, EnKF, ML-EnCMF, and LL-EnCMF, for the next 2000 steps using an incomplete observation model:

$$y_k^{\text{obs}} = \mathbf{H}q^{\text{tr}}(t_k) + \xi_k, \quad \xi_k \sim \mathcal{N}(0_{20}, 0.5 \mathbf{I}_{20}), \quad (46)$$

where $\mathbf{H} \in \mathbb{R}^{m \times n}$ with $m = 20$ is a linear operator that selects the even-indexed components $q_2^{\text{tr}}(t_k), q_4^{\text{tr}}(t_k), \dots, q_{40}^{\text{tr}}(t_k)$ from $q^{\text{tr}}(t_k)$. We examine the observation time interval $\Delta T_{\text{obs}} = 0.4$. For reference, $\Delta T_{\text{obs}} = 0.05$ is comparable to 6 h in a weather forecast model [31, 32]. Due to the

long observational period and the incomplete observation model, the relation between an observed component and an unobserved one is highly nonlinear. Moreover, the distribution of the forecast is no longer Gaussian. The considered setting is hence referred to as *the hard case* in [27].

The implemented ML-EnCMF in this example uses a one-hidden-layer ANN structure with 40 nodes in the middle layer for approximating the CM. Instead of a dense ANN structure, a *localized* one is used. There are two motivations for applying such structures: i) to explicitly remove correlation between components with a long spatial range — a widely applied technique for the EnKF and LL-EnCMF, ii) to reduce the number of hyperparameters required to be trained. Without using the localized structure, the performance of the ML-EnCMF may not be optimized for small ensemble sizes. The ANN structure designed for tracking the L96 system is described as follows. Let $u := g_{NN}(y)$, where $g_{NN} : \mathbb{R}^{20} \rightarrow \mathbb{R}^{40}$ (see Eq. (26)). The vector u is evaluated as

$$\begin{aligned} v_\alpha &= \sigma_1 \left(\sum_{\alpha'=2,4,\dots,40} w_{\alpha\alpha'}^1 y_{\alpha'/2} \mathbf{1}(z_{\alpha\alpha'} \leq l_1) + \gamma_\alpha^1 \right), \quad \alpha = 1, \dots, 40 \\ u_\alpha &= \sum_{\alpha'=1}^{40} w_{\alpha\alpha'}^2 v_{\alpha'} \mathbf{1}(z_{\alpha\alpha'} \leq l_2) + \gamma_\alpha^2, \quad \alpha = 1, \dots, 40 \end{aligned} \quad (47)$$

where $z_{\alpha\alpha'} = \min(|\alpha - \alpha'|, 40 - |\alpha - \alpha'|)$ is the spatial distance between components α and α' of the state vector, v_1, \dots, v_{40} are the output of the hidden layer, $w_{\alpha\alpha'}^1$ ($w_{\alpha\alpha'}^2$) and γ_α^1 (γ_α^2) are weights and biases of the hidden (output) layer, respectively, σ_1 is the activation function of the hidden layers, and l_1 and l_2 are the *localization lengths*. In a dense version, l_1 and l_2 are equal to 20. Here we choose $l_1 = l_2 = 3$ to incorporate sparsity. The number of trainable weights is hence reduced to 420 from 2400 in the dense version. For training the ANN, the variance reduction technique, see Eq. (32), is applied such that $N \times M \simeq 10000$. The other training settings are as described in the L63 example above.

The localization approach is also applied for the EnKF, the linear part of the ML-EnCMF, and the LL-EnCMF. For implementing the EnKF and the linear part in the ML-EnCMF, see Eq. (26), we use the Cohn-Gaspari covariance tapering [17] with $c = 20$ for $N \leq 300$ and $c = 30$ for $N = 400$. For the LL-EnCMF, we use the localization technique described in [27].

