



HAL
open science

Latent space data assimilation by using deep learning

Mathis Peyron, Anthony Fillion, Selime Gürol, Serge Gratton, Pierre Boudier,
Gael Goret

► **To cite this version:**

Mathis Peyron, Anthony Fillion, Selime Gürol, Serge Gratton, Pierre Boudier, et al.. Latent space data assimilation by using deep learning. Quarterly Journal of the Royal Meteorological Society, 2021, 147 (740, Part A), pp.3759-3777. 10.1002/qj.4153 . hal-03973243

HAL Id: hal-03973243

<https://hal.science/hal-03973243v1>

Submitted on 16 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Latent Space Data Assimilation by using Deep Learning

Mathis Peyron^{1,2,4}, Anthony Fillion^{5,1}, Selime Gürol^{4,1}, Victor Marchais¹, Serge Gratton^{5,1}, Pierre Boudier^{3,1}, and Gael Goret²

¹ANITI, Université de Toulouse

²Atos BDS R&D AI4Sim

³NVIDIA

⁴CERFACS

⁵UFTMIP

Abstract

Performing Data Assimilation (DA) at a low cost is of prime concern in Earth system modeling, particularly at the time of big data where huge quantities of observations are available. Capitalizing on the ability of Neural Networks techniques for approximating the solution of PDE's, we incorporate Deep Learning (DL) methods into a DA framework. More precisely, we exploit the latent structure provided by autoencoders (AEs) to design an Ensemble Transform Kalman Filter with model error (ETKF-Q) in the latent space. Model dynamics are also propagated within the latent space via a surrogate neural network. This novel ETKF-Q-Latent (thereafter referred to as ETKF-Q-L) algorithm is tested on a tailored instructional version of Lorenz 96 equations, named *the augmented Lorenz 96 system*: it possesses a latent structure that accurately represents the observed dynamics. Numerical experiments based on this particular system evidence that the ETKF-Q-L approach both reduces the computational cost and provides better accuracy than state of the art algorithms, such as the ETKF-Q.

1 Introduction

Data Assimilation estimates the state of a system $\mathbf{x} \in \mathbb{R}^n$ given two sources of information: a model that provides a background knowledge $\mathbf{x}^b \in \mathbb{R}^n$ and an observation vector $\mathbf{y} \in \mathbb{R}^p$. We can either observe the system directly - partially or entirely - or implicitly through an observation operator \mathcal{H} , yielding the observation relation $\mathbf{y} = \mathcal{H}(\mathbf{x})$ [3, 7].

Most variational and ensemble DA algorithms encounter two problems that will be considered in the present paper.

First, they are largely grounded on the Kalman filter which assumes that errors follow a Gaussian distribution and the physical model is mildly nonlinear. Another issue is the cost of applying the propagation model that can be computationally prohibitive. Thus, we propose to follow a data-driven approach which can efficiently represent nonlinear dynamics as long as a sufficient

amount of data is given. They can also make better use of big data and exploit better the future computing hardware.

In the last decade, machine learning (ML) techniques managed to outperform existing methods in image classification [29, 31, 26, 54], segmentation [43, 20], Natural Language Processing (NLP) [41, 21], language translation [55, 67] and in mastering go game [52, 53] for instance. It is now broadening to new areas like computational physics where it achieves interesting results too, as exemplified by the so-called physically informed artificial neural networks [44, 45, 40].

Nevertheless, the widespread applicability of ML to plethora of problems is hampered by significant concerns like the difficulty to incorporate physical knowledge in DL frameworks [64, 10] or the existence of adversarial examples [56, 14, 48] that reduces practical applications in sensitive fields (e.g. automated driving). We can also mention the absence of uncertainty estimation analysis [47].

As suggested by [22], Data Assimilation and Deep Learning can take advantage of each other as they are complementary and have also similarities. Therefore, coupling DA and DL is quite a natural approach. In this study, we are interested in utilizing DL tools within a DA framework to:

- reduce the computational cost and memory storage by performing calculations in a reduced space (latent space).
- get a better accuracy by exploiting the latent structure and a surrogate network.
- propose a novel and promising framework that could be extended to other DA methods and NNs. The proposed method simplicity is also a great advantage: there is no need to change the DA algorithm, only the operators.

1.1 Related work and our contributions

Optimizing DA methods is a key issue in Earth system modeling, particularly in the context of big data where huge quantities of observations are now available (thanks to remote sensing among other [30]), and where the increase of the complexity of physical models goes hand in hand with the augmentation of the computational resources needed in applications. For this reason, model reduction techniques and surrogate models have been investigated by the DA community.

As shown in [3, Chapter 5], DA reduction techniques rely on multiple mathematical tools in order to alleviate the computational cost: singular value decompositions, principal component analysis (PCA) or proper orthogonal decomposition (POD), spectral decomposition (*e.g.* Fourier series), wavelets and curvelets for instance. Reduction methods in DA can be divided into two groups: either they aim at reducing the numerical cost of the time integration of the model or they improve the DA method itself.

Probably the most direct approach in the first category consists in simplifying the model (see [38]), whereas more sophisticated approaches rely on POD [13, 2] or on wavelets [57]. Rank reduction techniques have been investigated as well for both variational and sequential methods: the core change is to introduce low rank approximations of the covariance matrices, without changing the complexity of the time integration. In this case, some methods focus on factorizing the model error covariance matrix by possibly relying on sparse structures (see [6] for a review). Propagating such a matrix through time is also expensive, hence the idea is to use a Truncated Singular Value Decomposition (TSVD) on the model linear operator [17].

When coming to improving the DA algorithm by space reduction, we find the Singular Evolutive Extended Kalman Filter (SEEK) [61] and, using sampling techniques, ensemble DA methods like Ensemble Kalman Filter (EnKF), the reduced-rank square root filter (RRSqrt), the stochastic EnKF and the deterministic EnKF.

Deep Learning also proposes reduction techniques along with surrogate models. It is also establishing itself in computational physics. Recent research shows that PDEs can be solved efficiently by using DL methods and the computational cost of simulations can be reduced by considering the latent space structure provided by NNs [44, 45, 32, 51, 25, 27]. For instance, the Physics Informed Neural Networks (PINNs) [44, 45] are capable to properly learn the solution of a nonlinear PDE while being computationally cheap. Also, one of the most promising research work, though not investigated here, consists in making the most of Fourier transforms to solve PDEs [32]: this technique has the advantage to be both accurate and computationally efficient.

Creating a latent space is mostly performed with well-known networks like autoencoders (AE), convolu-

tional autoencoders (CAE) or variational autoencoders (VAE). The aim in creating a latent space is either to get a better accuracy or to reduce the computational cost. For example, in [12] a convolutional VAE is used to construct a continuous parameterization for facies in order to preserve the geological realism of the model (*i.e.* predictions for oil and gas reservoirs). This leads to a better accuracy. In [65, 66], the authors are mainly interested in computational speed up and robust long-term predictions for fluid flows simulations. They therefore demonstrate the capability of data-driven approaches for modeling fluid dynamics: they implemented a CAE for spatial compression and stacked long-short term memory (LSTM) layers to define their surrogate network, *i.e.* a network that performs time propagation. They achieve a compression ratio of 256 enabling them to replace heavy fluid flows simulations. Similarly, [37] resorts to a CAE for space reduction and recurrent neural networks for the time propagation. About [35], in addition to Control Variable Transform (CVT) they take advantage of a CAE inspired by the image compression field that uses state-of-the-art deep learning techniques. In [46], fully-connected layers are considered to perform the reduction. Lastly, we can quote [19] where the authors utilize an AE to encode 3D virtual figurines, then deform their motions within the latent space and finally get the new pose in the full space. Time propagation within the latent space is performed with a mere matrix product given by the so-called Koopman operator. The reduced space obtained with this method is of much smaller dimension than linear model reduction techniques, therefore it yields faster and more accurate results along with improved robustness.

When including model reduction techniques within DA frameworks, we find [12] and [35] which both employ neural networks in their data assimilation architectures. Yet, they do not have the same purposes. Indeed, [12] trains a convolutional variational autoencoder (CVAE) with the aim to generate realistic geological facies by introducing a new continuous parameterization which was not properly achieved with prior techniques. They use a type of ensemble smoother based on multiple data assimilation and assimilate the data in the new parametrization (latent) space which leads to a better accuracy. They do not use a surrogate network in their DA method.

Regarding [35], the goal is to reduce the physical domain into a latent space to speed up computations in the context of an ensemble based DA without temporal aspect. Besides the temporal aspect, the main difference with our methodology is that they introduce an observation encoder network which includes an interpolation operator to map the observations to the full space. This observation network maps the observations to the latent space where data assimilation is performed. Having an interpolation operator may introduce additional errors during DA. Whereas in our method, observations stay in their original space and we do not need to apply any transformation. In their latent space DA approach, they also only focus on the computational cost gain.

As for surrogate networks coupled with DA algorithms, we have [42] and [11] which both consider Lorenz 96 system. In [42], Pawar and San model unresolved flow dynamics with a surrogate network and learn the correlation between resolved flow processes and unresolved subgrid variables thanks to a set of NNs. Whereas we both aim to more accurately forecast, their approach does not involve any space reduction technique. Their main motivation is not on reducing the computational cost either. Regarding [11], Brajard et al. use an ensemble Kalman filter algorithm combined to a convolutional neural network with *skipconnections* [26]. What is remarkable here is that the surrogate model is iteratively trained with the data assimilation algorithm, whereas in our case NNs are not aware of the existence of an outer data assimilation process: our networks are trained independently of the ETKF-Q algorithm. Nonetheless, Brajard et al. consider neither standard space reduction techniques nor latent spaces obtained with NNs like we do.

A first ingredient in the approach we introduce in the present paper is to replace the (supposedly expensive) time integration of the model with a NN surrogate. Time stepping methods based on surrogates have already been explored by several authors [34, 65, 66, 37, 63, 11, 42]. A key question that comes immediately is to ensure the time stability of the resulting scheme (see [24, 23] for an investigation within deep learning surrogate models) when the model is repeatedly called to propagate the state over several time steps. To this aim, we introduce in the training loss an explicit stabilization that involves a penalization of the growth of model iterations.

We also address the question of incorporating latent spaces in DA frameworks. In general, DA and ML are not really coupled, model reduction or surrogate models are used only for the time propagation within an existing DA algorithm. We believe our approach is original in that it performs DA directly in the NN latent space. This approach is appealing, since it fully exploits an underlying geometry and leads to perform the computations mostly in the space where they are cheap (the latent space). We shall see that to obtain good results, a special care has to be taken in the description of the dynamical and observational errors.

Therefore, in this paper we propose a new latent space DA methodology, that is the ETKF-Q-L one, which:

1. explores the ability of DL for creating a ℓ -dimensional reduced space, based on the assumption that a latent space of size ℓ which accurately represents the full dynamics exists. This is achieved with an autoencoder.
2. defines a surrogate network within the latent space that performs time propagation. An innovative iterative training enforces the surrogate to be stable over time.

3. implements an ensemble DA algorithm within the learned ℓ -dimensional latent space thanks to the AE and the surrogate network.

Interestingly, 1. and 2. can be performed in an all-at-once approach by training both the AE and the surrogate at the same time through a well-suited custom loss function. The training set is an ensemble of simulations of the physical system which lies in \mathbb{R}^N . For short, the proposed methodology provides the following advantages:

1. since any DA algorithm requires to store vectors lying in the model space, discovering a lower dimensional representation induces a reduction in memory needs and computational cost.
2. performing the DA linear analysis in the latent space obtained by AE is less susceptible to yield non-physical solutions since the decoder is a nonlinear transformation that fits the manifold where state trajectory statistically belongs, when such a structure does exist.

We show the relevance of our approach on a 400-dimensional system possessing an underlying dynamics that follows the 40-variables Lorenz 96 equations. The existence of this latent physics is ensured by construction of the full space dynamics which is built upon Lorenz 96 system.

1.2 Organization of this work

The remainder of this work is organized as follows: in [section 2](#), we precisely detail the role of the AE NN structure for the latent space computation. We also present the surrogate network used for time stepping and expose the loss function that allows the joint learning of these two networks.

In [section 3](#), we first remind some general DA concepts, then we explain an ensemble algorithm that takes model error into account that is the ETKF-Q algorithm. Lastly, we detail our latent DA framework namely the ETKF-Q-L method, based on the ETKF-Q algorithm associated with the AE and the surrogate.

In [section 4](#), we present the experimental context chosen to benchmark our approach, *i.e.* the augmented Lorenz 96 dynamics, the NN architectures and the training setting along with the DA context. We show the performance of the new DA method on the augmented Lorenz 96 system. A grid search algorithm is used to tune the DA algorithm parameters.

2 Latent Space Dynamics

In this section we explain how to approximate the model dynamics under consideration in a latent space. We first briefly present the autoencoder (AE) structures, which are widely used for dimension reduction [28], and then provide more details on the surrogate network which is nested in the AE ([figure 1](#)). We also discuss

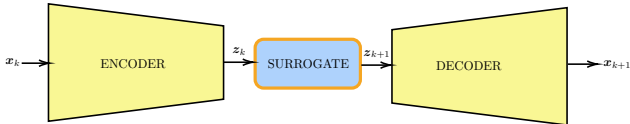


Figure 1: Training architecture of the autoencoder and the surrogate: the encoder maps from \mathbb{R}^N to \mathbb{R}^ℓ and the decoder performs the reverse operation. The surrogate, *i.e.* the time propagator operates within the latent space. Both networks are trained together.

how to obtain a stable trajectory of the dynamics by reformulating the loss function.

Resorting to a latent space data assimilation is also motivated by [60, 16, 58, 59]. In the case of nonlinear chaotic dynamics, the state space can be indeed divided into unstable and stable subspaces. It turns out that perturbations mostly live in the unstable manifold which is in general of much smaller dimension than the full space. In [8], these results are extended to the Kalman smoother and ensemble formulations of the Kalman filter and smoother are considered. This motivates the link we make between latent DA and unstable manifold theory; our goal is to learn a latent space that would capture the unstable manifold. Indeed, as shown by [62, 15], tracking the unstable directions is of primary relevance to properly perform the assimilation in possibly high-dimensional models such as the atmospheric or oceanic models.

2.1 Autoencoders

Autoencoders are a type of NNs that are trained to reproduce the input data by enforcing them to be accurately represented in a lower dimension [28]. They consist of an encoder and a decoder trained together. Given data lying in \mathbb{R}^N , the encoder maps from \mathbb{R}^N to \mathbb{R}^ℓ (with $\ell \ll N$) and is generally made of successive fully connected or convolutional layers with decreasing dimensions. The decoder performs the reverse operation and therefore mirrors encoder’s layers. A common loss function is the Mean Square Error (MSE):

$$\text{MSE}(\mathbf{x}, \mathcal{D}(\mathcal{E}(\mathbf{x}))) = \frac{1}{N} \sum_{i=1}^N \left(\mathbf{x}^{(i)} - [\mathcal{D}(\mathcal{E}(\mathbf{x}))]^{(i)} \right)^2 \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^N$, $\mathbf{x}^{(i)}$ denotes the i -th element of vector \mathbf{x} , \mathcal{E} and \mathcal{D} denote the encoder and the decoder, respectively. Note that unlike Principal Component Analysis (PCA), AEs leverage nonlinear transformations and are thus better suited to handling nonlinearities [28].

In order to apprehend how AEs work, let us consider MNIST dataset¹ made of hand-written digits stored as 28×28 images or vectors of size 784: a very simple encoder could contain 4 fully connected layers whose input dimensions could be as follows: 784, 300, 150, 100, 20. Then, the decoder would hold the same number of layers in the reverse order (*i.e.* with dimensions

20, 100, 150, 300, 784, respectively). Hence the bottleneck structure of AE networks with the reduced space at the encoder-decoder junction point.

While learning, encoder and decoder’s weights are modified so that the autoencoder can reconstruct the input digit with the strong requirement that data have to be well represented in a ℓ -dimensional space called the *latent space*.

AE’s quality is highly impacted by the type and the number of layers, their dimensions and the size of the latent space chosen. The major issue, that is used tackled numerically by hyperparameter tuning, is to be able to find the smallest latent space that enables to represent the data the more accurately as possible.

In this study, which is more of a proof of concept, we assume available a system of size N for which a latent space of lower dimension ℓ is deemed to exist and in which the observed dynamical system can be described. Our encoder maps from \mathbb{R}^N to \mathbb{R}^ℓ and the decoder performs the reverse operation. We want to emphasize that there is absolutely no reason for the latent space of dimension ℓ produced by the autoencoder to be unique. Indeed, the loss function used in the training promotes a coherence between the triplet consisting of the decoder, the encoder and the latent space on the one hand, and data on the other hand. Whenever one particular latent space is discovered (the network is completely free in the way it designs the latent space), other latent spaces exist as well, obtained by transformations such as rotations, or changes of scales.

2.2 Surrogate network and stability

We want to train a surrogate network such that time propagation of the model dynamics can be performed in the latent space (obtained by the AE). Therefore, our surrogate network is estimated using encoded data and outputs a transformation acting on, and producing latent vectors. Just like the AE, a first idea would consist in training the surrogate network through a MSE loss function as follows:

$$\text{MSE}(\mathbf{x}_{k+1}, \mathcal{T}(\mathbf{x}_k)) = \frac{1}{N} \sum_{i=1}^N \left(\mathbf{x}_{k+1}^{(i)} - [\mathcal{T}(\mathbf{x}_k)]^{(i)} \right)^2 \quad (2)$$

where $\mathbf{x}_k, \mathbf{x}_{k+1} \in \mathbb{R}^N$ are the state vectors at time t_k and t_{k+1} , respectively and operator \mathcal{T} is such that $\mathcal{T}(\cdot) = \mathcal{D}(\mathcal{S}(\mathcal{E}(\cdot)))$ with \mathcal{E} , \mathcal{S} and \mathcal{D} denoting the encoder, the surrogate and the decoder, respectively.

Nonetheless, training our surrogate with this loss function (equation (2)) does not yield a stable solution. This is especially easy to understand when the dynamics under consideration is chaotic, as often the case in data assimilation. In this case, if the non-vanishing components of the dynamics are not represented with enough accuracy, the surrogate dynamics is expected to be of insufficient quality. This is even worse in

¹see <http://yann.lecun.com/exdb/mnist/> for more details

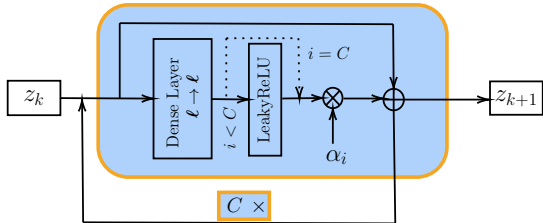


Figure 2: Surrogate network: it is based on skipconnections [26], and more precisely on an updated version of them [5]. The network consists of fully connected layers of dimension ℓ .

the case where the original dynamics would exhibit conservative components; if the surrogate dynamics does not capture these components accurately enough, it is easy to conceive that nonphysical unstable subspaces may occur, making the latent space time stepping with the surrogate inappropriate for DA.

Issues related to stable NNs approximation of time stepping methods have already been investigated in the literature though outside of our DA context.

They have been linked to exploding or vanishing gradients issues and NNs’ robustness as well. [24, 23] get some insights in this direction by proposing groundbreaking methods to make deep neural networks stable. However, the problem they address is not exactly the one we are looking at: they focus on Deep Neural Networks’ (DNNs) robustness to input perturbation, on their capability to distinguish between two initial vector states, *i.e.* not to bring both of them to 0 nor making them diverging.

Within the framework of DA, the presence of non-physical unstable dynamics components is controlled by using a simple penalty approach involving a technique we describe now.

Our method relies on a chained loss function, meaning that we train the surrogate to predict c successive states to enforce stability. In practice, given $\mathbf{x}_k \in \mathbb{R}^N$, the encoder yields $\mathbf{z}_k \in \mathbb{R}^\ell$. Then, the surrogate outputs $\mathbf{z}_{k+1}, \dots, \mathbf{z}_{k+c}$ which are all decoded afterwards and their distances to the ground truth states are measured through a custom loss function defined as follows:

$$\frac{1}{C} \sum_{c=1}^C \text{MSE}(\mathcal{T}^c(\mathbf{x}_k), \mathbf{x}_{k+c}) \quad (3)$$

where \mathcal{T}^c is a straightforward extension of operator \mathcal{T} : $\mathcal{T}^c(\cdot) = \mathcal{D}(\mathcal{S}^c(\mathcal{E}))$. Regarding \mathcal{S}^c , it means that the surrogate is applied c times in a row over the given data.

One remaining question is the number of iterations C we need to perform in order to achieve this stability criterion: according to our numerical experiments based on the augmented Lorenz 96 system, just 2 consecutive predictions already guarantee a stable behaviour. In the numerical tests, we pick this parameter in $\{2, 3, 4\}$.

In the surrogate’s architecture, we found decisive the use of so-called *skipconnections* [26] which are now a

common and good practice. It consists in adding the result of layer i to the one of layer $i - 1$ in the form $\mathbf{z} = \mathbf{z} + \text{layer}_i(\mathbf{z})$. This way we predict the increment needed to reach \mathbf{z}_{k+1} from \mathbf{z}_k rather than the raw output directly. We even go a step further as [5] proposed an updated version of *skipconnections* that performs better: $\mathbf{z} = \mathbf{z} + \alpha_i \text{layer}_i(\mathbf{z})$ where α_i are trainable parameters. Our surrogate learned with and without the modulation parameters α_i : they are clearly beneficial as we achieve better results while including them. Figure 2 exposes the architecture of the surrogate network with the fully connected layers of dimension ℓ and the associated *LeakyReLU* activations.

2.3 Learning strategy: training both networks together

As often the case when optimizing functions of several variables, performing sequential optimization by group of variables may be appealing since it reduces the search space of each optimization step. However it generally leads to a sub-optimal solution. In our case too, numerical experiments, not reported here, showed that training both the AE and the surrogate together gives better results than training the AE first and then the surrogate. Since AE’s quality influences surrogate’s performances, a combined training allows them to “communicate” and “share” information in order to more properly learn: the latent space is designed to fit the surrogate and vice versa. To do so, we define a custom loss function with a weighting parameter ρ that balances between equation (1) and equation (3):

$$\mathcal{L}(\mathbf{x}_{k:k+C}) = \mathcal{L}_{\text{AE}}(\mathbf{x}_{k+1:k+C}) + \rho \times \mathcal{L}_{\text{Sur}}(\mathbf{x}_{k:k+C}) \quad (4)$$

where:

$$\mathcal{L}_{\text{AE}}(\mathbf{x}_{k+1:k+C}) = \frac{1}{C} \sum_{c=1}^C \text{MSE}(\mathcal{D}(\mathcal{E}(\mathbf{x}_{k+c})), \mathbf{x}_{k+c})$$

and:

$$\mathcal{L}_{\text{Sur}}(\mathbf{x}_{k:k+C}) = \frac{1}{C} \sum_{c=1}^C \text{MSE}(\mathcal{T}^c(\mathbf{x}_k), \mathbf{x}_{k+c})$$

where $\mathbf{x}_{k:k+C}$ denotes the sequence $[\mathbf{x}_k, \mathbf{x}_{k+1}, \dots, \mathbf{x}_{k+C}]$.

3 Data Assimilation within a latent space

Before presenting in details our ETKF-Q-Latent algorithm, we first remind basic facts of sequential DA. The upcoming sections section 3.1 and section 3.2 are strongly inspired by [7, 4, 18].

3.1 Sequential Data Assimilation

Sequential or statistical DA is based on estimation theory and refers to a DA method for which observations are

sequentially assimilated as they become available.

Sequential data assimilation deals with the following stochastic-dynamical system:

$$\begin{cases} \mathbf{y}_k &= \mathcal{H}_k(\mathbf{x}_k) + \boldsymbol{\varepsilon}_k \\ \mathbf{x}_k &= \mathcal{M}_k(\mathbf{x}_{k-1}) + \boldsymbol{\eta}_k \end{cases} \quad (5)$$

where $\mathbf{x}_k \in \mathbb{R}^n$, $\mathbf{y}_k \in \mathbb{R}^p \forall k \in [0, K]$, \mathcal{M}_k is the nonlinear dynamical model used for time propagation of the state from time t_{k-1} to time t_k with the additive model error $\boldsymbol{\eta}_k$, and \mathcal{H}_k is the observation operator, mapping the state from the model space to the observation space with the observation error $\boldsymbol{\varepsilon}_k$. The errors are assumed to be all unbiased, uncorrelated in time and independent from \mathbf{x}_0 .

In sequential DA, the state estimation by using this stochastic-dynamical system is obtained based on the Bayesian approach which takes into account probability distributions of the errors. Available observation is used to update the conditional probability density function (pdf) (analysis step), and then this pdf is propagated to the next time step (forecast step).

A common choice for pdf is the Gaussian distribution since many processes are well described with it and it is algebraically convenient. Let us assume that in (equation (5)), the observation error follows a Gaussian distribution with zero mean and a covariance matrix \mathbf{R}_k and similarly model error follows a Gaussian distribution with zero mean and a covariance matrix \mathbf{Q}_k .

Using Gaussian error pdfs and under the assumption that the model and observation operators are linear (denoted by M_k and H_k , respectively), the Kalman filter recursively finds the analysis as the conditional mean of the posteriori pdf:

$$\mathbf{x}_k^a = \mathbf{x}_k^f + \mathbf{K}_k \left(\mathbf{y}_k - H_k \mathbf{x}_k^f \right). \quad (6)$$

Here, the state \mathbf{x}_k^f at time t_k represents the model prediction from the analysis at time t_{k-1} , *i.e.* $\mathbf{x}_k^f = M_k(\mathbf{x}_{k-1}^a)$ and \mathbf{K}_k denotes the Kalman gain matrix at time k :

$$\mathbf{K}_k = \mathbf{P}_k^f H_k^T \left(\mathbf{R}_k + H_k \mathbf{P}_k^f H_k^T \right)^{-1} \quad (7)$$

where \mathbf{P}_k^f is the error covariance matrix of the forecast \mathbf{x}_k^f . Note that the estimation given by (equation (6)) is also known as the **BLUE** (Best Linear Unbiased Estimator) estimate, which gives the minimum variance analysis with the choice of \mathbf{K}_k provided by (equation (7)).

Once the analysis is derived, the estimate and its error covariance matrix are propagated through time:

$$\mathbf{x}_{k+1}^f = M_{k+1}(\mathbf{x}_k^a) \quad (8)$$

$$\mathbf{P}_{k+1}^f = M_{k+1} \mathbf{P}_k^a M_{k+1}^T + \mathbf{Q}_{k+1} \quad (9)$$

where \mathbf{P}_k^a is the error covariance matrix of the analysis. \mathbf{P}_k^a is derived as follows:

$$\mathbf{P}_k^a = (\mathbf{I}_k - \mathbf{K}_k H_k) \mathbf{P}_k^f$$

with \mathbf{I}_k being the identity matrix of order k .

3.2 Ensemble Transform Kalman Filter with additive model error: ETKF-Q

Ensemble DA algorithms address downsides of the Kalman filter such as handling nonlinear models and storing and computing large matrices. For instance with a Kalman filter one has to store and manipulate error covariance matrices lying in $\mathbb{R}^{n \times n}$ which is often intractable in practice. Also, applying on both sides the model M_k in equation (9) to compute the forecast covariance matrix \mathbf{P}_k^f is prohibitively costly. Thus, ensembles enable to approximate the forecast covariance matrix thanks to a reduced set of sample vectors. In this section, we expose a tailored version of the widely used ensemble algorithm ETKF, namely ETKF-Q [18]. What we call ETKF-Q method precisely denotes the IEnKS-Q algorithm of [18, Algorithm 4.1] with parameters (L=0, K=0, S=1, G=0, one Gauss Newton loop, transform version).

Let us consider ensemble \mathbf{E}_k at time k such that: $\mathbf{E}_k = \{\mathbf{x}_k^1, \mathbf{x}_k^2, \dots, \mathbf{x}_k^m\} \in \mathbb{R}^{n \times m}$ where m is the number of members. Thus, we can empirically approximate its forecast covariance matrix:

$$\mathbf{P}_k^f = \frac{1}{m-1} \sum_{i=1}^m (\mathbf{x}_k^i - \bar{\mathbf{x}}_k) (\mathbf{x}_k^i - \bar{\mathbf{x}}_k)^T = \mathbf{X}_k^f \left(\mathbf{X}_k^f \right)^T \quad (10)$$

where superscript i denotes the i -th member of \mathbf{E}_k and $\bar{\mathbf{x}}_k$ the mean at time k (*i.e.* $\bar{\mathbf{x}}_k = \frac{1}{m} \sum_{i=1}^m \mathbf{x}_k^i$). As for $\mathbf{X}_k^f \in \mathbb{R}^{n \times m}$, it denotes the anomaly matrix such that $[\mathbf{X}_k^f]^i = \frac{\mathbf{x}_k^i - \bar{\mathbf{x}}_k}{\sqrt{m-1}}$.

Then the analysis \mathbf{x}_k^a can be written as:

$$\mathbf{x}_k^a = \bar{\mathbf{x}}_k + \mathbf{X}_k^f \mathbf{w}_k^a. \quad (11)$$

Substituting this equation into equation (6) and using Sherman-Morrison-Woodbury formula (see [3] for more details) yields:

$$\mathbf{w}_k^a = (\mathbf{I}_m + \mathbf{Y}_k \mathbf{R}^{-1} \mathbf{Y}_k)^{-1} \mathbf{Y}_k^T \mathbf{R}_k^{-1} \mathbf{d}_k \quad (12)$$

being an m -dimensional vector with $\mathbf{d}_k = \mathbf{y}_k - \overline{\mathcal{H}(\mathbf{x}_k)}$. In equation (12), \mathbf{Y}_k represents observation anomalies, *i.e.*

$$[\mathbf{Y}_k]^i = \frac{\mathcal{H}(\mathbf{x}_k^i) - \bar{\mathbf{y}}_k}{\sqrt{m-1}}$$

with $\bar{\mathbf{y}}_k = 1/m \sum_{i=1}^m \mathcal{H}(\mathbf{x}_k^i)$.

Note that the decomposition of \mathbf{x}_k^a given by equation (11) is not unique due to rank deficient matrix \mathbf{X}_k^f . This yields an ill-defined change of variables in the ensemble space that has to be fixed with the so called gauge-fixing term ([9]). As an alternative Fillion et al. [18] introduce deviation matrices to overcome this problem.

Definition Deviation matrix: a deviation matrix $\boldsymbol{\Delta}$ of a symmetric semi-definite positive matrix $\boldsymbol{\Sigma}$ is an **injective** factor verifying: $\boldsymbol{\Delta} \boldsymbol{\Delta}^T = \boldsymbol{\Sigma}$. A deviation matrix of an

ensemble is a deviation matrix of its sample covariance matrix.

Therefore, we aim to find a deviation matrix Δ_k of \mathbf{P}_k^f so that formulation in [equation \(11\)](#) yields a unique estimate \mathbf{x}_k^a . Hence, we apply [[18](#), Proposition 3.2]) to \mathbf{x}_k^a in order to ensure such a requirement. Since then, it exists a unique vector $\mathbf{w}_k^a \in \mathbb{R}^{m-1}$ such that:

$$\mathbf{x}_k^a = \bar{\mathbf{x}}_k + \Delta \mathbf{w}_k^a.$$

and

$$\begin{aligned} \mathbb{E}[\mathbf{w}_k^a] &= \mathbf{0}_{m-1}, \\ \mathbb{C}[\mathbf{w}_k^a] &= \mathbf{I}_{m-1}, \end{aligned}$$

where \mathbb{E} and \mathbb{C} are the expectation and covariance operator, respectively.

Remains the question of calculating a deviation matrix of \mathbf{P}_k^f . Again, we rely on another Fillion et al.'s proposition ([[18](#), Proposition 3.3]):

Proposition 1. (*Deviation matrix and ensemble construction*):

Let $n, m, l \in \mathbb{N}$ such that $n \geq m, l = m - 1$. Let $U_m \in \mathbb{R}^{m \times l}$ such that $\left[\frac{1}{\sqrt{m}} U_m\right] \in \mathbb{R}^{m \times m}$ be an orthonormal matrix. If $\mathbf{E} \in \mathbb{R}^{n \times m}$ is a full column rank ensemble, then the mean $\boldsymbol{\mu} \in \mathbb{R}^n$ and a deviation matrix $\Delta \in \mathbb{R}^{n \times l}$ of \mathbf{E} :

$$[\boldsymbol{\mu} \ \Delta] = \mathbf{E} \times \begin{bmatrix} \mathbf{1}_m & U_m \\ m & \sqrt{l} \end{bmatrix}. \quad (13)$$

Conversely, if $\boldsymbol{\mu} \in \mathbb{R}^n$ and $\Delta \in \mathbb{R}^{n \times l}$ then the ensemble $\mathbf{E} \in \mathbb{R}^{n \times m}$ defined by

$$\mathbf{E} = [\boldsymbol{\mu} \ \Delta] \times \begin{bmatrix} \mathbf{1}_m & \sqrt{l} U_m \end{bmatrix}^T \quad (14)$$

has $\boldsymbol{\mu}$ as sample mean and $\Delta \Delta^T$ as sample covariance matrix.

With [equation \(13\)](#), we can compute Δ_k , a deviation matrix of \mathbf{P}_k^f (we remind that $l = m - 1$):

$$\Delta_k = [x_k^1, x_k^2, \dots, x_k^m] \frac{U_m}{\sqrt{m-1}}$$

Regarding $U_m \in \mathbb{R}^{m \times (m-1)}$, it is a matrix such that $\left[\frac{1}{\sqrt{m}} U_m\right]$ is orthonormal (where $\mathbf{1}_m$ denotes the m -length vector $[1, 1, \dots, 1]^T$). It is worth mentioning that U_m can be constructed thanks to Householder's rotations.

When propagating through time, we know that our model \mathcal{M}_k is not perfect and has an intrinsic error denoted $\boldsymbol{\eta}_k$ (see [equation \(5\)](#)). However, up to now we have not included this particular knowledge in our analysis keeping the erroneous prediction as it is. Some approaches attempt to leverage this information in order to perform a model error correction and thus improve predictions' quality [[50](#), [39](#), [49](#), [36](#), [1](#)].

We now come to the core of the ETKF-Q algorithm, the variant of the ETKF one which takes model error into

account in the expression of the covariance matrix of \mathbf{x}_k (here \mathbf{x}_k denotes the real physical state):

$$\begin{aligned} \mathbf{x}_k &= \mathcal{M}_k(\mathbf{x}_{k-1}) + \boldsymbol{\eta}_k \\ \mathbb{C}[\mathbf{x}_k | \mathbf{y}_{0:k-1}] &= \mathbb{C}[\mathcal{M}_k(\mathbf{x}_{k-1}) + \boldsymbol{\eta}_k | \mathbf{y}_{0:k-1}] \end{aligned}$$

where $\mathbf{y}_{0:j}$ denotes the sequence of all the observations from time 0 to time j .

We have supposed that $\boldsymbol{\eta}_k \forall k \in [0, K]$ and \mathbf{x}_0 are mutually independent. Then, as $\mathcal{M}_k(\mathbf{x}_{k-1})$ is a function of \mathbf{x}_0 and of $\boldsymbol{\eta}_0, \boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_{k-1}$, it comes that $\mathcal{M}_k(\mathbf{x}_{k-1})$ and $\boldsymbol{\eta}_k$ are independent which yields that

$$\mathbb{C}[\mathbf{x}_k | \mathbf{y}_{0:k-1}] = \mathbb{C}[\mathcal{M}_k(\mathbf{x}_{k-1}) | \mathbf{y}_{0:k-1}] + \mathbb{C}[\boldsymbol{\eta}_k | \mathbf{y}_{0:k-1}]$$

We have also assumed that propagation and observation errors $\boldsymbol{\eta}_k$ and $\boldsymbol{\varepsilon}_k$ are mutually independent $\forall k \in [0, K]$. Then, we have $\mathbb{C}[\boldsymbol{\eta}_k | \mathbf{y}_{0:k-1}] = \mathbb{C}[\boldsymbol{\eta}_k] = \mathbf{Q}_k$. We get:

$$\mathbb{C}[\mathbf{x}_k | \mathbf{y}_{0:k-1}] = \mathbb{C}[\mathcal{M}_k(\mathbf{x}_{k-1}) | \mathbf{y}_{0:k-1}] + \mathbf{Q}_k$$

But, $\mathbb{C}[\mathcal{M}_k(\mathbf{x}_{k-1}) | \mathbf{y}_{0:k-1}]$ has been empirically approximated by $\mathbf{P}_k^f = \Delta_k \Delta_k^T$. Hence we obtain that

$$\mathbb{C}[\mathbf{x}_k | \mathbf{y}_{0:k-1}] \approx \Delta_k \Delta_k^T + \mathbf{Q}_k.$$

Deviation matrices of $\Delta_k \Delta_k^T + \mathbf{Q}_k$ are supposed to lie in $\mathbb{R}^{n \times n}$, but since a $n \times l$ deviation matrix is required for the next cycle, a reduction has to be performed. As $\Delta_k \Delta_k^T + \mathbf{Q}_k$ is symmetric (as a sum of symmetric matrices), its eigendecomposition by using the first $\ell (= m - 1)$ dominant eigenvectors yields $\mathbf{V}_k \in \mathbb{R}^{n \times (m-1)}$ and $\Lambda_k \in \mathbb{R}^{(m-1) \times (m-1)}$ such that:

$$\left(\Delta_k \Delta_k^T + \mathbf{Q}_k\right) \mathbf{V}_k \approx \mathbf{V}_k \Lambda_k$$

One could notice that this approximation is the best one in matrix Frobenius norm.

Therefore, a square root approximation of $\Delta_k \Delta_k^T + \mathbf{Q}_k$ is given by $\mathbf{V}_k \Lambda_k^{1/2}$. We hence update $\Delta_k = \mathbf{V}_k \Lambda_k^{1/2}$. Then, [equation \(14\)](#) enables to update ensemble \mathbf{E}_k according to this new statistic:

$$\mathbf{E}_k = \bar{\mathbf{x}}_k + \Delta_k \sqrt{m-1} U_m^T$$

Similarly, we apply [equation \(13\)](#) to the observation ensemble to produce \mathbf{Y}_k which is analogous to the observation anomalies in the regular ETKF algorithm:

$$\mathbf{Y}_k = [\mathcal{H}(\mathbf{x}_k^1), \mathcal{H}(\mathbf{x}_k^2), \dots, \mathcal{H}(\mathbf{x}_k^m)] \frac{U_m}{\sqrt{m-1}}$$

From now on, it is a straightforward application of the regular ETKF algorithm (see [[7](#), Section 5.3]). In [algorithm 1](#), we detail the ETKF-Q algorithm with the additional assumptions that $\mathbf{R}_k = \mathbf{R}$, $\mathbf{Q}_k = \mathbf{Q}$, $\forall k \in [0, K]$. In this algorithm, we mention that operator \mathcal{H} is a column-wise operator when applied to an ensemble, *i.e.* $\mathcal{H}(\mathbf{E}_k) = [\mathcal{H}(\mathbf{x}_k^1), \mathcal{H}(\mathbf{x}_k^2), \dots, \mathcal{H}(\mathbf{x}_k^m)]$.

Algorithm 1: ETKF-Q

Inputs:

Observation vector $\mathbf{y}_0 \in \mathbb{R}^p$;
 Observation operator $\mathcal{H} : \mathbb{R}^n \rightarrow \mathbb{R}^p$;
 Obs. error covariance matrix $\mathbf{R} \in \mathbb{R}^{p \times p}$;
 Ensemble $\mathbf{E}_0 = \{\mathbf{x}_0^1, \mathbf{x}_0^2, \dots, \mathbf{x}_0^m\} \in \mathbb{R}^{n \times m}$;
 Model operator $\mathcal{M} : \mathbb{R}^n \rightarrow \mathbb{R}^n$;
 Model error covariance matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$;
 Inflation parameter $\lambda \in \mathbb{R}$;

Initialization:

Construct U_m matrix such that $\begin{bmatrix} \frac{1}{\sqrt{m}} & U_m \end{bmatrix}$ is orthonormal ;

Define $\mathcal{U} := \begin{bmatrix} \frac{1}{m} & \frac{U_m}{\sqrt{m-1}} \end{bmatrix}$;

for $k = 1, 2, \dots$ **do**

Propagation step

1 $\mathbf{E}_k := \mathcal{M}(\mathbf{E}_{k-1}^a)$;
 2 $[\tilde{\mathbf{x}}_k \ \mathbf{\Delta}_k] := \mathbf{E}_k \times \mathcal{U}$;
 3 Calculate eigenpairs of $(\mathbf{\Delta}_k \mathbf{\Delta}_k^T + \mathbf{Q})$:
 $(\mathbf{\Delta}_k \mathbf{\Delta}_k^T + \mathbf{Q}) \mathbf{V}_k \approx \mathbf{V}_k \mathbf{\Lambda}_k$ with $\mathbf{V}_k \in \mathbb{R}^{n \times (m-1)}$
 and $\mathbf{\Lambda}_k \in \mathbb{R}^{(m-1) \times (m-1)}$;
 4 $\mathbf{\Delta}_k := \mathbf{V}_k \mathbf{\Lambda}_k^{1/2}$ (Update deviation matrix with
 model error) ;
 5 $\mathbf{E}_k := [\tilde{\mathbf{x}}_k \ \mathbf{\Delta}_k] \times \mathcal{U}^{-1}$ (Update ensemble with new
 statistics) ;

Analysis step

6 $[\tilde{\mathbf{y}}_k \ \mathbf{Y}_k] := \mathcal{H}(\mathbf{E}_k) \times \mathcal{U}$;
 7 Let $\mathbf{\Omega}_k \in \mathbb{R}^{(m-1) \times (m-1)}$ such that:
 $\mathbf{\Omega}_k \mathbf{\Omega}_k^T = (\mathbf{I}_{m-1} + \mathbf{Y}_k^T \mathbf{R}^{-1} \mathbf{Y}_k)^{-1}$;
 8 $\mathbf{w}_k^a := \mathbf{\Omega}_k \mathbf{\Omega}_k^T \mathbf{Y}_k \mathbf{R}^{-1} (\mathbf{y}_k - \tilde{\mathbf{y}}_k)$;
 9 $\mathbf{E}_k^a := \tilde{\mathbf{x}}_k \mathbf{1}_m^T + \lambda \times \mathbf{\Delta}_k (\mathbf{w}_k^a \mathbf{1}_m + \sqrt{m-1} \mathbf{\Omega}_k)$;

10 **end**

3.3 ETKF-Q-Latent algorithm

Our goal with the ETKF-Q-L algorithm is to perform DA analysis within the latent space of our autoencoder. Indeed, we now assume the existence of a N -dimensional system possessing a latent representation of lower dimension ℓ . From now on, variable N refers to the full space dimension whereas notation ℓ denotes the latent space dimension. [Algorithm 2](#) exposes the changes we made to do so. Here again operator \mathcal{H} is a column-wise operator when applied to an ensemble. We highlight that ensemble $\mathbf{E}_0 \in \mathbb{R}^{N \times m}$ is first encoded into ensemble $\mathbf{Z}_0 \in \mathbb{R}^{\ell \times m}$ and then all computations happen within the latent space. In order to calculate the misfit vector $\mathbf{d}_k = \mathbf{y}_k - \mathcal{H}(\mathbf{x}_k)$, first the decoder \mathcal{D} is used to map the ensemble from the latent space to the full space, then the observation operator \mathcal{H} maps the ensemble from the full space to the observation space. Therefore, we do not need to perform any operation with the observations. They stay in their original space. Since time propagation is performed in the latent space, we no longer refer to matrix \mathbf{Q} but we rather introduce \mathbf{Q}_ℓ . Instead of using $\mathbf{\Delta}$ to represent a deviation matrix, we refer to it as $\mathbf{\Gamma}$ in the case of the latent algorithm. For simplicity, we assume that $\mathbf{R} = \sigma_R^2 \mathbf{I}_p$ and $\mathbf{Q}_\ell = \sigma_{Q_\ell}^2 \mathbf{I}_\ell$. Since we cannot characterize the error committed by

Algorithm 2: ETKF-Q-Latent

Inputs:

Observation vector $\mathbf{y}_0 \in \mathbb{R}^p$;
 Observation operator $\mathcal{H} \circ \mathcal{D} : \mathbb{R}^\ell \rightarrow \mathbb{R}^N \rightarrow \mathbb{R}^p$;
 Obs. error covariance matrix $\mathbf{R} \in \mathbb{R}^{p \times p}$;
 Ensemble $\mathbf{E}_0 = \{\mathbf{x}_0^1, \mathbf{x}_0^2, \dots, \mathbf{x}_0^m\} \in \mathbb{R}^{N \times m}$;
 Surrogate model $\mathcal{S} : \mathbb{R}^\ell \rightarrow \mathbb{R}^\ell$;
 Model error covariance matrix $\mathbf{Q}_\ell \in \mathbb{R}^{\ell \times \ell}$;
 Encoder $\mathcal{E} : \mathbb{R}^N \rightarrow \mathbb{R}^\ell$;
 Inflation parameter $\lambda \in \mathbb{R}$;

Initialization:

Construct U_m matrix such that $\begin{bmatrix} \frac{1}{\sqrt{m}} & U_m \end{bmatrix}$ is orthonormal ;

Define $\mathcal{U} := \begin{bmatrix} \frac{1}{m} & \frac{U_m}{\sqrt{m-1}} \end{bmatrix}$;

Encode ensemble \mathbf{E}_0 : $\mathbf{Z}_0 := \mathcal{E}(\mathbf{E}_0) \in \mathbb{R}^{\ell \times m}$;

for $k = 1, 2, \dots$ **do**

Propagation Step

1 $\mathbf{Z}_k := \mathcal{S}(\mathbf{Z}_{k-1}^a)$;
 2 $[\tilde{\mathbf{z}}_k \ \mathbf{\Gamma}_k] := \mathbf{Z}_k \times \mathcal{U}$;
 3 Calculate eigenpairs of $(\mathbf{\Gamma}_k \mathbf{\Gamma}_k^T + \mathbf{Q}_\ell)$:
 $(\mathbf{\Gamma}_k \mathbf{\Gamma}_k^T + \mathbf{Q}_\ell) \mathbf{V}_k \approx \mathbf{V}_k \mathbf{\Lambda}_k$ with $\mathbf{V}_k \in \mathbb{R}^{\ell \times (m-1)}$
 and $\mathbf{\Lambda}_k \in \mathbb{R}^{(m-1) \times (m-1)}$;
 4 $\mathbf{\Gamma}_k := \mathbf{V}_k \mathbf{\Lambda}_k^{1/2}$ (Update deviation matrix with
 model error) ;
 5 $\mathbf{Z}_k := [\tilde{\mathbf{z}}_k \ \mathbf{\Gamma}_k] \times \mathcal{U}^{-1}$ (Update ensemble with new
 statistics) ;

Analysis step

6 $[\tilde{\mathbf{y}}_k \ \mathbf{Y}_k] := \mathcal{H}(\mathcal{D}(\mathbf{Z}_k)) \times \mathcal{U}$;
 7 Let $\mathbf{\Omega}_k \in \mathbb{R}^{(m-1) \times (m-1)}$ such that:
 $\mathbf{\Omega}_k \mathbf{\Omega}_k^T := (\mathbf{I}_{m-1} + \mathbf{Y}_k^T \mathbf{R}^{-1} \mathbf{Y}_k)^{-1}$;
 8 $\mathbf{w}_k^a := \mathbf{\Omega}_k \mathbf{\Omega}_k^T \mathbf{Y}_k \mathbf{R}^{-1} (\mathbf{y}_k - \tilde{\mathbf{y}}_k)$;
 9 $\mathbf{Z}_k^a := \tilde{\mathbf{z}}_k \mathbf{1}_m^T + \lambda \times \mathbf{\Gamma}_k (\mathbf{w}_k^a \mathbf{1}_m + \sqrt{m-1} \mathbf{\Omega}_k)$;

10 **end**

our surrogate network, the model correction step needs a tuned parameter σ_{Q_ℓ} that embodies the unknown surrogate error.

In order to picture the overall architecture of our DA framework, we can refer to [figure 3](#).

The strong difference between [algorithm 1](#) and [algorithm 2](#), is the reduction of the computational space from \mathbb{R}^N to \mathbb{R}^ℓ , which straightforwardly reduces both the computational cost and the memory storage. In practice, our latent space is 10 times smaller than our full space (see [section 4.4](#) for detailed results).

Numerical experiments also show an accuracy improvement when the assimilation is performed within the latent space. Indeed, decoding the analysis lying in the latent space outputs a state that is more likely to be on the system's trajectory while the linear analysis in the full space may deviate from it. [Figure 4](#) details how latent DA works compared to the regular DA: since a latent dynamics exists in \mathbb{R}^ℓ , latent DA leverages the nonlinear transformation provided by the encoder whereas full space DA might not capture the intrinsic dynamics and yields a poorer estimate.

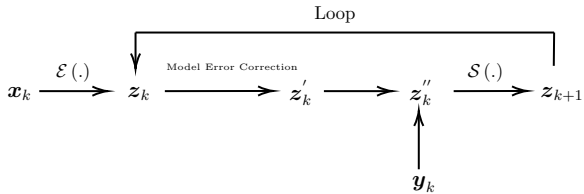


Figure 3: Outline of the Data Assimilation framework that includes the trained neural networks: z_k denotes the encoded input at time k , z'_k represents z_k corrected with respect to the model error, z''_k is the estimate yielded by assimilating observation y_k , and $\mathcal{E}(\cdot)$ and $\mathcal{S}(\cdot)$ denote the encoder and the surrogate operator, respectively.

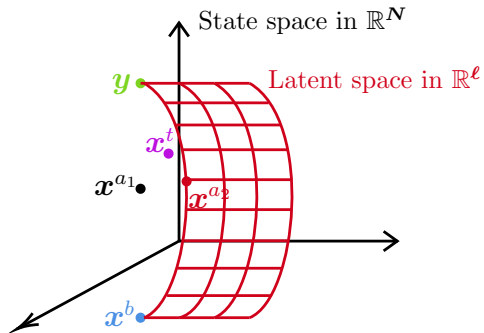


Figure 4: Comparison between full space DA and latent space DA (temporal subscript is dropped here). Variables x^b , x^{a1} , x^{a2} and x^t denote the background knowledge, the full space estimate, the latent estimate and the ground truth state, respectively. We also assume $\mathcal{H} = I_N$.

4 Numerical experiments

4.1 Choice of a physical system: The augmented Lorenz 96 model

The Lorenz 96 model [33] is widely used as a dynamical system [11, 63, 42] in ensemble data assimilation in particular for weather prediction.

It is defined as follows:

$$\frac{d\mathbf{x}[i]}{dt} = (\mathbf{x}_{[i+1]} - \mathbf{x}_{[i-2]}) \mathbf{x}_{[i-1]} - \mathbf{x}_{[i]} + F, \quad \forall i = 1, \dots, L \quad (15)$$

with $\mathbf{x}_{[-1]} = \mathbf{x}_{[L-1]}$, $\mathbf{x}_{[0]} = \mathbf{x}_{[L]}$, $\mathbf{x}_{[L+1]} = \mathbf{x}_{[1]}$, and $L \geq 4$. Subscript $[i]$ denotes the i -th variable.

In this equation, quadratic terms represent the advection that conserves the total energy, linear term represents the damping through which the energy decreases, and the constant term represents external forcing keeping the total energy away from zero. The L variables may be thought of as values of some atmospheric quantity in L sectors of a latitude circle. For $F = 8$, the system is known to have a chaotic behavior [33].

We consider the Lorenz 96 dynamics with $L = 40$ and then construct an augmented model based on this latent space representation. Doing so, we guarantee the existence of a latent space in which the observed dynamical system can be accurately expressed. Hence, by construc-

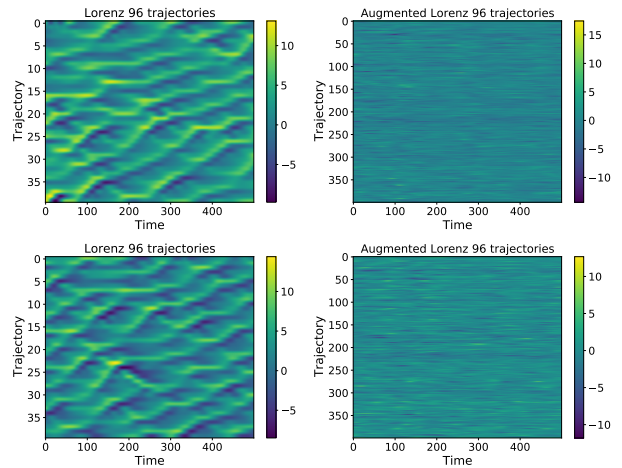


Figure 5: An illustration of two 40-variables Lorenz dynamics (top and bottom left) and their associated augmented Lorenz models in dimension 400 (top and bottom right). By construction, the augmented Lorenz systems has a latent representation in dimension 40.

tion our augmented Lorenz dynamics has a latent representation of dimension 40.

The definition of the augmented Lorenz system is given as follows:

Definition Augmented Lorenz model: it is a N -dimensional system for which there exists a function \mathcal{F} such that \mathcal{F} transforms the augmented Lorenz model into a ℓ -dimensional system ($\ell \ll N$) that follows Lorenz 96 equations.

Let us consider simulations of a $\ell (= 40)$ -dimensional Lorenz 96 state generated by integrating the well-known Lorenz equations (with a 4th order Runge-Kutta scheme for instance). Then, we apply a non-linear function $\mathbf{F} : \mathbb{R}^\ell \rightarrow \mathbb{R}^N$ to generate the augmented Lorenz 96 state of dimension N . In applying the nonlinear function, the ℓ -dimensional state is first mapped into a N -dimensional state via an orthonormal matrix $\mathbf{O} \in \mathbb{R}^{\ell \times N}$ and then an element-wise non-linear function \mathbf{f} (which is basically an invertible 3rd degree polynomial) is applied, i.e.

$$\mathbf{F} : \mathbf{x} \rightarrow \mathbf{f}(\mathbf{O}(\mathbf{x}))$$

Figure 5 shows two Lorenz 96 dynamics (generated from two very close but different initial points) in dimension 40 and their associated augmented Lorenz 96 models in \mathbb{R}^{400} . This figure helps in visualising how these dynamics behave in their respective spaces.

4.2 Networks' architectures

The encoder is made of 4 fully connected layers each one followed by a 0.2 slope *LeakyReLU* activation except for the last layer whose activation function is *tanh*. These 4 fully connected layers maps from \mathbb{R}^{400} to \mathbb{R}^{40} going by dimensions 300, 200 and 150. The decoder performs the reverse operation but there is no activation function in the last layer. Several variants have been tested but

these consecutive layers give the best achievements.

The surrogate network consists of 5 fully connected layers each one being followed by a 0.2 slope *LeakyReLU* activation except the last one. Data remain in \mathbb{R}^{40} through this network.

4.3 Neural Networks training

We have generated 1000 Lorenz 96 simulations with 500 time steps and 40 trajectories each, *i.e.* a thousand of 500×40 images. We precise that in this paper a trajectory denotes the time evolution of one variable $x[i]$ (see equation (15)), *i.e.* of a row in a 500×40 image, whereas a simulation represents the image itself.

All simulations come from the same distribution defined as follows (Python code):

```

1 X = np.zeros((time, nb_sim, dim))
2
3 for i in range(nb_sim):
4     x0 = np.random.randn(dim) * 0.01 + F0
5     initial_perturbation = np.random.randn(dim)
6     x0 += initial_perturbation
7     X[:, i] = new_trajectory(x0, dim, burn, time, F0, deltaT)

```

where $x0$ denotes the initial state, $F0$ the forcing term in Lorenz 96 system (denoted F in equation (15)), dim is the number of variables (here 40), $burn$ the number of burned states, $time$ the number of time steps, nb_sim the number of simulations (here 1000) and $deltaT$ represents a single time step that is 0.01 in our case.

We have then transformed the 1000 Lorenz 96 simulations into 1000 augmented Lorenz 96 ones, *i.e.* into a thousand 400-dimensional data.

The training set represents 95% of all available data and the test set the remaining ones. The batch size is set to 32 and we chose an Adam optimizer with a learning rate of 10^{-3} . The number of epochs is first arbitrarily fixed to 20 for computation time purposes. Then, once a network shows satisfactory performances, it is retrained with 40 epochs. Networks' weights are saved each time we reach a lower loss score on the testset. Also, it is worth reminding that both the autoencoder and the surrogate are trained together: parameter ρ of our custom loss function (see equation (4)) is set to 5 and the number of iterations C is set to 2. Figure 6 consists of four curves that validate NNs effective learning: indeed, the smooth exponential decreasing of the loss function (see top left graph) indicates that NNs are performing the task they are assigned to better and better over training. Second and third plots (top right and bottom left, respectively) confirm that both the AE and the surrogate do learn, *i.e.* that none of them is left behind during the training stage. Regarding the last plot (bottom right), it is close to the third one as it also measures the efficiency of the surrogate but without chaining, meaning that states are encoded, propagated only once and decoded afterwards. One could have noticed that \mathcal{L}_{AE} and \mathcal{L}_{Sur} have almost the same values over learning: it would suggest to set ρ to 1 rather than to 5 in order to define a fair loss function. However, it turns out that weighting more \mathcal{L}_{Sur} gives

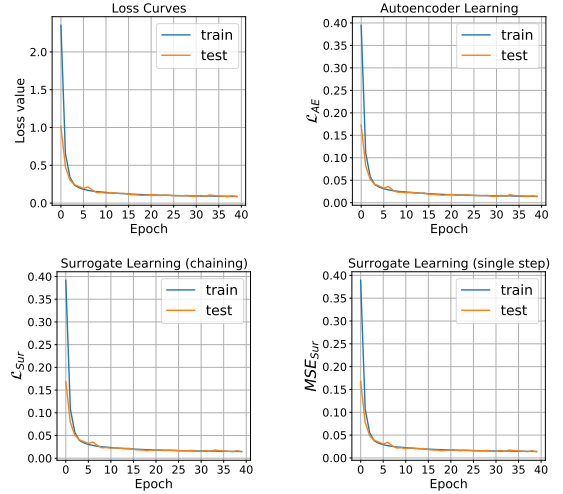


Figure 6: First plot (top left) represents the loss function (see equation (4)) over learning. Second and third plots (top right and bottom left, respectively) separately show the two parts of equation (4). Last plot (bottom right) measures surrogate's efficiency but without chaining.

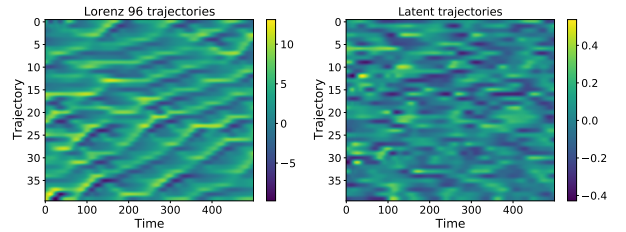


Figure 7: First graph (left) represents the original Lorenz 96 dynamics in dimension 40. Second graph (right) is the associated latent data yielded by the trained autoencoder.

better scores, meaning that more effort is needed for the surrogate to properly learn than for the AE. Along with the aforementioned stability issues, it confirms that the surrogate network is the more challenging to train.

In section 2.1, we pointed out the fact that the latent space produced by the autoencoder and the original 40-variables Lorenz 96 dynamics have absolutely no reason to be alike: indeed, we remind that no constraint is added to the autoencoder in this sense. Figure 7 confirms this comment by putting aside both the Lorenz 96 data in dimension 40 and the associated latent transformation: the autoencoder does not reproduce the original dynamics within its latent space. One could note that since the last activation function of the encoder is *tanh* we could not expect the two plots to have the same order of magnitude and range of values.

4.4 Assessing the performance of our DA framework: the ETKF-Q-L

We have proposed a new latent space DA algorithm coupled with an AE and a surrogate network. In this section we perform benchmark tests against the ETKF-Q-L algorithm presented in section 3.3. For the comparison, the

Root Mean Square Error (RMSE) is computed in the full space of dimension 400:

$$RMSE(\mathbf{x}_k, \mathcal{D}(\mathbf{z}_k)) = \sqrt{\frac{1}{N} \sum_{i=1}^N \left(\mathbf{x}_k^{(i)} - \mathcal{D}(\mathbf{z}_k^{(i)}) \right)^2} \quad (16)$$

where \mathcal{D} denotes the decoder, \mathbf{z}_k is the latent prediction at time k and \mathbf{x}_k is the truth at the same time.

Here is the list of the other approaches we benchmark our ETKF-Q-L method against:

- *ETKF-Q*: we perform the standard ETKF-Q algorithm [18] over the augmented Lorenz 96 data without resorting to any neural network. Propagation is performed by applying a standard Lorenz 96 propagator based on the Runge-Kutta fourth-order (RK4) scheme: this requires first to bring back the 400-dimensional data to the 40 dimensional space by applying transformation \mathcal{F} (see section 4.1). Data are mapped to the full space afterwards and the method loops back.
- *ETKF-Q-P (ETKF-Q-Physical)*: it remains pretty close to the regular ETKF-Q algorithm (see algorithm 1) as only the propagation step differs from it: instead of the standard propagator, we apply the encoder, the surrogate and the decoder, respectively. Notation *-P* indicates that data assimilation is performed within the full - *i.e.* physical - space.
- *PCA-S-P (Principal Component Analysis - Surrogate - Physical)*: it works exactly like *ETKF-Q-P* except that the encoder and the decoder are switched for a Principle Component Analysis (PCA) which can be seen as the simplest linear space reduction technique.
- *PCA-S-L (Principal Component Analysis - Surrogate - Latent)*: this approach and ours are very alike, the only difference is that encoding and decoding stages are performed with a PCA.
- *PCA-LinReg-P (Principal Component Analysis - Linear Regression - Physical)*: same as *PCA-S-P* but here the surrogate is replaced with the *scikit-learn* linear regression predictor.
- *PCA-LinReg-L (Principal Component Analysis - Linear Regression - Latent)*: same as *PCA-S-L* but here again the surrogate is replaced with the *scikit-learn* linear regression predictor.

Table 1 summarizes the benchmark context by specifying the DA space (*i.e.* the space where analysis is performed) along with the propagation method used.

In the standard ETKF algorithm, the multiplicative inflation parameter needs to be tuned. For simplicity, we chose \mathbf{R} and \mathbf{Q} as diagonal matrices: $\mathbf{R} = \sigma_R^2 \mathbf{I}_p$ and $\mathbf{Q} = \sigma_Q^2 \mathbf{I}_N$. Then, in the ETKF-Q algorithm [18], in addition to the multiplicative inflation parameter, there

Table 1: Data Assimilation benchmark context. DA space denotes the analysis space.

Name	DA space	Propagation	Correction
<i>ETKF-Q</i>	Full space	RK4	σ_Q
<i>ETKF-Q-P</i>	Full space	($\mathcal{D}, S, \mathcal{E}$)	σ_Q
<i>ETKF-Q-L</i>	Latent space	S	σ_{Q_ℓ}
<i>PCA-S-P</i>	Full space	(<i>PCA</i> , S , <i>reverse PCA</i>)	σ_Q
<i>PCA-S-L</i>	Latent space	S	σ_{Q_ℓ}
<i>PCA-LinReg-P</i>	Full space	(<i>PCA</i> , <i>LinReg</i> , <i>reverse PCA</i>)	σ_Q
<i>PCA-LinReg-L</i>	Latent space	<i>LinReg</i>	σ_{Q_ℓ}

is another parameter namely the standard deviation of the additive model error, σ_Q . The trajectories of the ensemble states are corrected using this model error at the beginning of each cycle (see algorithm 1 and algorithm 2). These updated states are used as the forecast states. However, due to the sampling error, this model error at each step can be thought of as an additive type of inflation. Thus, this model error parameter may correct the sampling error as well. With this idea in mind, the values of the model error during propagation and the update of the deviation matrix may be different. According to our experiments, indeed taking different values gives better RMSE scores.

As for the observation operator \mathcal{H} , it is set to \mathbf{I}_p , meaning that observation and state spaces are the same *i.e.* $N = p$.

For the latent space algorithms, we introduce $\mathbf{Q}_\ell = \sigma_{Q_\ell}^2 \mathbf{I}_\ell$. Since we do not know the propagation error of the surrogate model, we need to iteratively test dozens of σ_{Q_ℓ} to find the one that best suits.

After some hand-made experiments, we decided to find the best parameters combination in a straightforward manner through a grid search. The context of experiment is as follows:

- Parameters:
 - 40 ensemble members
 - Observation error covariance matrix $\sigma_R^2 \mathbf{I}_{400}$ with $\sigma_R = \mathbf{1.0}$.
 - Initial forecast error covariance matrix $\sigma_B^2 \mathbf{I}_{400}$ at time $k = 0$ with $\sigma_B = \mathbf{0.3}$.
 - 1000 iterations/time steps.
 - Model error covariance matrix $\sigma_M^2 \mathbf{I}_{40}$ (for *ETKF-Q* method only): $\sigma_M = \mathbf{0.3}$.
- Grid search:
 - *inflation* ranges from **0.99** to **1.9**.
 - σ_Q/σ_{Q_ℓ} ranges from 10^{-7} to **0.9**.

We conducted an experiment that consists in tuning *inflation* and σ_Q or σ_{Q_ℓ} parameters. Distinction between σ_Q and σ_{Q_ℓ} only aims at differentiating full space data assimilation from latent space data assimilation, respectively. Results are given in table 2.

They draw the conclusion that our approach is competitive with other methods that rely on a simple space

Table 2: Data Assimilation results (mean RMSE). Only two parameters are tuned: *inflation* and σ_Q/σ_{Q_ℓ} .

Name	RMSE	Inflation	σ_Q/σ_{Q_ℓ}
<i>ETKF-Q</i>	0.194	1.12	0.07
<i>ETKF-Q-P</i>	0.217	1.08	0.1
<i>ETKF-Q-L</i>	0.168	1.004	5.10^{-5}
<i>PCA-S-P</i>	0.383	1.145	0.5
<i>PCA-S-L</i>	0.383	1.13	0.6
<i>PCA-LinReg-P</i>	0.429	1.24	0.7
<i>PCA-LinReg-L</i>	0.426	1.2	0.9

reduction technique like PCA or on a very straightforward propagator such as the linear predictor from the *scikit-learn* library. As best RMSE score is reached for our solution, it suggests that the autoencoder fully leverages nonlinear transformations to compress the information hold by the augmented Lorenz 96 dynamics (see figure 4). Besides, it proves that for this particular tailored system, a mere Multi-Layer Perceptron (MLP) can properly perform the time propagation step in a suitable latent space: more precisely, the experiments do not show that a simple MLP can propagate through time, but it rather demonstrates that surrounded with an encoder and a decoder, a MLP can replace the model. Our training probably finds a latent space where the MLP surrogate can perform well the time propagation. Although the proposed approach seems both simple and efficient, we remind that it was nonetheless not straightforward and we needed to resort to a custom loss function with an iterative training scheme that enhances stability in order to achieve a satisfactory solving.

As priorly mentioned in the introduction, our method also aims at tackling time computing issues faced by regular DA algorithms. Therefore, all our benchmark algorithms have been run 100 times on a virtual machine composed of four virtual CPUs, one NVIDIA Tesla P100 GPU and 15 GB of RAM. Table 3 evidences that our method computes 2.4 to 2.5 times faster than the standard ETKF-Q algorithm. More generally, we notice that as long as an algorithm utilizes a latent space structure to perform DA analysis, a significant computational gain is obtained. We do not observe a major difference between GPU and CPU computations except for *ETKF-Q-P* algorithm. Yet, standard deviations seem to be larger when using CPU. Crossing these results with table 2 clearly reveals that our approach is the best one in our benchmark context on both the accuracy and the computational cost criteria.

5 Conclusion and Outlook

Broadening DL algorithms to solve physics is justified by its efficiency in space reduction and time propagation tasks. Whereas prior aims like performing image classification, segmentation or language processing have been addressed by NNs solely, today’s challenges involve using DL in more global processes. In this paper, we proposed to study how to incorporate DL in a DA framework. For

Table 3: GPU and CPU average computational times over 100 runs for all the algorithms tested in this case study. Resorting to a GPU device is only meaningful for methods that include NNs. Std denotes the standard deviation.

Name	GPU		CPU	
	Avg. Time	Std	Avg. Time	Std
<i>ETKF-Q</i>	16.32s	0.20s	16.23s	0.31s
<i>ETKF-Q-P</i>	13.60s	0.19s	17.12s	0.84s
<i>ETKF-Q-L</i>	6.52s	0.12s	6.89s	0.62s
<i>PCA-S-P</i>	13.22s	0.18s	12.18s	0.31s
<i>PCA-S-L</i>	5.93s	0.10s	5.35s	0.16s
<i>PCA-LinReg-P</i>	12.09s	0.29s	11.62s	0.27s
<i>PCA-LinReg-L</i>	5.02s	0.12s	4.94s	0.15s

this proof of concept, we supposed that we were given a dynamical system lying in \mathbb{R}^N exhibiting a latent representation of lower dimension ℓ . Under this assumption, we proved that it is possible to perform the ETKF-Q algorithm within a latent space of same size ℓ produced by an autoencoder. While performing the ETKF-Q algorithm in a latent space, we take proper account of the necessary model calibration needed in the evolution model description in latent space.

The motivation for performing the latent space DA was both reducing the computational cost and also getting a better accuracy. Reducing the computational cost is a natural gain due to performing DA in a reduced space. The accuracy gain on the other hand depends on the fact that DA linear analysis in the latent space obtained by the AE is less susceptible to yield non-physical solutions.

We have therefore trained an AE and a surrogate network through a single learning thanks to a chained custom loss function. In addition to allowing a better training since both the AE and the surrogate modify their weights according to each other, this particular loss function also enhances stability. Indeed, the surrogate is called several times in a row and thus has to produce stable results at least for two successive time steps. Experience demonstrates that this is already enough to get a quite satisfactory stable behaviour on longer time windows. Then, given these two networks we can perform latent DA.

We have shown the potential of our methodology on the instructional augmented Lorenz 96 system which is designed such that we ensure the existence of a latent dynamics. We have compared our methodology to the existing ETKF-Q algorithm and numerical results confirmed that the proposed methodology performs better than the usual full space strategy. Our methodology can be utilized as long as a system lying in \mathbb{R}^N is accurately representable in a lower dimension.

We believe that the proposed proof of concept is encouraging and that, as such, it stimulates several tracks for future research. For instance, to before considering a use in more operational situations, a number of theoretical question should be considered. For instance, it would be interesting to further investigate the sensibility of the methodology to several hyper-parameters of the method including the latent space dimension, obser-

vation frequency, inflation etc. It is also important to better understand how the properties of the full model dynamics behave in the latent space including the model error. We remind that in this study we consider very simple NN architectures. Investigating more sophisticated networks and introducing constraints to enforce latent space’s structure could be of course important. This structure may be motivated by underlying properties of the physics, such as the coupling between variables, or other structures. This being done, considering other toy problems would necessary before going for more complex problems.

Finally, the latent space strategy that we propose in this paper for an ensemble Kalman filter type method is quite general in the sense that it can be easily adapted to other DA algorithms like the variational approaches.

Acknowledgements

We are very thankful to Atos, ANITI, CERFACS, University of Toulouse, and NVIDIA for supporting and being involved in Mathis Peyron’s PhD. They have provided us with computational resources, materials, funds and experts’ advice that have been core in this research.

References

- [1] J. Amezcua, M. Goodliff, and P. J. V. Leeuwen. A weak-constraint 4densemblevar. part i: formulation and simple model experiments. *Tellus A: Dynamic Meteorology and Oceanography*, 69(1):1271564, 2017.
- [2] G. Artana, A. Cammilleri, J. Carlier, and E. Mémin. Strong and weak constraint variational assimilations for reduced order fluid flow modeling. *Journal of Computational Physics*, 213(8):3264–3288, Apr. 2012.
- [3] M. Asch, M. Bocquet, and M. Nodet. *Data assimilation: methods, algorithms, and applications*, chapter 1: Introduction to data assimilation and inverse problems, pages 3–23. Fundamentals of Algorithms. SIAM, 2016.
- [4] M. Asch, M. Bocquet, and M. Nodet. *Data assimilation: methods, algorithms, and applications*. Fundamentals of Algorithms. SIAM, 2016.
- [5] T. Bachlechner, B. P. Majumder, H. H. Mao, G. W. Cottrell, and J. McAuley. Rezero is all you need: Fast convergence at large depth, 2020.
- [6] R. N. Bannister. A review of forecast error covariance statistics in atmospheric variational data assimilation. i: Characteristics and measurements of forecast error covariances. *Quarterly Journal of the Royal Meteorological Society*, 134(637):1951–1970, 2008.
- [7] M. Bocquet. Lecture notes, 2014, last revision: January 2019.
- [8] M. Bocquet and A. Carrassi. Four-dimensional ensemble variational data assimilation and the unstable subspace. *Tellus A*, 69(1):1304504, Mar. 2017.
- [9] M. Bocquet and P. Sakov. An iterative ensemble kalman smoother. *Quarterly Journal of the Royal Meteorological Society*, 140(682):1521–1535, 2014.
- [10] S.-A. Boukabara, V. Krasnopolsky, S. G. Penny, J. Q. Stewart, A. McGovern, D. Hall, J. E. T. Hoeve, J. Hickey, H.-L. A. Huang, J. K. Williams, K. Ide, P. Tissot, S. E. Haupt, K. S. Casey, N. Oza, A. J. Geer, E. S. Maddy, and R. N. Hoffman. Outlook for exploiting artificial intelligence in the earth and environmental sciences. *Bulletin of the American Meteorological Society*, pages 1 – 53, 20 Nov. 2020.
- [11] 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, Jul 2020.
- [12] S. W. Canchumuni, A. A. Emerick, and M. A. C. Pacheco. Towards a robust parameterization for conditioning facies models using deep variational autoencoders and ensemble smoother. *Computers & Geosciences*, 128:87–102, Jul 2019.
- [13] Y. Cao, J. Zhu, I. M. Navon, and Z. Luo. A reduced-order approach to four-dimensional variational data assimilation using proper orthogonal decomposition. *International Journal for Numerical Methods in Fluids*, 53(10):1571–1583, 2007.
- [14] N. Carlini and D. A. Wagner. Towards evaluating the robustness of neural networks. *CoRR*, abs/1608.04644, 2016.
- [15] A. Carrassi, A. Trevisan, L. Descamps, O. Talagrand, and F. Uboldi. Controlling instabilities along a 3dvar analysis cycle by assimilating in the unstable subspace: a comparison with the enkf. *Nonlinear Processes in Geophysics*, 15(4):503–521, 2008.
- [16] A. Carrassi, A. Trevisan, and F. Uboldi. Adaptive observations and assimilation in the unstable subspace by breeding on the data-assimilation system. *Tellus A*, 59(1):101–113, 2007.
- [17] S. E. Cohn and R. Todling. