

Bayesian inference of chaotic dynamics by merging data assimilation, machine learning and expectation-maximization

Article

Published Version

Bocquet, M., Brajard, J., Carrassi, A. ORCID: <https://orcid.org/0000-0003-0722-5600> and Bertino, L. (2020) Bayesian inference of chaotic dynamics by merging data assimilation, machine learning and expectation-maximization. *Foundations of Data Science*, 2 (1). pp. 55-80. ISSN 2639-8001 doi: <https://doi.org/10.3934/fods.2020004> Available at <https://centaur.reading.ac.uk/89798/>

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

To link to this article DOI: <http://dx.doi.org/10.3934/fods.2020004>

Publisher: American Institute of Mathematical Sciences

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

BAYESIAN INFERENCE OF CHAOTIC DYNAMICS BY MERGING DATA ASSIMILATION, MACHINE LEARNING AND EXPECTATION-MAXIMIZATION

MARC BOCQUET*

CEREA, joint laboratory École des Ponts ParisTech and EDF R&D
Université Paris-Est, Champs-sur-Marne, France

JULIEN BRAJARD

Nansen Environmental and Remote Sensing Center, Bergen, Norway
and Sorbonne University, CNRS-IRD-MNHN, LOCEAN, Paris, France

ALBERTO CARRASSI

Department of Meteorology, University of Reading and NCEO, United-Kingdom
and Mathematical Institute, Utrecht University, The Netherlands

LAURENT BERTINO

Nansen Environmental and Remote Sensing Center, Bergen, Norway

(Communicated by Kody J. H. Law)

ABSTRACT. The reconstruction from observations of high-dimensional chaotic dynamics such as geophysical flows is hampered by (i) the partial and noisy observations that can realistically be obtained, (ii) the need to learn from long time series of data, and (iii) the unstable nature of the dynamics. To achieve such inference from the observations over long time series, it has been suggested to combine data assimilation and machine learning in several ways. We show how to unify these approaches from a Bayesian perspective using expectation-maximization and coordinate descents. In doing so, the model, the state trajectory and model error statistics are estimated all together. Implementations and approximations of these methods are discussed. Finally, we numerically and successfully test the approach on two relevant low-order chaotic models with distinct identifiability.

1. Introduction.

1.1. Data assimilation. The primary aim of data assimilation (DA), as used in the geosciences, is to accurately infer the state of a physical system from observations and theoretical knowledge of the system, such as the governing equations [12]. DA techniques are particularly useful to forecast chaotic geofluids, such as the atmosphere, the ocean, and the climate because of (i) the intrinsic dynamical instability of the flows and, (ii) the errors made in the governing equations and their numerical implementation [31]. Like in machine learning (ML), the observation sets can be huge (especially in operational weather forecast centers). However,

2010 *Mathematics Subject Classification.* Primary: 49J15, 65C60; Secondary: 86-08.

Key words and phrases. Data assimilation, machine learning, neural networks, chaotic dynamical systems, expectation-maximization, coordinate descent.

* Corresponding author: Marc Bocquet.

as opposed to machine learning, the models are built on physical principles and are often computationally very costly.

A Bayesian view on inference is universally adopted in DA [2]. Typically, one looks for the unknown trajectory $\mathbf{x}_{0:K} = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K\}$ of the model state, i.e. the collection of all model states $\mathbf{x}_k \in \mathbb{R}^{N_x}$, from time t_0 to time t_K . The physical system is assumed to be observed: \mathbf{y}_k is the vector of observations at time t_k , while $\mathbf{y}_{0:K} = \{\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_K\}$ denotes the full set of observations within the time window $[t_0, \dots, t_K]$. Bayes' inference then yields

$$p(\mathbf{x}_{0:K}|\mathbf{y}_{0:K}) = \frac{p(\mathbf{y}_{0:K}|\mathbf{x}_{0:K})p(\mathbf{x}_{0:K})}{p(\mathbf{y}_{0:K})}, \quad (1)$$

where $p(\mathbf{x}_{0:K}|\mathbf{y}_{0:K})$ is the conditional probability density function (pdf), $p(\mathbf{y}_{0:K}|\mathbf{x}_{0:K})$ is the likelihood and is usually specified, and $p(\mathbf{x}_{0:K})$ is the prior pdf which depends on the dynamical model.

Because the observations are assimilated usually soon after being acquired and because the dynamics are often unstable, the inference is usually sequential in the geosciences. In that case, the inference can be seen as a hidden Markov problem and can be written, assuming Markovian dynamics, as a sequence of an *analysis* and a *forecast* step [12]:

$$\text{Analysis step:} \quad p(\mathbf{x}_k|\mathbf{y}_{0:k}) \propto p(\mathbf{y}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{y}_{0:k-1}), \quad (2a)$$

$$\text{Forecast step:} \quad p(\mathbf{x}_{k+1}|\mathbf{y}_{0:k}) = \int d\mathbf{x}_k p(\mathbf{x}_{k+1}|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{y}_{0:k}). \quad (2b)$$

For this inference to be scalable to high-dimensional models, approximations have been proposed, which yield a wealth of powerful DA techniques [2]. For instance, the 4D-Var method seeks the maximum a posteriori of (1), while Gaussian filtering solutions such as the ensemble Kalman filter seek an approximate solution following (2). See [24, 7] for a comparison of their algorithmic merits.

In a challenging context with real observations, the inference bears not only on $\mathbf{x}_{0:K}$, but often also on some coefficients tuning the error statistics. These are needed to accurately specify, for instance the likelihood in (2a), or the Chapman-Kolmogorov transition pdf in (2b). It is also common to infer a selection of physical parameters of an established numerical model. But the physical model remains key in any DA inference.

1.2. The model as a control variable. Recently, fostered by the success of data-driven ML techniques, it has been suggested to partially or entirely learn a model of the physical system dynamics [26] from observations. There is certainly a long way before any data-driven techniques can become a substitute to a full realistic model in the geosciences. However, the subject is definitely open and fruitful. Moreover, the approach can specifically be applied to the uncertain parameterizations used in the model [40].

Focusing on the inference of dynamics, there are critical questions that need a clarification such as the use of partial and noisy observations, the required amount of data, or the short term predictive versus asymptotic properties of the retrieved model. These questions will be addressed and discussed in this paper, in the wake of [6, 8].

Incorporating the *surrogate model*, represented here by a set of static coefficients \mathbf{A} , as a control variable in the Bayesian formalism yields [1, 6]

$$p(\mathbf{x}_{0:K}, \mathbf{A} | \mathbf{y}_{0:K}) = \frac{p(\mathbf{y}_{0:K} | \mathbf{A}, \mathbf{x}_{0:K}) p(\mathbf{x}_{0:K} | \mathbf{A}) p(\mathbf{A})}{p(\mathbf{y}_{0:K})}. \quad (3)$$

Consequently, the inference of the model can be seen as a generalized DA problem. The problem is formally close to that of the weak-constraint 4D-Var [43] but where \mathbf{A} is also inferred [6].

In this paper, we will generalize this approach, primarily by incorporating the statistics of model error in the inference. By definition, these statistics are not known a priori and, yet, are expected to have a strong influence on the quality of the estimate of \mathbf{A} .

1.3. Surrogate model representations from machine learning. Considering solutions of (3) requires to build practical and numerically efficient representations of the model. Such representation could be defined by the flow rate of the dynamics:

$$\frac{d\mathbf{x}}{dt} = \phi_{\mathbf{A}}(\mathbf{x}), \quad (4)$$

or by the resolvent, i.e. the integration of the flow rate from one time step t_{k-1} to the next one t_k :

$$\mathbf{x}_k = \mathbf{F}_{\mathbf{A}}(\mathbf{x}_{k-1}). \quad (5)$$

Let us mention broad categories of representations that have been proposed in the literature. The resolvent $\mathbf{F}_{\mathbf{A}}$ could be expanded on a set of nonlinear regressors [33, 9] or it could be given by a neural network (NN) [34, 16, 8, 44], or an echo state network using reservoir modeling [36, 35, 44]. As opposed to the resolvent, unveiling the flow rate (4) may offer an explicit representation of the model, closer to an ordinary (or partial) differential equation system. This could be efficiently achieved through a simple expansion on monomial regressors [6] or using a NN [45, 19, 28]. It has been found that NN architectures for the dynamics that mimic numerical integration schemes, and which take the form of residual NNs, are particularly adequate [17, 13]. Note that most of these papers only consider fully observed systems with none or little noise in the observations.

1.4. Problem and objectives. To address partial and noisy observations, these ML representations can be plugged into an overarching DA formalism as demonstrated in [6]. Hence both DA and ML techniques have to be considered to solve this problem. This is critical for realistic applications.

In this paper, we will generalize this DA framework, incorporating ML representations, but also accounting for model error inherent to the retrieval of an unknown model. The surrogate model, the state trajectory and model error statistics will be estimated all together, which, to the best of our knowledge, has never been achieved before. Moreover, we will unify what has been proposed so far in the literature that aims at learning model dynamics from partial and noisy data.

The main assumptions used in this study are:

1. the governing equations of the model are autonomous, such that the set of parameters \mathbf{A} does not depend on time,
2. the physical system to be learned from is possibly chaotic such that only sequential DA schemes are viable over long training time windows. As in geophysics, the system is spatially extended and possibly high-dimensional,
3. the observations one learns from are often partial and noisy.

For the sake of simplicity, other assumptions of lesser importance will be made. For instance, the observation error statistics will be assumed to be known.

Because model errors are accounted for in the inference, our surrogate model will actually be based on the flow rate

$$\frac{d\mathbf{x}}{dt} = \phi_{\mathbf{A}}(\mathbf{x}) + \boldsymbol{\eta}(t) \quad (6)$$

as a generalization to (4), where $\boldsymbol{\eta}(t)$ is a N_x -dimensional Wiener process. Nonetheless, since in practice the DA system is discretized in time, model error is only injected in between two consecutive observation updates, as is usually done in weak-constraint variational DA, which slightly differs from the continuous stochastic differential equation view offered by (6).

1.5. Outline. In Section 2, we will first describe possible representations for the surrogate dynamics. We will discuss symmetries to be exploited in order for our methods to be scalable to high-dimensional systems. In Section 3, we will define the Bayesian DA framework needed to address the inference of the dynamics, jointly or not with the model state trajectory, while model noise will be accounted for. In Section 4, we will review and discuss the optimization strategies in the case where the model noise statistics are known, while in Section 5 we will propose and discuss the optimization strategies in the case where they are not known. Expectation-maximization and coordinate descent will be the two key techniques for the optimization and they will help unify the approaches proposed so far. The formalism will be illustrated and numerically tested in Section 6.

2. Model representation. Representations required for arbitrary dynamics are likely to have too many degrees of freedom to be tractable in high-dimension, i.e. typically when the flow rate of one variable is a function of all the other variables. However, chaotic geophysical flows have properties and symmetries that could make their representation sparse and practical, despite exhibiting complex behavior.

2.1. Symmetries and constraints. The main hypothesis that we will use is that the governing equations are local. This means that the partial differential equations of the dynamics can be written in terms of point-wise algebraic and differential operators. If the dynamical system is defined in terms of ordinary differential discretized equations, like in this paper, *local* means that the flow rate of a given state variable only depends on state variables defined within the stencil of the state variable. This property is called *locality*. Specific examples will be given below. For high-dimensional systems, this property is critical to avoid more than linear (typically quadratic) dependence in the number of regressors, and more generally degrees of freedom in \mathbf{A} , needed in the representation. Locality is indirectly connected, in many realistic physical systems, to the existence of fast decreasing error correlation in space, which led to the success of ensemble DA localization techniques in these systems.

One can further assume that the flow rate representation is translationally invariant in space. This drastically reduces the number of required regressors in the representation. This property will be called (statistical) *homogeneity*. This property is not as fundamental as locality since most geophysical model are inhomogeneous. However, important physical models such as homogeneous turbulence satisfy this

property. In fact, geofluids usually entail homogeneous fluid dynamics with heterogeneous forcings. Hence an appropriate hybrid representation could involve a homogeneous one for the dynamics and a heterogeneous one for the forcings.

2.2. Sleek representation of the dynamics. A simple representation of the surrogate dynamics is given by $\phi_{\mathbf{A}}(x) = \mathbf{A}\mathbf{r}(\mathbf{x})$, where \mathbf{A} is a matrix of scalar coefficients of size $N_x \times N_p$, and $\mathbf{r}(\mathbf{x})$ is a vector of regressors of size N_p . This has been discussed in detail in [6] and we only summarize the construction of the representation here.

The simplest nonlinear and non-trivial set of regressors is given by the constant, linear and bilinear monomials in the $\{\mathbf{x}_n\}_{0 \leq n \leq N_x-1}$ variables. Note that through the time integration of the flow, higher-order terms are implicitly generated, along with longer distance correlations resulting in non-trivial resolvents for the dynamics.

However, the number of such regressors scales as $N_x^2/2$ which is prohibitive for high-dimensional problems. As discussed above, we can assume locality and demand that the monomials are built with variables whose locations belong to a small stencil of typical size L . The matrix \mathbf{A} then becomes sparse and can be redefined into a dense matrix, which we still denote \mathbf{A} but now of size $N_x \times N_a$. N_a is proportional to L^{2d} up to a geometrical factor, where d is the dimension of the underlying physical space. The linear dependence, instead of quadratic, of the matrix size in N_x makes this representation potentially affordable. If, in addition, homogeneity is assumed, then \mathbf{A} just becomes a vector of size N_a . This yields a minimal but efficient representation of the flow rate.

2.3. Neural network representation of the dynamics. An alternative is to choose a NN representation for the flow rate $\phi_{\mathbf{A}}$. We choose here to mimic the sleek representation described earlier. To do so, we could first (i) build a dense layer DNN_1 from the input $\mathbf{x} \in \mathbb{R}^{N_x}$, then (ii) form the bilinear product from this layer output DNN_1 , then (iii) concatenate this bilinear product with DNN_1 to form an intermediate layer, and (iv) apply to this layer a final dense layer DNN_2 which outputs a flow rate $\phi_{\mathbf{A}}(\mathbf{x})$ in \mathbb{R}^{N_x} , which depends on the weights \mathbf{A} of the DNNs. Not only do we generate linear and bilinear terms but constant terms are also generated with the layers' biases. To mimic the sleek representation, the activation can be chosen linear. But other choices are possible. A batch norm layer can also be applied first to rescale the variability in the NN. Note however that, as explained above, the number of weights should scale like $N_x^2/2$.

When accounting for locality and homogeneity, we use convolutional layers (CNNs) instead of dense layers but follow the same architecture for the whole NN. It is represented in the first row of Figure 1. In this case, the number of weights is proportional to the volume of each convolution in state space and to the number of required convolutions, and has a similar scaling to the sleek homogeneous and local representation.

When accounting for locality only, without homogeneity, we use locally connected layers instead, so that the coefficients of the convolution depend on the location. In this case, the number of weights scales like N_x .

Contrary to [19], we use convolutional layers, so that, owing to locality, we can apply our techniques to high-dimensional systems. We use a simpler NN for the flow rate than in [8], which we will integrate later on, whereas [8] built a NN for the resolvent.

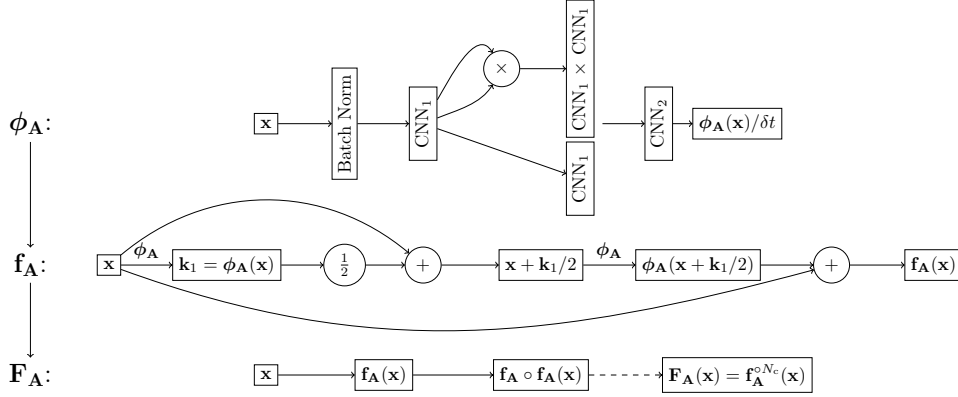


FIGURE 1. From top to bottom: representation of the flow rate $\phi_{\mathbf{A}}$ with a NN, integration of the flow rate into $\mathbf{f}_{\mathbf{A}}$ using an explicit integration scheme (here a second-order Runge Kutta scheme), and N_c -fold composition up to the full resolvent $\mathbf{F}_{\mathbf{A}}$. δt is the integration time step corresponding to the resolvent $\mathbf{f}_{\mathbf{A}}$.

2.4. Time integration and resolvent. We integrate those flow rates to build up a resolvent needed to propagate the state in between observations' times. This integration follows the scheme proposed in [6]. We first use an explicit numerical scheme, typically a fourth-order Runge-Kutta scheme which is based on the flow rate $\phi_{\mathbf{A}}$. This yields the resolvent $\mathbf{f}_{\mathbf{A}}$ over δt . Incorporating a neural network into an integration schema like Runge-Kutta has the property of recurrent neural networks to compose several successive neural blocks with shared weights. A major difference with structures like Long-Short Term Memory (LSTM) or Gated Recurrent Unit (GRU) is that our approach does not have a memory because we rely on the assumption that the dynamics are Markovian.

We then compose this resolvent as many N_c times as required to integrate the surrogate model from t_k to t_{k+1} , assuming $t_{k+1} - t_k$ is a multiple of the integration time step δt . Note that N_c could be made dependent on k if the t_k were not evenly spaced. The full integration scheme, including the NN flow rate, is schematized in Figure 1.

For the sleek representation, its numerical integration, and its composition, we had derived in [6] the adjoint of the tangent linear model needed to formally compute the gradients of a cost function built with the resolvent. With the NN representation, the adjoint of the resolvent needed to compute the gradients are generated by a ML tool library (typically TensorFlow or PyTorch). Note that we are also able to efficiently build the resolvent of the sleek representation using a specific NN representation and a ML library, although this approach is not pursued here.

In the numerical illustrations of Section 6, the NNs, and their resolvent are implemented using Keras 2.2 and TensorFlow 1.5 [14].

3. Bayesian analysis. At this point, we have obtained a complete representation of the resolvent of the surrogate model and its explicit dependence on the set of weights \mathbf{A} . Adopting a Bayesian view on the problem, we can now build the conditional pdf and associated cost function that generalized (3) as introduced in Section 1, by additionally accounting for the model error statistics.

3.1. Prior error statistics. Let us make assumptions on the statistics of the errors, meant to build a tractable cost function. However, other error distributions could be considered. The observation error pdf is assumed Gaussian:

$$p(\mathbf{y}_k | \mathbf{x}_k, \mathbf{R}_k) = \frac{\exp\left(-\frac{1}{2}\|\mathbf{y}_k - \mathbf{H}_k(\mathbf{x}_k)\|_{\mathbf{R}_k^{-1}}^2\right)}{\sqrt{(2\pi)^{N_y} |\mathbf{R}_k|}}, \quad (7)$$

where the vector norm is generically defined by $\|\mathbf{x}\|_{\mathbf{B}}^2 = \mathbf{x}^\top \mathbf{B} \mathbf{x}$. In this paper, the observation error covariance matrices $\mathbf{R}_{0:K} = \{\mathbf{R}_0, \mathbf{R}_1, \dots, \mathbf{R}_K\}$ are supposed to be known. We could incorporate them into the hyperparameters to control, but this would distract us from the main focus which is model error statistics. We also make the assumption of a Gaussian distribution for model error:

$$p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{A}, \mathbf{Q}_k) = \frac{\exp\left(-\frac{1}{2}\|\mathbf{x}_k - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1})\|_{\mathbf{Q}_k^{-1}}^2\right)}{\sqrt{(2\pi)^{N_x} |\mathbf{Q}_k|}}, \quad (8)$$

where $\mathbf{Q}_{1:K} = \{\mathbf{Q}_1, \mathbf{Q}_2, \dots, \mathbf{Q}_K\}$ are not necessarily known. We further assume that these Gaussian errors are white in time and that the observation and model errors are mutually independent.

As explained in [6], there are two types of goals, one that focuses on the joint estimation of \mathbf{A} and $\mathbf{x}_{0:K}$, and the other that focuses on the estimation of \mathbf{A} only.

3.2. Joint estimation of the model, its error statistics and the state trajectory. With the first objective, one looks for the maximum a posterior (MAP) of the joint pdf in \mathbf{A} and $\mathbf{x}_{0:K}$ conditioned on the observations. However, here we wish to estimate model error statistics as well. As a consequence, we focus on the distribution of $\mathbf{A}, \mathbf{Q}_{1:K}, \mathbf{x}_{0:K} | \mathbf{y}_{0:K}, \mathbf{R}_{0:K}$. The generalized conditional pdf to consider in this approach can be expressed using the hierarchy:

$$p(\mathbf{A}, \mathbf{Q}_{1:K}, \mathbf{x}_{0:K} | \mathbf{y}_{0:K}, \mathbf{R}_{0:K}) = \frac{p(\mathbf{y}_{0:K} | \mathbf{x}_{0:K}, \mathbf{R}_{0:K}) p(\mathbf{x}_{0:K} | \mathbf{A}, \mathbf{Q}_{1:K}) p(\mathbf{A}, \mathbf{Q}_{1:K})}{p(\mathbf{y}_{0:K}, \mathbf{R}_{0:K})}, \quad (9)$$

where the mutual independence of the observation and model error was used. The first term in the numerator of the right-hand side is the likelihood of the observations and can be obtained using (7). The second term in the numerator is the prior on the trajectory when the model is known, which can be obtained using the Markov property of the surrogate dynamics and (8), together with the absence of correlation of model error in time. The final term of the numerator is the joint prior of the model and the model error statistics. Here we will not make any prior assumption on \mathbf{A} . However, we will discuss this point to a limited extent in Section 5.4, together with the use of an hyperprior for $\mathbf{Q}_{1:K}$.

As a consequence, looking for the MAP of this conditional pdf, a cost function can be derived:

$$\begin{aligned} \mathcal{J}(\mathbf{A}, \mathbf{x}_{0:K}, \mathbf{Q}_{1:K}) &= -\ln p(\mathbf{A}, \mathbf{Q}_{1:K}, \mathbf{x}_{0:K} | \mathbf{y}_{0:K}, \mathbf{R}_{0:K}) \\ &= \frac{1}{2} \sum_{k=0}^K \left\{ \|\mathbf{y}_k - \mathbf{H}_k(\mathbf{x}_k)\|_{\mathbf{R}_k^{-1}}^2 + \ln |\mathbf{R}_k| \right\} \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{2} \sum_{k=1}^K \left\{ \left\| \mathbf{x}_k - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1}) \right\|_{\mathbf{Q}_k^{-1}}^2 + \ln |\mathbf{Q}_k| \right\} \\
& - \ln p(\mathbf{x}_0, \mathbf{A}, \mathbf{Q}_{1:K}),
\end{aligned} \tag{10}$$

up to an irrelevant constant term. Note the resemblance of (10) with the weak-constraint 4D-Var cost function of classical DA [43].

Very importantly, this Bayesian formulation allows for a rigorous treatment of partial and noisy observations. The classical ML loss function that uses noiseless complete observations of the physical system can be derived from this cost function, assuming that \mathbf{Q}_k is known, $\mathbf{H}_k \equiv \mathbf{I}_x$ and letting \mathbf{R}_k goes to $\mathbf{0}$, so that $\mathcal{J}(\mathbf{A}, \mathbf{x}_{0:K}, \mathbf{Q}_{1:K})$ becomes in this limit

$$\mathcal{J}(\mathbf{A}) = \frac{1}{2} \sum_{k=1}^K \left\| \mathbf{y}_k - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{y}_{k-1}) \right\|_{\mathbf{Q}_k^{-1}}^2 - \ln p(\mathbf{y}_0, \mathbf{A}, \mathbf{Q}_{1:K}). \tag{11}$$

This view of the connections and similarities between DA and ML is similar to and extends to broader configurations than the DA and ML connections put forward in [22, 1, 6].

3.3. Joint estimation of the model and the error statistics. The second approach with a different objective is to obtain a MAP for the surrogate model, irrespective of any model state realization, i.e. we are interested in the MAP of the marginal conditional pdf

$$p(\mathbf{A}, \mathbf{Q}_{1:K} | \mathbf{y}_{0:K}, \mathbf{R}_{0:K}) = \int d\mathbf{x}_{0:K} p(\mathbf{A}, \mathbf{Q}_{1:K}, \mathbf{x}_{0:K} | \mathbf{y}_{0:K}, \mathbf{R}_{0:K}), \tag{12}$$

which is theoretically obtained by minimizing

$$\mathcal{J}(\mathbf{A}, \mathbf{Q}_{1:K}) = -\ln p(\mathbf{A}, \mathbf{Q}_{1:K} | \mathbf{y}_{0:K}, \mathbf{R}_{0:K}). \tag{13}$$

As pointed out in [6], the marginal pdf (12) can be approximately related to the joint pdf through a Laplace approximation of the integral. Here, however, we are interested in the full solution to this problem. The main optimization technique described in this paper is designed to solve this marginal problem, and the numerical experiments will use this approach.

4. Optimization with known model error statistics. We now focus on efficient numerical optimization techniques to obtain the MAP of either (10) or (13). In this section, we specifically focus on the joint estimation of \mathbf{A} and $\mathbf{x}_{0:K}$, i.e. on the cost function (10) assuming that we know $\mathbf{Q}_{1:K}$, i.e. a cost function of the form $\mathcal{J}(\mathbf{A}, \mathbf{x}_{0:K})$ is minimized. We unify the approaches advocated in [6, 8, 32] by seeing them as many approximate optimization techniques for the same Bayesian problem. We discuss the pros and cons of each approach.

4.1. Joint variational solution. In [6], (10) is jointly minimized on \mathbf{A} and on $\mathbf{x}_{0:K}$, considering $\mathbf{Q}_{1:K}$ to be known. The quasi-Newton L-BFGS technique [10] is chosen to implement the minimization. The technique is very fast and efficient, provided the training window is not too long. Indeed, the technique is well suited for nonlinear optimization problems. However, when the training window increases, i.e. when K gets larger, the approximation of the inverse Hessian formed by the L-BFGS algorithm becomes a large matrix, which heavily weights on the computational efficiency. However, if we assume ergodicity of the model state trajectory, we can

assume that distant in time pieces of the trajectory are independent. This was not exploited in [6] and this explains why it becomes impractical as K becomes large.

4.2. Coordinate descent. More generally, because of the difference in the nature of the variables \mathbf{A} and $\mathbf{x}_{0:K}$, we advocate a coordinate descent [48], which alternates minimization on \mathbf{A} and $\mathbf{x}_{0:K}$. This is likely to become numerically efficient as K increases.

Here, we focus on optimization strategies that use coordinate descent making explicit the differences in the \mathbf{A} and $\mathbf{x}_{0:K}$ variables. The key idea put forward in [8], is that the minimization on \mathbf{A} can be addressed using ML techniques, while the minimization on $\mathbf{x}_{0:K}$ can be carried out using classical DA techniques. Note however, that the framework is that of DA. Using ML for the minimization over \mathbf{A} merely means that we rely on ML libraries to solve a variational problem because we leverage the fact that the surrogate model is a NN that can conveniently be implemented with those libraries. Hence the use of ML is here more technical than conceptual. Let us discuss in the following subsections of the implementation of this coordinate descent optimization.

4.3. Minimization over \mathbf{A} . The minimization of $\mathcal{J}(\mathbf{A}, \mathbf{x}_{0:K})$ over \mathbf{A} , considering $\mathbf{x}_{0:K}$ fixed, is a variational subproblem. Because the resolvent will be implemented in Section 6 with Keras/Tensorflow, we could use stochastic gradient descent techniques, since they are implemented in these software tools. This is what was used in [8]. We also tested this approach in the context of this paper. However, as expected with stochastic gradient descent methods, we found the minimization to exhibit a slow convergence close to the solution. This is fine in ML applications since accuracy is usually not an issue, or adjusted accuracy is a means to enforce regularization. Here, however, the reconstruction of the flow rate of the dynamics and its quality are sensitive to the accuracy. We found it to be more efficient to implement an L-BFGS minimizer but using the gradients generated by Keras/TensorFlow. In particular, the convergence is faster and a better accuracy is ultimately and systematically achieved. All the numerical experiments in Section 6 use this scheme for the minimization over \mathbf{A} .

4.4. Minimization over $\mathbf{x}_{0:K}$. The minimization of $\mathcal{J}(\mathbf{A}, \mathbf{x}_{0:K})$ over $\mathbf{x}_{0:K}$, considering \mathbf{A} , i.e. the surrogate model, to be fixed, is a classical DA problem. Because the model is known up to some model noise, it is actually a weak-constraint data assimilation problem [43]. It can be approximately solved using weak-constraint 4D-Var optimization techniques or using sequential schemes such as an ensemble Kalman smoother that accounts for model errors.

4.4.1. With variational subproblems. As explained above, using a global quasi-Newton method is prohibitive for long training window $[t_0, \dots, t_K]$. There are however techniques that have been proposed to parallelize the problem in time [20, 39]. We have actually tested numerically a variant where L-BFGS is applied in parallel over subwindows of $[t_0, \dots, t_K]$ with possible overlapping. This enables training over long time windows. However, we do not wish to emphasize this approach here since it does not allow to account for uncertain model error statistics, without modifications, while ensemble techniques do.

4.4.2. *With ensemble data assimilation.* An alternative is to use an ensemble Kalman filter or smoother to estimate $\mathbf{x}_{0:K}$. In [8], an ensemble Kalman filter (EnKF, [18]) is chosen. It accounts for sampling and model error through adaptive inflation [5] and additive stochastic model noise. The solution of this subproblem is the posterior mean trajectory of the EnKF. Firstly, however, as opposed to what will be done in Section 5, the posterior ensemble is later only used through its mean. Secondly, to improve the accuracy of the trajectory estimate, one can use a smoother in place of the filter. Thirdly, to account for model error one can use a more accurate scheme than stochastic perturbations. These aspects will be discussed in Section 5.

With the coordinate descent approach, we assumed that the model and observation error statistics are known. The DA techniques used to find $\mathbf{x}_{0:K}$ in the coordinate descent could be supplemented by estimation techniques for $\mathbf{Q}_{1:K}$ and/or $\mathbf{R}_{0:K}$, which have been recently reviewed in [42]. However, the main focus of this paper being on the marginal problem (which implicitly involves the state trajectory), this key issue is addressed differently, as described in the following Section.

5. Full joint optimization with expectation-maximization. In this section, we focus on the joint estimation of \mathbf{A} and $\mathbf{Q}_{1:K}$, knowing the observations $\mathbf{y}_{0:K}$ and the statistics of their errors $\mathbf{R}_{0:K}$. Hence, we look for the MAP of (12), i.e. a minimum of (13). Hereafter, we assume $\mathbf{Q}_k \equiv \mathbf{Q}$, i.e. that the model error statistics do not depend on time. Estimating \mathbf{Q}_k at each time step would yield a severely unconstrained estimation problem. This assumption can also be seen as a direct consequence of the hypothesis of autonomous dynamics (flawed or not), see Section 1.4. In the following, \mathbf{Q} will either be assumed non-parameterized (i.e. full \mathbf{Q}) or proportional to the identity matrix $\mathbf{Q} = q\mathbf{I}_x$. In the absence of statistical homogeneity, a diagonal \mathbf{Q} would also be a convenient parameterization [27], but is not tested here. It has been suggested by [21], in the context where both model error statistics and a model are to be identified, that the MAP of pdfs similar to (12) could be found using the expectation-maximization (EM) technique. More recently [37] have used it in a DA problem to find model and observation error statistics in the case where model error is parameterized and stochastic.

5.1. Principle of expectation-maximization. The EM method [15] is ideally suited to find the maximum likelihood or MAP of a marginal pdf, over some parameters. The principle of the method is as follows. We are interested in a local maximum of the conditional pdf depending on a vector $\boldsymbol{\theta}$ of parameters:

$$p(\boldsymbol{\theta}|\mathbf{y}) = \frac{p(\mathbf{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{y})} = \frac{p(\boldsymbol{\theta})}{p(\mathbf{y})} \int d\mathbf{x} p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) p(\mathbf{x}|\boldsymbol{\theta}), \quad (14)$$

where an expression for $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$ and for $p(\mathbf{x}|\boldsymbol{\theta})$ is known whereas the integral being intractable, an analytic expression for $p(\mathbf{y}|\boldsymbol{\theta})$ is not known. Usually, the method focus on a likelihood $p(\mathbf{y}|\boldsymbol{\theta})$ whereas here, as in [27], we are interested in applying EM to the conditional pdf $p(\boldsymbol{\theta}|\mathbf{y})$. As seen in (14), the only difference is in the hyperprior $p(\boldsymbol{\theta})$. The classical approach implicitly chooses it to be $p(\boldsymbol{\theta}) \equiv 1$. We will see in Section 5.4 that other choices are possible.

The EM algorithm, meant to find the maximum of $p(\boldsymbol{\theta}|\mathbf{y})$, is iterative and alternates the following steps until convergence to a local maximum of $\boldsymbol{\theta}$.

5.1.1. *The expectation step:* Given the iterate $\boldsymbol{\theta}^{(j)}$ of the parameters at the iteration j , one computes the function $\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)}) = \mathbb{E}_{\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}^{(j)}} [\ln p(\mathbf{x}, \mathbf{y}, \boldsymbol{\theta})]$, where $\mathbb{E}_{\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}^{(j)}}$ is

the expectation operator over the distribution of $\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}^{(j)}$. Note that $p(\mathbf{x}, \mathbf{y}, \boldsymbol{\theta}) = p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta})$ corresponds to the whole integrand in (14) together with the hyperprior for $\boldsymbol{\theta}$.

5.1.2. The maximization step: Once $\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)})$ has been computed, one can look for one of its local maxima and set it to be the new vector of parameters: $\boldsymbol{\theta}^{(j+1)} = \operatorname{argmax}_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)})$.

Iterating, the EM method converges to a local maximum provided there is at least one. An analytical expression is sometimes known for \mathcal{L} , for example when looking for the optimal mixing coefficients of a Gaussian mixture model [3]. Unfortunately, there is no known expression in our case. An approximation meant to compute $\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)})$ estimates the expectation from a Monte Carlo sampling [46]. If there are N_e such samples, an estimator for \mathcal{L} is

$$\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)}) \approx \frac{1}{N_e} \sum_{i=1}^{N_e} \ln p(\mathbf{x}_i^{(j)}, \mathbf{y}, \boldsymbol{\theta}), \quad (15)$$

where $\mathbf{x}_i^{(j)}$ is a sample of $\mathbf{x}|\mathbf{y}, \boldsymbol{\theta}^{(j)}$.

Alternatively, one can look for a local minimum of $-\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)})$. This is the convention adopted in the following and we redefine $-\mathcal{L}$ as \mathcal{L} , and deal with minima instead of maxima.

5.2. Full scheme. Hence, getting the MAP of (12) can in principle be numerically solved using an EM-like algorithm. In our problem, $\boldsymbol{\theta}$ is the set of parameters (\mathbf{A}, \mathbf{Q}) , while \mathbf{y} is $\mathbf{y}_{0:K}$ and the marginal integration is over all possible trajectories $\mathbf{x}_{0:K}$.

We choose for the initial $\mathbf{A}^{(0)}$ a small set of random coefficients. We specifically want these to be small so that $\mathbf{A}^{(0)}$ lies not too far from the domain of stability of the generated dynamics; see Section 3.2 of [6] for a discussion. We choose $\mathbf{Q}^{(0)}$ to be the matrix $q^{(0)} \mathbf{I}_x$ where $q^{(0)}$ will be chosen large enough so as not to underestimate model error and avoid a premature divergence in the ensemble DA step. For sequential DA schemes, an initial $\mathbf{x}^{(0)}$ also has to be chosen. The final iterates should not depend on these initial values except for picking the local minimum towards which the algorithm converges.

5.2.1. The expectation step: We assume \mathbf{A}, \mathbf{Q} to be given by their estimates at iteration j , i.e. $\mathbf{A}^{(j)}, \mathbf{Q}^{(j)}$. The expectation step requires an ensemble DA technique, which later enables the use of Monte Carlo EM, the ensemble providing the samples. Specifically, we will rely in the numerical experiments of Section 6 on the classical ensemble Kalman smoother (EnKS, [18]) based on the ensemble transform formulation [4, 23]. The smoother has two steps: the first one is a filtering pass; the second one is the smoother pass which traces back L time steps in the past, where L is known as the lag of the smoother. The model used in the EnKS is by construction the surrogate model specified by $\mathbf{A}^{(j)}$.

In order not to be annoyed by difficult technical details of the implementation which are not key to the method presented here, we will assume that the size of the ensemble N_e is large enough so that (i) localization is not necessary (ii) model error is well represented by an ensemble with a large span. For instance if $N_e - 1$ is chosen to be the dimension of the surrogate model state vector, one can expect on the basis of the experiments presented in [38] that the posterior error is minimal. In

realistic, high-dimensional problems where only a small ensemble can be afforded, localization would have to be used and would be efficient as demonstrated in [41]. Model error is accounted for in the filtering pass of the EnKS using the deterministic SQRT-CORE scheme [38] using $\mathbf{Q}^{(j)}$. With a large ensemble size, we expect it to be more accurate than the addition of stochastic noise as in [8]. The output of this step is the smoothed ensemble trajectory $\mathbf{E}_{0:K}^{(j)} \in \mathbb{R}^{K \times N_x \times N_e}$ as computed by the EnKS.

5.2.2. *The maximization step:* At iteration j , we assume the ensemble trajectory

$$\mathbf{E}_{0:K}^{(j)} = \left\{ \mathbf{x}_{i,k}^{(j)} \right\}_{1 \leq i \leq N_e, 0 \leq k \leq K} \quad (16)$$

to be given by the expectation/DA step. Following the Monte Carlo EM algorithm, we would like to minimize

$$\begin{aligned} \mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q}) &= -\frac{1}{N_e} \sum_{i=1}^{N_e} \ln p(\mathbf{E}_i^{(j)}, \mathbf{y}_{0:K}, \mathbf{A}, \mathbf{Q}, \mathbf{R}_{0:K}) \\ &= \frac{1}{2N_e} \sum_{i=1}^{N_e} \sum_{k=1}^K \left\{ \left\| \mathbf{x}_{k,i}^{(j)} - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1,i}^{(j)}) \right\|_{\mathbf{Q}^{-1}}^2 + \ln |\mathbf{Q}| \right\} \\ &\quad - \ln p(\mathbf{x}_{0,i}^{(j)}, \mathbf{A}, \mathbf{Q}) + \dots, \end{aligned} \quad (17)$$

up to an irrelevant constant term.

5.2.3. *Joint optimization of \mathbf{A} and \mathbf{Q} .* A straightforward solution of the maximization step consists in optimizing $\mathcal{L}(\mathbf{A}, \mathbf{Q})$ on both \mathbf{A} and \mathbf{Q} using L-BFGS of [10] and the gradients generated by an ML library such as Keras/Tensorflow. Note that this requires to store the full ensemble of trajectories over the training window.

The full scheme algorithm is summarized in Algorithm 1. The criteria for convergence can be rather complex, having to monitor both \mathbf{A} and \mathbf{Q} . For the sake of simplicity, a fixed number of iterations N_{iter} will be imposed in Section 6.

Algorithm 1 Full scheme Monte Carlo EM algorithm for estimating \mathbf{A} and \mathbf{Q}

- 1: $j = 0$
 - 2: Set $\mathbf{A}^{(0)}$ and $\mathbf{Q}^{(0)}$
 - 3: **while** ($\mathbf{A}^{(j)}$ or $\mathbf{Q}^{(j)}$ are not converged) and $j < N_{\text{iter}}$ **do**
 - 4: Expectation/DA step: EnKS with SQRT-CORE using $\mathbf{A}^{(j)}, \mathbf{Q}^{(j)}$
 Output: $\mathbf{E}_{0:K}^{(j)}$
 - 5: Maximization/ML step: Minimize $\mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q})$ knowing $\mathbf{E}_{0:K}^{(j)}$
 Output: $\mathbf{A}^{(j+1)}$ and $\mathbf{Q}^{(j+1)}$
 - 6: $j \leftarrow j + 1$
 - 7: **end while**
-

We tested this option of jointly optimizing on \mathbf{A} and \mathbf{Q} with success. However, the minimization of $\mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q})$ over \mathbf{Q} alone has an analytic expression which can efficiently be leveraged upon. To do so, one could use a coordinate descent on \mathbf{A} and \mathbf{Q} . At iteration j of the EM algorithm, one first minimizes $\mathcal{L}(\mathbf{A}, \mathbf{Q}^{(j)})$ using

$\mathbf{A}^{(j)}$ as starting point of the minimization. Note that this cost function is essentially the classical least squares loss:

$$\mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q}^{(j)}) = \frac{1}{2N_e} \sum_{i=1}^{N_e} \sum_{k=1}^K \left\| \mathbf{x}_{k,i}^{(j)} - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1,i}^{(j)}) \right\|_{\mathbf{Q}^{(j)-1}}^2, \quad (18)$$

where the prior on \mathbf{A} is neglected. The local minimum which is obtained is noted $\mathbf{A}^{(j,1)}$. Then one can solve for the exact minimum of $\mathcal{L}(\mathbf{A}^{(j,1)}, \mathbf{Q})$ over the non-parameterized \mathbf{Q} :

$$\mathbf{Q}^{(j,1)} = \frac{1}{KN_e} \sum_{i=1}^{N_e} \sum_{k=1}^K \left(\mathbf{x}_{k,i}^{(j)} - \mathbf{F}_{\mathbf{A}^{(j,1)}}^{k-1}(\mathbf{x}_{k-1,i}^{(j)}) \right) \left(\mathbf{x}_{k,i}^{(j)} - \mathbf{F}_{\mathbf{A}^{(j,1)}}^{k-1}(\mathbf{x}_{k-1,i}^{(j)}) \right)^\top. \quad (19)$$

Similar expressions can be obtained in the case where \mathbf{Q} is diagonal or of the specific form $q\mathbf{I}_x$ (see e.g., [27]). With this new $\mathbf{Q}^{(j,1)}$, one can go back to the minimization on \mathbf{A} with the initial guess being $\mathbf{A}^{(j,1)}$ and iterate the coordinate descent until convergence. This ultimately yields $\mathbf{A}^{(j+1)}$ and $\mathbf{Q}^{(j+1)}$.

However, we do not expect the number of iterations of this coordinate descent to be large. Indeed, $\mathcal{L}(\mathbf{A}, \mathbf{Q}^{(j)})$ as seen in (18) should not be very sensitive to $\mathbf{Q}^{(j)}$ once $\mathbf{Q}^{(j,\star)}$ is close to convergence because it only plays the role of a Mahalanobis metric in the inner product. To support this argument, note that in the case where $\mathbf{Q} = q\mathbf{I}_x$, the solution of the minimization of $\mathcal{L}(\mathbf{A}, \mathbf{Q}^{(j)})$ does not actually depends on \mathbf{Q} at all, since it only represents a scaling factor of the loss function. In this case, at iteration j of the EM algorithm, the coordinate descent simplifies to a single minimization of $\mathcal{L}(\mathbf{A}, \mathbf{Q}^{(j)})$ on \mathbf{A} which yields $\mathbf{A}^{(j+1)}$, followed by (19) which yields $\mathbf{Q}^{(j+1)}$.

5.3. Approximate scheme. We now present an approximation of the maximization step, which avoids to store the full ensemble of smoothed trajectories over the training window (as required by (17) or (18)).

At iteration j of the EM algorithm, we first implement the EnKS with SQR-CORE using $\mathbf{A}^{(j)}$ and $\mathbf{Q}^{(j)}$. Because we want to avoid computing $\mathbf{Q}^{(j+1)}$ at the maximization step through

$$\mathbf{Q}^{(j+1)} = \frac{1}{KN_e} \sum_{i=1}^{N_e} \sum_{k=1}^K \left(\mathbf{x}_{k,i}^{(j)} - \mathbf{F}_{\mathbf{A}^{(j+1)}}^{k-1}(\mathbf{x}_{k-1,i}^{(j)}) \right) \left(\mathbf{x}_{k,i}^{(j)} - \mathbf{F}_{\mathbf{A}^{(j+1)}}^{k-1}(\mathbf{x}_{k-1,i}^{(j)}) \right)^\top, \quad (20)$$

which requires the full ensemble, we rather compute it online in the expectation step by accumulating terms of the sum as the EnKS unfolds. Moreover, instead of extracting the full ensemble from the EnKS, we only store and use the smoothed ensemble trajectory mean $\bar{\mathbf{x}}_{0:K}^{(j)}$.

Still at iteration j of the EM algorithm, focusing on the maximization step, we carry out the minimization of

$$\begin{aligned} \mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q}^{(j+1)}) &= -\ln p(\bar{\mathbf{x}}_{0:K}^{(j)}, \mathbf{y}_{0:K}, \mathbf{A}, \mathbf{Q}^{(j+1)}, \mathbf{R}_{0:K}) \\ &= \frac{1}{2} \sum_{k=1}^K \left\{ \left\| \bar{\mathbf{x}}_k^{(j)} - \mathbf{F}_{\mathbf{A}}^{k-1}(\bar{\mathbf{x}}_{k-1}^{(j)}) \right\|_{\mathbf{Q}^{(j+1)-1}}^2 + \ln |\mathbf{Q}^{(j+1)}| \right\} \\ &\quad - \ln p(\bar{\mathbf{x}}_0^{(j)}, \mathbf{A}, \mathbf{Q}^{(j+1)}) + \dots, \end{aligned} \quad (21)$$

with respect to \mathbf{A} , up to terms not depending on \mathbf{A} . The result is a local minimum: $\mathbf{A}^{(j+1)} = \operatorname{argmin}_{\mathbf{A}} \mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q}^{(j+1)})$. We cannot iterate the minimization on \mathbf{A} and

\mathbf{Q} as the $\mathbf{Q}^{(j+1)}$ has been computed once and for all in the expectation step, which is a significant approximation of the coordinate descent. However, as argued in Section 5.2.3, this approximation should be mild since, if $\mathbf{Q} = q\mathbf{I}_x$, the coordinate descent would not require iterating.

The approximate algorithm is summarized in Algorithm 2.

Algorithm 2 Approximate Monte Carlo EM algorithm for estimating \mathbf{A} and \mathbf{Q}

```

1:  $j = 0$ 
2: Set  $\mathbf{A}^{(0)}$  and  $\mathbf{Q}^{(0)}$ 
3: while ( $\mathbf{A}^{(j)}$  or  $\mathbf{Q}^{(j)}$  are not converged) and  $j < N_{\text{iter}}$  do
4:   Expectation/DA step: EnKS with SQRT-CORE using  $\mathbf{A}^{(j)}, \mathbf{Q}^{(j)}$ 
     Output:  $\bar{\mathbf{x}}_{0:K}^{(j)}$  and  $\mathbf{Q}^{(j+1)}$  computed online
5:   Maximization/ML step: Minimize  $\mathcal{L}^{(j)}(\mathbf{A}, \mathbf{Q}^{(j+1)})$  knowing  $\bar{\mathbf{x}}_{0:K}^{(j)}$ 
     Output:  $\mathbf{A}^{(j+1)}$ 
6:    $j \leftarrow j + 1$ 
7: end while

```

The authors of [32] have nicely suggested to use the EM algorithm to learn a model from data following [21], in particular using an EnKS for the expectation step. However, in their algorithms and applications, they (i) do not estimate the model error noise, (ii) do not make use of the ensemble output of the EnKS in the maximization step, which is key to EM. As a result, even though their algorithms have the structure of EM, their algorithms are merely coordinate descents applied to $\mathcal{J}(\mathbf{A}, \mathbf{x}_{0:K})$, which amounts to the algorithm proposed in [8].

5.4. Hyperpriors.

5.4.1. *Hyperprior on \mathbf{Q} .* An hyperprior for \mathbf{Q} is allowed by the general formalism (14). We expect such an hyperprior to be impactful only in the case of short training windows; it should be overridden by increasing data. If \mathbf{Q} is parameterized as $q\mathbf{I}_x$, we can use a scalar Jeffreys hyperprior on the variance, which is $p(q) = q^{-1}$. Focusing only on the occurrences of q , the main cost function is

$$\mathcal{L}(\mathbf{A}, \mathbf{Q}) = -\ln p(q) + \frac{Ks}{2q} + \frac{KN_x}{2} \ln(q) + \dots \quad (22)$$

where

$$s = \frac{1}{KN_e} \sum_{i=1}^{N_e} \sum_{k=1}^K \|\mathbf{x}_{k,i} - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1,i})\|^2. \quad (23)$$

Minimizing on q , this yields for the maximization step:

$$q = \frac{K}{KN_x + 2} s. \quad (24)$$

Clearly, the influence of the hyperprior on the estimate of q (+2 in the denominator) is minor, especially for large system size or training window.

If \mathbf{Q} is full, we can use a matrix Jeffreys hyperprior $p(\mathbf{Q}) = |\mathbf{Q}|^{-\frac{N_x+1}{2}}$. Focusing only on the occurrences of \mathbf{Q} , the main cost function is

$$\mathcal{L}(\mathbf{A}, \mathbf{Q}) = -\ln p(\mathbf{Q}) + \frac{K}{2} \text{Tr}(\mathbf{S}\mathbf{Q}^{-1}) + \frac{K}{2} \ln(|\mathbf{Q}|) + \dots \quad (25)$$

where

$$\mathbf{S} = \frac{1}{KN_e} \sum_{i=1}^{N_e} \sum_{k=1}^K (\mathbf{x}_{k,i} - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1,i})) (\mathbf{x}_{k,i} - \mathbf{F}_{\mathbf{A}}^{k-1}(\mathbf{x}_{k-1,i}))^\top. \quad (26)$$

Minimizing on \mathbf{Q} , this yields for the maximization step:

$$\mathbf{Q} = \frac{K}{K + N_x + 1} \mathbf{S}. \quad (27)$$

The influence of the hyperprior on the estimate of \mathbf{Q} can be stronger in this case, in particular when the training window length is smaller or of the order of the state space dimension.

5.4.2. Hyperprior on \mathbf{A} . The design of the hyperprior for \mathbf{A} is primarily driven by physical modeling and numerical stability. Its role has been discussed to some extent in Section 3.1 of [6]. Practically, such an hyperprior could be implemented by adding a regularization term (typically L1 or L2 norm) on the coefficients of \mathbf{A} , corresponding to specific distribution assumptions for \mathbf{A} . We avoid such regularization here by mostly considering very long training windows, and because \mathbf{A} is rather well constrained by locality and/or homogeneity. However, with higher dimensional physical models, larger \mathbf{A} , deeper NN representations, and shorter training windows by comparison, methods used in machine learning and deep learning to regularize and avoid overfitting could be used, for instance dropouts and stochastic optimization techniques (e.g., Section 1.4 of [25]).

6. Numerical illustrations. In this section, we test the algorithms introduced in Section 5. The algorithms reinterpreted as coordinate descent in Section 5, and for which model error statistics were prescribed have been tested in [6, 8].

6.1. Two low-order reference models. The new methods are tested on two low-order models: the Lorenz-96 model (L96, [30]) and the two-scale Lorenz model, sometimes known as Lorenz-05III (L05III, [29]). Both are spatially distributed one-dimensional models defined over circles. In the following experiments, they are called the *reference* models.

6.1.1. Lorenz-96. The L96 model is defined by a set of ODEs over a periodic domain with variables indexed by $n = 0, \dots, N_x - 1$:

$$\frac{dx_n}{dt} = (x_{n+1} - x_{n-2})x_{n-1} - x_n + F, \quad (28)$$

where $N_x = 40$, $x_{N_x} = x_0$, $x_{-1} = x_{N_x-1}$, $x_{-2} = x_{N_x-2}$, and $F = 8$. This model is an idealized representation of a one-dimensional latitude band of the Earth atmosphere. Its Lyapunov time (inverse of the largest Lyapunov exponent) is 0.60. The standard deviation of any model state variable is 3.62. The truth run of the L96 model is integrated using the fourth-order Runge Kutta (RK4) scheme and with the time step $\delta t_r = 0.05$.

A classical DA configuration consists in the observation of every variable of the model every $\Delta t_r = \delta t_r = 0.05$. The observations are perturbed with random white in time normal noise of covariance matrix $\mathbf{R} = \sigma_y^2 \mathbf{I}_x$. In this classical configuration, we have $\sigma_y = 1$.

6.1.2. *Lorenz-05III*. The two-scale Lorenz model L05III is given by the following two-scale set of ODEs:

$$\frac{dx_n}{dt} = \psi_n^+(\mathbf{x}) + F - h \frac{c}{b} \sum_{m=0}^9 u_{m+10n}, \quad (29a)$$

$$\frac{du_m}{dt} = \frac{c}{b} \psi_m^-(b\mathbf{u}) + h \frac{c}{b} x_{m/10}, \quad (29b)$$

$$\psi_n^\pm(\mathbf{x}) = x_{n\mp 1}(x_{\pm 1} - x_{n\mp 2}) - u_n, \quad (29c)$$

for $n = 0, \dots, N_x - 1$ with $N_x = 36$, and $m = 0, \dots, N_u - 1$ with $N_u = 360$. The indices are defined periodically over their domain; $m/10$ stands for the integer division of m by 10. The other parameters: $c = 10$ for the time-scale ratio, $b = 10$ for the space-scale ratio, $h = 1$ for the coupling, and $F = 10$ for the forcing, are set to their original values. When uncoupled ($h = 0$), the Lyapunov time of the slow variables \mathbf{x} sector of the model (29) is 0.72, which is the key time scale when focusing on the slow variables [11]. The standard deviation of any of the slow variables is 3.54.

The vector \mathbf{u} represents unresolved scales and hence model error when only considering the slow variables \mathbf{x} . It is integrated with an RK4 scheme and the time step $\delta t_r = 0.005$ since it is stiffer than the L96 model.

A classical DA configuration is defined by the observation of every slow variable of the model, with a time step of $\Delta t_r = 10 \delta t_r = 0.05$. The observations are perturbed with random white in time normal noise of covariance matrix $\mathbf{R} = \sigma_y^2 \mathbf{I}_x$. In this classical configuration, we have $\sigma_y = 1$.

The goal is to learn surrogate dynamics from observing these models. Note that in the L05III case, only the slow variables are observed and only them are represented in the surrogate model and forecasted. Given that only its slow scale is observed and reconstructed, the L05III reference model is not identifiable in that it does not belong to the space generated by all possible surrogate models. It is vain to attempt to exactly match the ODEs description of the reference models. Although less obvious, the L96 model may neither be identifiable by the NN representation. Indeed, we cannot obtain a perfect numerical reconstruction of the dynamics to numerical precision through the NN representation, as opposed to when using the sleek representation [6].

6.2. Surrogate model. For both reference models, we will use a surrogate model with a NN architecture based on the RK4 scheme. Its architecture is the same as the one depicted in Figure 1, except for RK2 replaced by RK4. We choose $N_c = 1$, for the sake of simplicity and because the accuracy improvement offered by $N_c > 1$, which has been tested, is often small.

The CNN_1 layer is a convolutional layer with a width of 5 sites, and comprises 6 filters (or *channels*). We have checked that a larger width or more filters does not lead to noticeable improvements. This is consistent with the theory developed in [6] which emphasizes that only a few key parameters are required (ultimately those of the ODEs of the true model). For CNN_1 , we use linear activation functions. We have also tested relu and tanh activation functions, which turned out to have no significant impact or can degrade the accuracy of the surrogate model. All occurrences of the CNN_1 layer within the NN share the same weights.

The CNN_2 layer is a point-wise convolutional layer with one filter that outputs a state vector. It always has linear activation functions. Note that in the following experiments, batch normalization is not used.

This minimal NN (i.e. assuming locality and homogeneity) only has 49 weights (for both reference models). With locally connected convolutional layers, hence assuming locality but breaking translational invariance, it has 1960 and 1764 weights for the L96 and L05III experiments, respectively.

6.3. Evaluation metrics. Once the surrogate models are learned from the observations, their quality must be evaluated considering both the short-term and asymptotic properties of the dynamics.

Our evaluation metrics for the numerical experiments will be:

- Forecast skill (FS): the normalized root mean square error (NRMSE) is the root mean square difference between a forecast of the reference model and a forecast of the surrogate model starting from the same initial condition (up to machine precision), and normalized by the standard deviation of the reference model variables. This score depends on the lead time. The FS is averaged over 5×10^3 model runs with as many independent initial conditions (not taken from the training window) to obtain a satisfactory convergence. Above a NRMSE of 1, the forecast of the surrogate model has become useless.
- Lyapunov spectrum (LS): It measures the asymptotic properties of the model and is computed numerically for both the reference and surrogate model. The spectrum is computed over 10^4 time steps, using either finite differences or a Jacobian indirectly obtained from Keras/TensorFlow.
- Power spectrum density (PSD): It measures the energy content of the dynamics in spectral space, as a function of the wave frequency. The PSD is computed as an average on all grid points using the Welch method [47].

Hovmöller diagrams can also be exploited to visualize differences of trajectories of the reference and the surrogate models. Some of them have been displayed in [6, 8].

Several scalar indicators will also be used. The average validation prediction time when the NRMSE first reaches 0.5 will be denoted $\pi_{\frac{1}{2}}$ (in units of Lyapunov time). The average diagnosed residual model error will be measured by

$$\sigma_q = \sqrt{\frac{1}{N_x} \text{tr}(\mathbf{Q}^*)}, \quad (30)$$

where \mathbf{Q}^* is the model error covariance matrix estimated in the training phase. The average first Lyapunov exponent will be noted λ_1 .

6.4. First results.

6.4.1. Preliminary screening of hyperparameters. We have carried out a series of experiments with the parameters described above, but taking the EnKS lag in the set $L = \{0, 2, 4, 6, 8, 10, 12\}$ and the training window length in the set $K = \{50, 100, 200, 400, 800, 1600, 3200, 6400\}$. From these experiments, we picked the *nominal* configuration $L = 4, K = 5000$ as a configuration with satisfactory forecast skills, i.e. which maximizes $\pi_{\frac{1}{2}}$. Yet, such a criterion does not guarantee that the asymptotic properties indicators are optimal as well.

The fact that $L = 4$, a rather small value, is close to optimal may be due to the balance between (i) a large L which in principle offers samples that are more representative of the conditional pdf, and (ii) a degradation of the smoother due to

Model	N_y	σ_y	K	L	$\pi_{\frac{1}{2}}$	σ_q	λ_1
L96	40	1	5000	4	4.56 ± 0.06	$0.08790 \pm 2 \cdot 10^{-5}$	1.66 ± 0.02
L05III	36	1	5000	4	4.06 ± 0.21	$0.07720 \pm 2 \cdot 10^{-5}$	1.03 ± 0.05

TABLE 1. Scalar indicators for nominal experiments based on L96 and L05III. Key hyperparameters are recalled. The statistics of the indicators are obtained over 10 samples.

the presence of model noise as L gets larger and (iii) the fact we used the classical smoother and not iterative, more optimal ones.

Some of the indicators may display variability when changing the seed of the random generator used to compute the reference model trajectory and to perturb the observations. This point will be investigated in the following for each of the three indicators by running 10 experiments with a new trajectory of the reference model each, and hence new observations, due to the choice of a new seed.

As mentioned before, the stopping criterion of the EM algorithm will simply be the number of iterations chosen to be $N_{\text{iter}} = 25$. In all the numerical experiments, we observe that the sequence $\mathbf{Q}^{(j)}$ usually converges in less than 10 iterations, but that the model weights sequence $\mathbf{A}^{(j)}$ keeps improving a little bit more beyond this point; hence the safe choice $N_{\text{iter}} = 25$.

We have first compared the results of configurations using either $\mathbf{Q} = q\mathbf{I}_x$ or the full \mathbf{Q} . We found a slight improvement due to the full \mathbf{Q} representation. Since the overhead computational cost due to the use of \mathbf{Q} is mild for our experiments, we chose it for all the following reported experiments. Note that the MAP \mathbf{Q} that we obtain in the subsequent experiments are systematically close to diagonal, for both models, which supports the little difference noticed between $\mathbf{Q} = q\mathbf{I}_x$ and the full \mathbf{Q} .

6.4.2. Results for the L96 reference model. The results for the L96 reference model in the nominal configuration with the approximate learning scheme described in Section 5.3 are reported in Figure 2, left column, where the FS, the LS, and the PSD are plotted.

We first observe some variability of the FS with the sample, which is small but non negligible. The asymptotic properties (LS and PSD) are however essentially insensitive to it. Then LS and PSD of the surrogate model are very close to the reference model ones, which shows that the surrogate model have very reliable asymptotic properties. The only discrepancy is seen for the PSD at high frequency where the surrogate model PSD saturates. This is due to the fact that, for the PSD, we implemented the surrogate model with the additive stochastic noise inferred from the training phase. This quantifies how reliable is the surrogate model depending on the frequency. Had we only use the deterministic part of the surrogate model, the PSD curves would have actually coincided even better at higher frequency (checked but not shown) but this match at high frequencies would be misleading.

Several scalar indicators are also given in Table 1.

6.4.3. Results for the L05III reference model. The results for the L05III reference model in the nominal configuration with the approximate learning scheme are reported in Figure 2, right column, where the FS, the LS, and the PSD are plotted. We observe a higher variability of all three indicators compared to L96. Even though the sensitivity of LS and FSD remain weak, it is significantly higher for the FS.

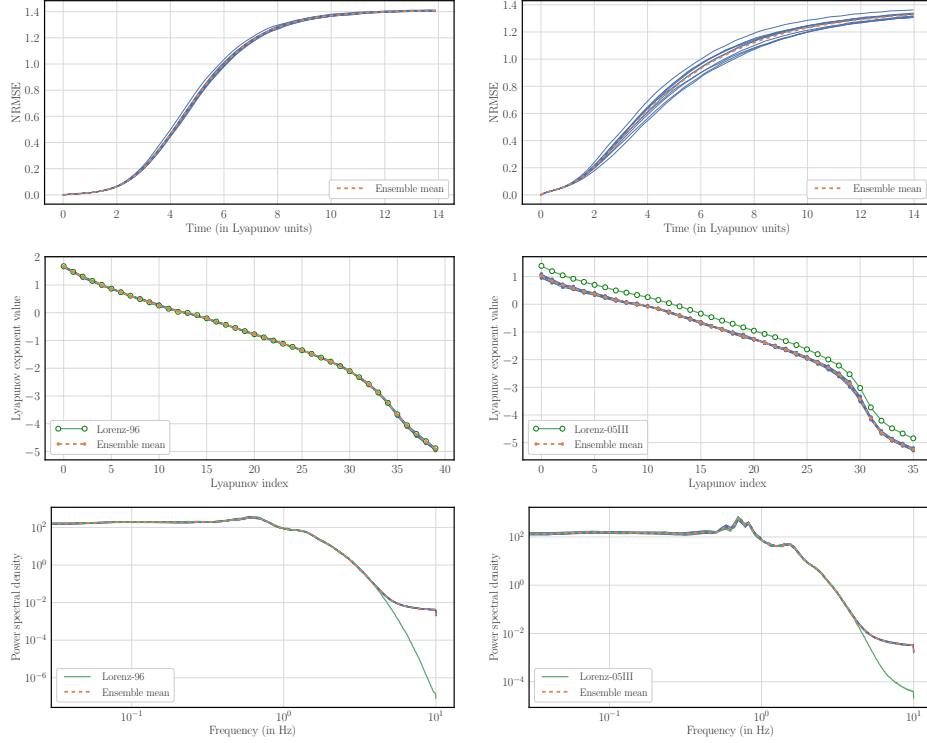


FIGURE 2. On the left hand side: Properties of the surrogate model obtained from full but noisy observation of the L96 model in the nominal configuration ($L = 4$, $K = 5000$, $\sigma_y = 1$, $N_y = N_x = 40$). On the right hand side: Properties of the surrogate model obtained from full but noisy observation of the L05III model in the nominal configuration ($L = 4$, $K = 5000$, $\sigma_y = 1$, $N_y = N_x = 36$). From top to bottom, are plotted the FS (NRMSE as a function of lead time in Lyapunov time), the LS (all exponents), and the PSD (in log-log-scale). A total of 10 experiments have been performed for both configurations. The curves corresponding to each member are drawn with thin blue lines while the mean of each indicator over the ensemble are drawn in thick dashed orange line.

The FS shows a steeper curve for short term prediction compared to L96, which is explained by the weaker identifiability of the L05III reference model. However it relaxes to the climatology at a slower rate compared to L96. This is due to the heterogeneity of the prediction: in the absence of fast mode burst (in space and time) into the slow compartment, the predictability can be long, or it can be shortened whenever the fast variables impact the slow variables. This has been explained and illustrated in [6].

The LS of the surrogate model follows closely that of the reference model but is systematically smaller. This may be due to the lack of variability of the surrogate model, where only the average impact of the fast variables are learned from the reference model.

Model	Scheme	$\pi_{\frac{1}{2}}$	σ_q	λ_1
L96	Approximate	4.56 ± 0.06	$0.08790 \pm 2 \cdot 10^{-5}$	1.66 ± 0.02
L96	Full	4.24 ± 0.07	0.09152	1.66 ± 0.02
L05III	Approximate	4.06 ± 0.21	$0.07720 \pm 2 \cdot 10^{-5}$	1.03 ± 0.05
L05III	Full	3.97 ± 0.17	0.08024	1.03 ± 0.04

TABLE 2. Scalar indicators for L96 and L05III in their nominal configuration, using either the full or the approximate schemes. The statistics of the indicators are obtained over 10 samples.

The PSD of the surrogate and reference model are close to each other except for the high frequencies. Similarly to the L96 case, at high frequency, the surrogate model PSD saturates, because the additive stochastic noise inferred from the training phase is accounted for. In the absence of the noise, we checked that there would be a saturation at much higher frequency, but that the PSD of the surrogate model would be offset compared to that of the reference model. This confirms that the stochastic correction plays the important role of stochastic subgrid parameterization in the L05III case.

Several scalar indicators are also given in Table 1.

6.5. Comparison of full and approximate scheme. We have compared the efficiency of the full and approximate schemes for both reference models in their nominal configuration. First of all, the full scheme is computationally much more demanding: it requires a storage at least N_e larger than the approximate one and it requires the computation of the cost function (17) which is N_e times more demanding than cost function (21) of the approximate scheme. In the L96 case, the results for the full and approximate schemes are very close to each other. The FS of the approximate scheme turns out to be slightly better, while the LS and PSD are almost identical. Compared scalar indicators are given in Table 2.

In the L05III case, the LS and PSD curves are also very close to each other. However, the FS of the approximate scheme is distinctively better than the FS of the full scheme as summarized by $\pi_{\frac{1}{2}}$ in Table 2.

The edge that the approximate scheme has on the full scheme as far as FS is concerned can be explained as follows. The approximate scheme looks for \mathbf{A} as the MAP rather than the average as the full scheme would. Hence, it has the tendency to improve the short term deterministic FS as much as possible and hence minimizes σ_q , at the risk of overfitting to the FS criterion. But it might be less reliable in terms of the representation of the uncertainty. This effect is more noticeable in the L96 case. It could be due to the better identifiability of the L96 reference model by the surrogate model.

6.6. Sensitivity on the length of the training window. One key advantage of the schemes proposed in Section 5 resides in their ability to handle long training windows. Note that a clear tendency when increasing the length K of the training windows was difficult to exhibit in [6, 8]. Here, we consider the nominal configurations for L96 and L05III, i.e. with $L = 4$, but with

$$K \in \{50, 100, 200, 400, 800, 1600, 3200, 6400\}. \quad (31)$$

Moreover, for each configuration, average indicators are obtained from 10 sample runs (each one with a different random seed).




FIGURE 3. Same as Figure 2 but for several values of the training window length K . Each curve is the mean over 10 experiments with different sets of observations. The LS and PSD of the reference models are also plotted for comparison.

The results are plotted in Figure 3, following the same structure as Figure 2, except that we now only plot the means but for several K values. It can be checked in the L96 case that the FS improves for all lead times with increasing K , as expected. The tendency is also verified in the L05III case, but not as clearly for large K . It may be due to the higher variability of the FS in the L05III case observed previously. But this is more likely due to a saturation of the performance as K increases, due to the weaker identifiability of the L05III model compared to that of the L96 model. For both models, the surrogate model LS has converged with good accuracy whenever $K \geq 400$, and the PSD converges to that of the reference model with good accuracy whenever $K \geq 200$, except for the high frequencies. Let us remark that the high frequency PSD is recovered better and better as K increases.

6.7. Sensitivity on the lag of the smoother. We have carried out experiments keeping $K = 5000$, but with the lag of the smoother chosen in

$$L \in \{0, 2, 4, 6, 8, 10, 12\}. \quad (32)$$

For each configuration, 10 independent samples are obtained. In the case of L05III, $L = 4$ yields the best FS, while $L = 2, 4$ are optimal for the FS in the case of L96. Note that in both cases, the configuration $L = 0$, i.e. filtering in place of smoothing,




FIGURE 4. On the left hand side: Properties of the surrogate model obtained from full but noisy observation of the L96 model in the nominal configuration ($L = 4$, $K = 5000$, $N_y = N_x = 40$ and with several σ_y). On the right hand side: Properties of the surrogate model obtained from full but noisy observation of the L05III model in the nominal configuration ($L = 4$, $K = 5000$, $\sigma_y = 1$, $N_y = N_x = 36$ and with several σ_y). From top to bottom, are plotted the FS (NRMSE as a function of lead time in Lyapunov time) and the PSD (in log-log-scale), averaged over an ensemble of 10 samples.

significantly degrades the FS. The LS and PSD are improved with $L \geq 2$ but only marginally. The reasons for these results had been tentatively given in Section 6.4.1.

6.8. Sensitivity on the observation noise. We also investigate the dependence on the observation noise magnitude. For both models, in the nominal configurations, the observation error standard deviation (recall $\mathbf{R} = \sigma_y^2 \mathbf{I}_x$) takes value in

$$\sigma_y = \{2, 1, 2^{-1}, 2^{-2}, 2^{-3}, 2^{-4}, 2^{-5}\}. \quad (33)$$

The results for the FS and PSD are reported in Figure 4. In the L96 case, the FS keeps increasing as σ_y decreases. This confirms the good identifiability of the L96 model by the surrogate model's architecture. In contrast, in the L05III case, the FS saturates with decreasing observation noise magnitude as soon as $\sigma_y = 1$, which is due to the mild identifiability given the chosen structure for the surrogate model. For both reference models, however, the PSD keeps improving at high frequencies as σ_y decreases.

6.9. Sensitivity on the density of the observations. Finally the dependence of the quality of the surrogate model as a function of the density of the monitoring network is investigated. For both models, we consider the nominal configuration but where the monitoring network at t_k , consists in the k -dependent random uniform draw of N_y sites where the variables are observed with an unbiased Gaussian noise of standard deviation $\sigma_y = 1$. The number of observations at each observation time




FIGURE 5. On the left hand side: Properties of the surrogate model obtained from partial and noisy observation of the L96 model in the nominal configuration ($L = 4$, $K = 5000$, $\sigma_y = 1$, $N_x = 40$) where N_y is varied. On the right hand side: Properties of the surrogate model obtained from partial and noisy observation of the L05III model in the nominal configuration ($L = 4$, $K = 5000$, $\sigma_y = 1$, $N_x = 36$) where N_y is varied. From top to bottom, are plotted the mean FS (NRMSE as a function of lead time in Lyapunov time), the mean LS (all exponents), and the mean PSD (in log-log-scale). A total of 10 experiments have been performed for both configurations.

step is taken in the set $N_y \in \{36, 32, 28, 24, 20, 16, 12\}$, for L05III and in the set $N_y \in \{40, 36, 32, 28, 24, 20, 16\}$, for L96. The results for the FS, LS and PSD are reported in Figure 5. As expected, the performance degrades significantly with the increase of sparsity. When N_y is smaller than $N_x/2$, some of the training steps may fail. This is most likely due to the inability of the EnKF/EnKS to be stable with higher sparsity, as is well known (see e.g., Section 4.3 in [6]). Localization is known to improve the robustness of the filter/smoother with higher sparsity but to a point.

As in the previous experiments, there is a saturation of the L05III FS performance for dense networks, as opposed to the L96 case which shows a significant improvement even from $N_y = 36$ to $N_y = 40$. Note that, in contrast with the previous results, the quality of LS and PSD are significantly dependent on the density of the monitoring network, especially in the L05III case.

7. Conclusion. In this paper, we have proposed a unifying Bayesian view on the extended data assimilation problem where not only the state trajectory is a control variable, but the dynamical model and its model error statistics are as well. This enables to learn about these control variables from partial and noisy observations of the physical system. The surrogate dynamical model is formalized as a neural network built with convolutional layers leveraging locality and possibly homogeneity assumptions. Techniques proposed so far to optimize the cost function associated to this problem have been re-interpreted as coordinate descent schemes. The EM technique was used to solve the marginal problem where the most likely model and model error statistics are both solved for. The algorithm alternates a traditional data assimilation step based on an EnKS and where a posterior ensemble is obtained, with an optimization step where new iterates of the model and the model error statistics are derived using machine learning tools, a quasi-Newton optimizer and analytical formula. The scheme enables to successfully handle very long training windows. Note that the outcome is not only a deterministic surrogate model but also its associated stochastic correction, representative of the uncertainty attached to the deterministic part, which accounts for residual model errors.

An approximate scheme where model error statistics are computed online from the ensemble during the EnKS step is proposed and compared to the full EM scheme which requires the ensemble to be stored. These two schemes are successfully tested on two low-order chaotic models with distinct identifiability, investigating the sensitivity to the length of the training window, the observation error magnitude, the density of the monitoring network, the lag of the EnKS, and using statistical indicators that probe short-term and asymptotic properties of the surrogate model.

This study has answered several of the questions originally raised when these combined data assimilation and machine learning methods were introduced in [6, 8]. A few others remain. How do we incorporate a priori information on the dynamical model? How do we build the corresponding online learning system as in sequential data assimilation? There are also technical challenges to address, such as those arising when applying these techniques to more complex, higher-dimensional models.

Acknowledgments. The authors are grateful to two reviewers for their comments and suggestions that helped improve the paper. The authors are thankful to Alban Farchi and Quentin Malartic for their suggestions on the original manuscript. This work was granted access to the HPC resources of IDRIS under the allocation AD011011184 made by GENCI. AC, JB and LB have been funded by the project REDDA (#250711) of the Norwegian Research Council. AC was also supported by the Natural Environment Research Council (Agreement PR140015 between NERC and the National Centre for Earth Observation). CERE and LOCEAN are members of Institut Pierre-Simon Laplace (IPSL).

REFERENCES

- [1] H. D. I. Abarbanel, P. J. Rozdeba and S. Shirman, [Machine learning: Deepest learning as statistical data assimilation problems](#), *Neural Computation*, **30** (2018), 2025–2055.
- [2] M. Asch, M. Bocquet and M. Nodet, [Data Assimilation: Methods, Algorithms, and Applications](#), Fundamentals of Algorithms, SIAM, Philadelphia, 2016.
- [3] C. M. Bishop (ed.), *Pattern Recognition and Machine Learning*, Springer-Verlag New-York Inc, 2006.
- [4] C. H. Bishop, B. J. Etherton and S. J. Majumdar, [Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects](#), *Mon. Wea. Rev.*, **129** (2001), 420–436.

- [5] M. Bocquet, [Ensemble Kalman filtering without the intrinsic need for inflation](#), *Nonlin. Processes Geophys.*, **18** (2011), 735–750.
- [6] M. Bocquet, J. Brajard, A. Carrassi and L. Bertino, [Data assimilation as a learning tool to infer ordinary differential equation representations of dynamical models](#), *Nonlin. Processes Geophys.*, **26** (2019), 143–162.
- [7] M. Bocquet and P. Sakov, [Joint state and parameter estimation with an iterative ensemble Kalman smoother](#), *Nonlin. Processes Geophys.*, **20** (2013), 803–818.
- [8] J. Brajard, A. Carrassi, M. Bocquet and L. Bertino, [Combining data assimilation and machine learning to emulate a dynamical model from sparse and noisy observations: a case study with the Lorenz 96 model](#), *J. Comput. Sci.*, 2020, <http://arxiv.org/pdf/2001.01520.pdf>.
- [9] S. L. Brunton, J. L. Proctor and J. N. Kutz, [Discovering governing equations from data by sparse identification of nonlinear dynamical systems](#), *PNAS*, **113** (2016), 3932–3937.
- [10] R. H. Byrd, P. Lu and J. Nocedal, [A limited memory algorithm for bound constrained optimization](#), *SIAM Journal on Scientific and Statistical Computing*, **16** (1995), 1190–1208.
- [11] M. Carlu, F. Ginelli, V. Lucarini and A. Politi, [Lyapunov analysis of multiscale dynamics: The slow bundle of the two-scale Lorenz 96 model](#), *Nonlin. Processes Geophys.*, **26** (2019), 73–89.
- [12] A. Carrassi, M. Bocquet, L. Bertino and G. Evensen, [Data assimilation in the geosciences: An overview on methods, issues, and perspectives](#), *WIREs Climate Change*, **9** (2018), e535.
- [13] B. Chang, L. Meng, E. Haber, F. Tung and D. Begert, [Multi-level residual networks from dynamical systems view](#), in *Proceedings of ICLR 2018*, 2018.
- [14] F. Chollet, *Deep Learning with Python*, Manning Publications Company, 2017.
- [15] A. P. Dempster, N. M. Laird and D. B. Rubin, [Maximum likelihood from incomplete data via the EM algorithm](#), *Journal of the Royal Statistical Society. Series B*, **39** (1977), 1–38.
- [16] P. D. Dueben and P. Bauer, [Challenges and design choices for global weather and climate models based on machine learning](#), *Geosci. Model Dev.*, **11** (2018), 3999–4009.
- [17] W. E, [A proposal on machine learning via dynamical systems](#), *Commun. Math. Stat.*, **5** (2017), 1–11.
- [18] G. Evensen, *Data Assimilation: The Ensemble Kalman Filter*, 2nd edition, Springer-Verlag Berlin Heidelberg, 2009.
- [19] R. Fablet, S. Ouala and C. Herzet, [Bilinear residual neural network for the identification and forecasting of dynamical systems](#), in *EUSIPCO 2018, European Signal Processing Conference*, Rome, Italy, 2018, 1–5.
- [20] M. Fisher and S. Gürol, [Parallelization in the time dimension of four-dimensional variational data assimilation](#), *Q. J. R. Meteorol. Soc.*, **143** (2017), 1136–1147.
- [21] Z. Ghahramani and S. T. Roweis, [Learning nonlinear dynamical systems using an EM algorithm](#), in *Advances in neural information processing systems*, 1999, 431–437.
- [22] W. W. Hsieh and B. Tang, [Applying neural network models to prediction and data analysis in meteorology and oceanography](#), *Bull. Amer. Meteor. Soc.*, **79** (1998), 1855–1870.
- [23] B. R. Hunt, E. J. Kostelich and I. Szunyogh, [Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter](#), *Physica D*, **230** (2007), 112–126.
- [24] E. Kalnay, H. Li, T. Miyoshi, S.-C. Yang and J. Ballabrera-Poy, [4-D-Var or ensemble Kalman filter?](#), *Tellus A*, **59** (2007), 758–773.
- [25] Y. A. LeCun, L. Bottou, G. B. Orr and K.-R. Müller, [Efficient backprop](#), in *Neural networks: Tricks of the trade*, Springer, 2012, 9–48.
- [26] R. Lguensat, P. Tandeo, P. Ailliot, M. Pulido and R. Fablet, [The analog data assimilation](#), *Mon. Wea. Rev.*, **145** (2017), 4093–4107.
- [27] Y. Liu, J.-M. Haussaire, M. Bocquet, Y. Roustan, O. Saunier and A. Mathieu, [Uncertainty quantification of pollutant source retrieval: comparison of Bayesian methods with application to the Chernobyl and Fukushima-Daiichi accidental releases of radionuclides](#), *Q. J. R. Meteorol. Soc.*, **143** (2017), 2886–2901.
- [28] Z. Long, Y. Lu, X. Ma and B. Dong, [PDE-Net: Learning PDEs from Data](#), in *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- [29] E. N. Lorenz, [Designing chaotic models](#), *J. Atmos. Sci.*, **62** (2005), 1574–1587.
- [30] E. N. Lorenz and K. A. Emanuel, [Optimal sites for supplementary weather observations: Simulation with a small model](#), *J. Atmos. Sci.*, **55** (1998), 399–414.
- [31] L. Magnusson and E. Källén, [Factors influencing skill improvements in the ECMWF forecasting system](#), *Mon. Wea. Rev.*, **141** (2013), 3142–3153.

- [32] V. D. Nguyen, S. Ouala, L. Drumetz and R. Fablet, EM-like learning chaotic dynamics from noisy and partial observations, arXiv preprint, [arXiv:1903.10335](https://arxiv.org/abs/1903.10335).
- [33] J. Paduart, L. Lauwers, J. Swevers, K. Smolders, J. Schoukens and R. Pintelon, [Identification of nonlinear systems using polynomial nonlinear state space models](#), *Automatica*, **46** (2010), 647–656.
- [34] D. C. Park and Y. Zhu, [Bilinear recurrent neural network](#), in *IEEE World Congress on Computational Intelligence., 1994 IEEE International Conference on*, **3** (1994), 1459–1464.
- [35] J. Pathak, B. Hunt, M. Girvan, Z. Lu and E. Ott, [Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach](#), *Phys. Rev. Lett.*, **120** (2018), 024102.
- [36] J. Pathak, Z. Lu, B. R. Hunt, M. Girvan and E. Ott, [Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data](#), *Chaos*, **27** (2017), 121102, 9pp.
- [37] 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*, **70** (2018), 1–17.
- [38] P. N. Raanes, A. Carrassi and L. Bertino, [Extending the square root method to account for additive forecast noise in ensemble methods](#), *Mon. Wea. Rev.*, **143** (2015), 3857–3873.
- [39] V. Rao and A. Sandu, [A time-parallel approach to strong-constraint four-dimensional variational data assimilation](#), *J. Comp. Phys.*, **313** (2016), 583–593.
- [40] M. Reichstein, G. Camps-Valls, B. Stevens, J. Denzler, N. Carvalhais and Prabhat, [Deep learning and process understanding for data-driven Earth system science](#), *Nature*, **566** (2019), 195–204.
- [41] P. Sakov, J.-M. Haussaire and M. Bocquet, [An iterative ensemble Kalman filter in presence of additive model error](#), *Q. J. R. Meteorol. Soc.*, **144** (2018), 1297–1309.
- [42] P. Tandeo, P. Ailliot, M. Bocquet, A. Carrassi, T. Miyoshi, M. Pulido and Y. Zhen, A review of innovation based approaches to jointly estimate model and observation error covariance matrices in ensemble data assimilation, 2020, <https://arxiv.org/abs/1807.11221>, Submitted.
- [43] Y. Trémolet, Accounting for an imperfect model in 4D-Var, *Q. J. R. Meteorol. Soc.*, **132** (2006), 2483–2504.
- [44] P. R. Vlachas, J. Pathak, B. R. Hunt, T. P. Sapsis, M. Girvan, E. Ott and P. Koumoutsakos, Forecasting of spatio-temporal chaotic dynamics with recurrent neural networks: a comparative study of reservoir computing and backpropagation algorithms, arXiv preprint, [arXiv:1910.05266](https://arxiv.org/abs/1910.05266).
- [45] Y.-J. Wang and C.-T. Lin, Runge-Kutta neural network for identification of dynamical systems in high accuracy, *IEEE Transactions on Neural Networks*, **9** (1998), 294–307.
- [46] G. C. G. Wei and M. A. Tanner, [A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms](#), *Journal of the American Statistical Association*, **85** (1990), 699–704.
- [47] P. Welch, [The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms](#), *IEEE Transactions on Audio and Electroacoustics*, **15** (1967), 70–73.
- [48] S. J. Wright, [Coordinate descent algorithms](#), *Mathematical Programming*, **151** (2015), 3–34.

Received for publication January 2020.

E-mail address: marc.bocquet@enpc.fr

E-mail address: julien.brajjard@nersc.no

E-mail address: n.a.carrassi@reading.ac.uk

E-mail address: laurent.bertino@nersc.no