The inflation technique is applied for all tested filters in this *hard case* example. The selected inflation coefficient for the EnKF is $\delta = 1.005$. For the ML-EnCMF and LL-EnCMF, we test different inflation coefficients in ranges $[1., 1.05]$ and $[1., 1.3]$, respectively, and choose the best results. The selected inflation coefficients for the ML-EnCMF are 1.05 ($N = 100$), 1.03 ($N \in \{100, 200\}$), 1.01 ($N = 300$), and 1.005 ($N = 400$). Those of the LL-EnCMF are 1.2 ($N = 100$), 1.18 ($N \in \{150, 200\}$), and 1.1 ($N \in \{300, 400\}$).

We run four data assimilation procedures with different generating seeds. The average performance metrics are reported on Tab 3. For cross-reference, in [27], the average RMSEs obtained using the LL-EnCMF and EnKF with $N = 400$ are 0.68 and 0.79, respectively. These values are slightly different from our EnKF and LL-EnCMF results due to the non-deterministic of the synthesized measurement data and statistical errors when evaluating the average values. The LL-EnCMF performances for smaller ensemble sizes ($N < 400$) are not reported in [27].

We observe a pattern similar to the L63 system. Indeed, the ML-EnCMF consistently outperforms the EnKF in terms of average RMSE and average spread thanks to its ability to approximate the nonlinearity of the CM. In terms of 95%-coverage probability, the ML-EnCMF and EnKF have comparable performance.

Compared with the LL-EnCMF, the ML-EnCMF and LL-EnCMF exhibit comparable performance for $N = 400$. However, the ML-EnCMF outperforms the LL-EnCMF for smaller ensemble sizes due to the significant statistical error when approximating the CM via the likelihood function

Table 3: L96 system: performance metrics, *i.e.*, average RMSE - average spread (average coverage probability f_{cv}), of the ML-EnCMF compared with the LL-EnCMF and EnKF with $\Delta T_{\text{obs}} = 0.4$.

N	100	150	200	300	400
ML-EnCMF	0.84 - 0.70 (0.94)	0.79 - 0.67 (0.94)	0.75 - 0.67 (0.94)	0.72 - 0.64 (0.94)	0.69 - 0.63 (0.94)
EnKF	0.88 - 0.71 (0.92)	0.85 - 0.73 (0.93)	0.83 - 0.74 (0.94)	0.84 - 0.76 (0.94)	0.83 - 0.78 (0.95)
LL-EnCMF	1.21 - 0.74 (0.89)	0.92 - 0.73 (0.93)	0.88 - 0.76 (0.95)	0.74 - 0.68 (0.95)	0.70 - 0.69 (0.96)

with a small ensemble size. Notably, the LL-EnCMF shows worse performance compared to the EnKF for $N \leq 200$.

5.4 Discussion on large-scale applications

The ML-EnCMF practicality requires that the computational time for ANN training and other computational tasks is shorter than the prediction range. Data assimilation applications have varying temporal prediction ranges in practice. For example, the prediction ranges of weather and climate applications can be classified into short-range (one to two days), medium-range (days to weeks), and (sub)seasonal-range.

In our numerical experiments, the typical training time for approximating the CM is 22-27 (s) and 32-37 (s) for L63 with an ensemble size of 200 and L96 with an ensemble size of 400, respectively, on one core of a 2.8 GHz Quad-Core Intel Core i7 CPU. The training time in the L96 case does not significantly grow owing to the sparse ANN structure and the applied early stopping. Furthermore, the sizes of augmented datasets are identical for the L63 and L96 experiments.

We project that the ML-EnCMF fulfills the computational time constraint for applications with moderately high-dimensional states with a short (or longer) prediction range and for larger systems with medium to seasonal ranges. Moreover, the ANN training process is accelerated by using advanced hardware, *e.g.*, GPU, as well as improving the algorithm, *e.g.*, the ANN can be pre-trained to learn the manifold of the states to reduce the online-training time on that ANN [28, 42]. In [11], Chantry *et al.* discuss in more detail the opportunities and challenges for machine learning in weather and climate modeling. We will investigate applications of the ML-EnCMF for larger systems with the vector-field state in our future research.

6 Conclusion

This work analyzes the properties of the CMF and develops the ML-based ensemble method for its implementation. The CMF’s updated mean matches that of the posterior, obtained by applying Bayes’ rule on the filter’s forecast distribution. Moreover, the CMF’s updated covariance coincides with the expected conditional covariance.

We develop ML-EnCMF using ANNs and based on the orthogonal projection property of the CM. The filter naturally inherits the robust performance of the EnKF in closely linear Gaussian scenarios, owing to the combination of ANNs with the EnKF’s updating map. A systematical methodology for integrating machine learning approaches into the EnCMF is developed. We apply a variance reduction technique to augment the dataset and reduce statistical errors. Moreover, we perform a model selection procedure at each updating step for selecting element-wisely the applied

filter, *i.e.*, either the EnKF or the ML-EnCMF. Notably, the model selection ensures that our filter outperforms the EnKF.

We demonstrate the ML-EnCMF performance using the Lorenz-63 and Lorenz-96 systems. Notably, for the Lorenz-96 system, we propose a localized ANN structure to avoid overestimated long-range correlation and reduce computational costs for ANN training. In summary, the proposed filtering method exhibits considerable improvement compared with the commonly used methods, such as the EnKF and the LL-EnCMF, particularly for small ensemble sizes.

For future work, the following developments will be considered: i) developing the ML-EnCMF for large-scale systems with vector-field state, ii) enhancing the computational efficiency of the ML-EnCMF and iii) combining the filter with the method of A-optimal design of experiments.

Acknowledgment

This publication was supported by funding from the Alexander von Humboldt Foundation and King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under award numbers

URF/1/2281-01-01 and URF/1/2584-01-01 in the KAUST Competitive Research Grants Programs, respectively. R. Tempone is a member of the KAUST SRI Center for Uncertainty Quantification in Computational Science and Engineering. Simulations were performed with computing resources granted by RWTH Aachen University under project rwth0632.

A Kalman filters

This appendix summarizes the Kalman filter and its ensemble version.

A.1 Kalman filter

In the linear-Gaussian setting, the RVs Q_k^f and Y_k^f are considered Gaussian. Consequently, the conditional PDF stated in Eq. (7) is simplified to a normal distribution PDF. Let $h(q) = \mathbf{H}q$, where $\mathbf{H} \in \mathbb{R}^{m \times n}$. The closed-form of the RV Q_k^a , whose probability PDF is identical to that of the local posterior [22], is expressed as follows:

$$Q_k^a = Q_k^f + \mathbf{K}_k^l (y_k^{\text{obs}} - Y_k^f), \quad (48a)$$

$$\mathbf{K}_k^l = \Sigma_{Q_k^f} \mathbf{H}^\top (\Sigma_{\Xi_k} + \mathbf{H} \Sigma_{Q_k^f} \mathbf{H}^\top)^{-1}, \quad (48b)$$

where Q_k^f is evaluated using the model Ψ_k assumed to be linear, \mathbf{K}_k^l is the Kalman gain, and $\Sigma_{Q_k^f}$ and Σ_{Ξ_k} are the covariance matrices of the RVs Q_k^f and Ξ_k , respectively. The transformation in Eq. (48a) is a linear version of the map \mathcal{T}_k (Eq.(6c)). In the linear-Gaussian setting, the RV Q_k^a obtained using Eq. (48a) is a Gaussian RV, hence, fully characterized by its mean and covariance. From Eq. (48a), the formulations for updating these statistical moments can be straightforwardly derived, known as the KF [22].

A.2 Ensemble Kalman filter

The EnKF is an ensemble implementation of the KF for dealing with nonlinear dynamical systems, in which the forecast and updated RVs are represented using the ensembles of their samples, thereby allowing the approximation of non-Gaussian distributions [15]. However, the EnKF applicability

still requires a linear observation map. Let $\{q_{k-1}^{a(i)}\}_{i=1}^N$ be the N -sample ensemble of the RV Q_{k-1}^a . The ensembles $\{q_k^{f(i)}\}_{i=1}^N$ and $\{y_k^{f(i)}\}_{i=1}^N$ of the RVs Q_k^f and Y_k^f , respectively, are obtained using Eq. (4). For the update step, we first estimate the mean \bar{q}_k^f and the covariance matrix $\Sigma_{Q_k^f}$ of the RV Q_k^f using the forecast ensembles:

$$\begin{aligned}\bar{q}_k^f &\approx \frac{1}{N} \sum_{i=1}^N q_k^{f(i)}, \\ \Sigma_{Q_k^f} &\approx \frac{1}{N} \sum_{i=1}^N \left(q_k^{f(i)} - \bar{q}_k^f \right) \left(q_k^{f(i)} - \bar{q}_k^f \right)^\top.\end{aligned}\tag{49}$$

The obtained covariance matrix is then used to compute the Kalman gain by evaluating Eq. (48b). Subsequently, the EnKF uses the linear update map of the KF, as expressed in Eq. (48a), and employs the ensemble technique to approximate the updated ensemble:

$$q_k^{a(i)} = q_k^{f(i)} + \mathbf{K}_k^l (y_k^{\text{obs}} - y_k^{f(i)}), \quad i = 1, \dots, N,\tag{50}$$

where $\{q_k^{a(1)}, \dots, q_k^{a(N)}\}$ is the updated ensemble representing the RV Q_k^a .

B Likelihood-based approximation of the CM and the LL-EnCMF

Combining Eqs. (3) and (9), the conditional mean ϕ_{Q^f} can be evaluated via the likelihood function as

$$\phi_{Q^f}(y) = \int_{\mathbb{R}^n} q \frac{\pi_{Q^f}(q) \pi_\Xi(y - h(q))}{\int_{\mathbb{R}^n} \pi_{Q^f}(q') \pi_\Xi(y - h(q')) dq'} dq.\tag{51}$$

Similarly the conditional expectation of the second moment $\mathbb{E} \left[(Q^f)^2 | Y^f \right]$ can be expressed as

$$\mathbb{E} \left[(Q^f)^2 | Y^f = y \right] = \int_{\mathbb{R}^n} q^2 \frac{\pi_{Q^f}(q) \pi_\Xi(y - h(q))}{\int_{\mathbb{R}^n} \pi_{Q^f}(q') \pi_\Xi(y - h(q')) dq'} dq.\tag{52}$$

We recall from Sec. 3.2 that $\{q^{f(i)}\}_{i=1}^N$ and $\{y^{f(i)}\}_{i=1}^N$ are the forecast ensembles of the state and observation, respectively. Implementing the EnCMF, see Eq. (23), requires evaluating map ϕ_{Q^f} for $\{y^{f(i)}\}_{i=1}^N$ and y^{obs} . Using the likelihood-based approach, the values of $\phi_{Q^f}(y^*)$ with $y^* \in \{y^{f(1)}, \dots, y^{f(n)}, y^{\text{obs}}\}$ are respectively approximated as

$$\widehat{\phi}_{Q^f}(y^*) = \frac{1}{N} \sum_{i=1}^N q^{f(i)} \frac{\pi_\Xi(y^* - h(q^{f(i)}))}{\frac{1}{N} \sum_{i=1}^N \pi_\Xi(y^* - h(q^{f(i)}))}, \quad y^* \in \{y^{f(1)}, \dots, y^{f(n)}, y^{\text{obs}}\},\tag{53}$$

which is the ensemble-based version of the Eq. (51). The updated ensemble of the LL-EnCMF is then obtained as

$$q^{a(i)} = q^{f(i)} + \widehat{\phi}_{Q^f}(y^{\text{obs}}) - \widehat{\phi}_{Q^f}(y^{f(i)}).\tag{54}$$

C Analysis of the variance reduced estimator

To motivate that \widehat{J}^{vr} may provide an estimation of J with reduced variance, we note that

$$\begin{aligned} \lim_{M \rightarrow \infty} \widehat{J}^{\text{vr}}(\boldsymbol{\theta}|\mathcal{D}_T) &=_{\text{a.s.}} \frac{1}{N_T} \sum_{i=1, \dots, N_T} \mathbb{E} \left[\|q^{f(i)} - g_1(h(q^{f(i)} + \Xi)) \right. \\ &\quad \left. - g_{\text{NN}}(h(q^{f(i)} + \Xi; \boldsymbol{\theta})) \|^2 \right] \\ &= \frac{1}{N_T} \sum_{i=1, \dots, N_T} \mathbb{E} \left[A(Q^f, \Xi; \boldsymbol{\theta}) | Q^f = q^{f(i)} \right], \end{aligned} \quad (55)$$

\mathbb{P} -almost surely, where

$$A(q, \xi; \boldsymbol{\theta}) \equiv \|q - g_1(h(q) + \xi) - g_{\text{NN}}(h(q) + \xi; \boldsymbol{\theta})\|^2. \quad (56)$$

Let $\widehat{J}^{\text{vr}*}(\boldsymbol{\theta}|\mathcal{D}_T)$ denote the right-hand-side term in Eq. (55). To gain some further insight, let us assume that the ensemble members $q^{f(i)}$, $i = 1, 2, \dots, N$, are i.i.d. samples for the sake of simplicity. We can then approximately quantify the statistical errors of the approximation $\widehat{J}(\boldsymbol{\theta}|\mathcal{D}_T)$ and $\widehat{J}^{\text{vr}}(\boldsymbol{\theta}|\mathcal{D}_T)$, respectively, as

$$\text{Var} \left[\widehat{J}(\boldsymbol{\theta}|\mathcal{D}_T) - J(\boldsymbol{\theta}) \right] = \frac{\text{Var} \left[A(Q^f, \Xi; \boldsymbol{\theta}) \right]}{N_T}, \quad (57a)$$

$$\text{Var} \left[\widehat{J}^{\text{vr}*}(\boldsymbol{\theta}|\mathcal{D}_T^a) - J(\boldsymbol{\theta}) \right] = \frac{\text{Var} \left[\mathbb{E} \left[A(Q^f, \Xi; \boldsymbol{\theta}) | Q^f \right] \right]}{N_T}. \quad (57b)$$

Using the law of total variance,

$$\text{Var} \left[A(Q^f, \Xi; \boldsymbol{\theta}) \right] = \text{Var} \left[\mathbb{E} \left[A(Q^f, \Xi; \boldsymbol{\theta}) | Q^f \right] \right] + \mathbb{E} \left[\text{Var} \left[A(Q^f, \Xi; \boldsymbol{\theta}) | Q^f \right] \right], \quad (58)$$

we obtain

$$\text{Var} \left[\mathbb{E} \left[A(Q^f, \Xi; \boldsymbol{\theta}) | Q^f \right] \right] \leq \text{Var} \left[A(Q^f, \Xi; \boldsymbol{\theta}) \right]. \quad (59)$$

References

- [1] H. D. Abarbanel, P. J. Rozdeba, and S. Shirman. Machine learning: deepest learning as statistical data assimilation problems. *Neural Computation*, 30(8):2025–2055, 2018.
- [2] M. Ades and P. J. van Leeuwen. An exploration of the equivalent weights particle filter. *Quarterly Journal of the Royal Meteorological Society*, 139(672):820–840, 2013.
- [3] A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas. A fast and scalable method for A-optimal design of experiments for infinite-dimensional bayesian nonlinear inverse problems. *SIAM Journal on Scientific Computing*, 38(1):A243–A272, 2016.
- [4] S. Amari. Backpropagation and stochastic gradient descent method. *Neurocomputing*, 5(4-5):185–196, 1993.
- [5] M. Asch, M. Bocquet, and M. Nodet. *Data assimilation: methods, algorithms, and applications*. SIAM, 2016.

- [6] P. Bauer, B. Stevens, and W. Hazeleger. A digital twin of Earth for the green transition. *Nature Climate Change*, 11(2):80–83, 2021.
- [7] A. Bobrowski. *Functional analysis for Probability and Stochastic Process*. Cambridge University press, 2005.
- [8] M. Bocquet, J. Brajard, A. Carrassi, and L. Bertino. Bayesian inference of chaotic dynamics by merging data assimilation, machine learning and expectation-maximization. *Foundations of Data Science*, 2(1):55–80, 2020.
- [9] J. Brajard, A. Carrassi, M. Bocquet, and L. Bertino. Combining data assimilation and machine learning to emulate a dynamical model from sparse and noisy observations: A case study with the Lorenz 96 model. *Journal of Computational Science*, 44:101171, 2020.
- [10] F. Burden and D. Winkler. Bayesian regularization of neural networks. In D. J. Livingstone, editor, *Artificial Neural Networks: Methods and Applications*, pages 23–42. Humana Press, Totowa, NJ, 2009.
- [11] M. Chantry, H. Christensen, P. Dueben, and T. Palmer. Opportunities and challenges for machine learning in weather and climate modelling: hard, medium and soft ai. *Philosophical Transactions of the Royal Society A*, 379(2194):20200083, 2021.
- [12] B. Crestel, A. Alexanderian, G. Stadler, and O. Ghattas. A-optimal encoding weights for non-linear inverse problems, with application to the Helmholtz inverse problem. *Inverse problems*, 33(7):074008, 2017.
- [13] R. Durrett. *Probability: theory and examples*. Cambridge University press, 2019.
- [14] O. G. Ernst, B. Sprungk, and H.-J. Starkloff. Bayesian inverse problems and Kalman filters. In *Extraction of Quantifiable Information from Complex Systems*, pages 133–159. Springer, 2014.
- [15] G. Evensen. *Data assimilation: The ensemble Kalman filter*. Springer Berlin Heidelberg, 2009.
- [16] D. J. Gagne II, H. M. Christensen, A. C. Subramanian, and A. H. Monahan. Machine learning for stochastic parameterization: Generative adversarial networks in the Lorenz '96 model. *Journal of Advances in Modeling Earth Systems*, 12(3):e2019MS001896, 2020. e2019MS001896 10.1029/2019MS001896.
- [17] G. Gaspari and S. E. Cohn. Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(554):723–757, 1999.
- [18] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [19] H. Hoel, K. J. H. Law, and R. Tempone. Multilevel ensemble Kalman filtering. *SIAM Journal on Numerical Analysis*, 54(3):1813–1839, 2016.
- [20] P. L. Houtekamer and H. L. Mitchell. Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, 126(3):796–811, 1998.
- [21] A. H. Jazwinski. *Stochastic processes and filtering theory*. Courier Corporation, 2007.
- [22] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45, 03 1960.

- [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- [24] K. J. H. Law, A. Stuart, and K. Zygalakis. *Data assimilation: : A Mathematical Introduction*, volume 214. Springer, 2015.
- [25] K. J. H. Law, H. Tembine, and R. Tempone. Deterministic mean-field ensemble Kalman filtering. *SIAM Journal on Scientific Computing*, 38(3):A1251–A1279, 2016.
- [26] Y. Lee and A. J. Majda. State estimation and prediction using clustered particle filters. *Proceedings of the National Academy of Sciences*, 113(51):14609–14614, 2016.
- [27] J. Lei and P. Bickel. A moment matching ensemble filter for nonlinear non-Gaussian data assimilation. *Monthly Weather Review*, 139(12):3964–3973, 2011.
- [28] T. Lin and H. Zha. Riemannian manifold learning. *IEEE transactions on pattern analysis and machine intelligence*, 30(5):796–809, 2008.
- [29] E. N. Lorenz. Deterministic nonperiodic flow. *Journal of Atmospheric Sciences*, 20(2):130–141, 1963.
- [30] E. N. Lorenz. Predictability: A problem partly solved. In *Proc. Seminar on Predictability*, volume 1, 1996.
- [31] E. N. Lorenz and K. A. Emanuel. Optimal sites for supplementary weather observations: simulation with a small model. *Journal of the Atmospheric Sciences*, 55(3):399–414, 1998.
- [32] A. J. Majda and J. Harlim. *Filtering complex turbulent systems*. Cambridge University Press, 2012.
- [33] H. G. Matthies, E. Zander, B. V. Rosić, and A. Litvinenko. Parameter estimation via conditional expectation: a Bayesian inversion. *Advanced Modeling and Simulation in Engineering Sciences*, 3(1):1–21, 2016.
- [34] S. Pathiraja, S. Reich, and W. Stannat. Mckean–Vlasov SDEs in nonlinear filtering. *SIAM Journal on Control and Optimization*, 59(6):4188–4215, 2021.
- [35] H. K. Pradhan, C. Völker, S. N. Losa, A. Bracher, and L. Nerger. Global assimilation of ocean-color data of phytoplankton functional types: Impact of different data sets. *Journal of Geophysical Research: Oceans*, 125(2):e2019JC015586, 2020.
- [36] A. Rasheed, O. San, and T. Kvamsdal. Digital twin: values, challenges and enablers from a modeling perspective. *IEEE Access*, 8:21980–22012, 2020.
- [37] S. Reich. A dynamical systems framework for intermittent data assimilation. *BIT Numerical Mathematics*, 51(1):235–249, 2011.
- [38] S. Reich. Data assimilation: the Schrödinger perspective. *Acta Numerica*, 28:635–711, 2019.
- [39] S. Reich and C. Cotter. *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press, 2015.
- [40] S. Reich and C. J. Cotter. Ensemble filter techniques for intermittent data assimilation. In M. Cullen, M. A. Freitag, S. Kindermann, and R. Scheichl, editors, *Large Scale Inverse Problems: Computational Methods and Applications in the Earth Sciences*, pages 91–134. De Gruyter, 2013.

- [41] R. H. Reichle. Data assimilation methods in the earth sciences. *Advances in Water Resources*, 31(11):1411–1418, 2008.
- [42] C. Soize and R. Ghanem. Probabilistic learning on manifolds constrained by nonlinear partial differential equations for small datasets. *Computer Methods in Applied Mechanics and Engineering*, 380:113777, 2021.
- [43] A. Spantini, R. Baptista, and Y. Marzouk. Coupling techniques for nonlinear ensemble filtering. *arXiv preprint arXiv:1907.00389*, 2019.
- [44] Q. Tang, L. Mu, D. Sidorenko, H. Goessling, T. Semmler, and L. Nerger. Improving the ocean and atmosphere in a coupled ocean–atmosphere model by assimilating satellite sea-surface temperature and subsurface profile data. *Quarterly Journal of the Royal Meteorological Society*, 146(733):4014–4029, 2020.
- [45] P. J. van Leeuwen, H. R. Künsch, L. Nerger, R. Potthast, and S. Reich. Particle filters for high-dimensional geoscience applications: A review. *Quarterly Journal of the Royal Meteorological Society*, 145(723):2335–2365, 2019.
- [46] M. Verlaan and A. W. Heemink. Nonlinearity in data assimilation applications: a practical method for analysis. *Monthly Weather Review*, 129(6):1578–1589, 2001.
- [47] S. Vetra-Carvalho, P. J. van Leeuwen, L. Nerger, A. Barth, M. U. Altaf, P. Brasseur, P. Kirchgessner, and J.-M. Beckers. State-of-the-art stochastic data assimilation methods for high-dimensional non-Gaussian problems. *Tellus A: Dynamic Meteorology and Oceanography*, 70(1):1–43, 2018.
- [48] J. Vondřejc and H. G. Matthies. Accurate computation of conditional expectation for highly nonlinear problems. *SIAM/ASA Journal on Uncertainty Quantification*, 7(4):1349–1368, 2019.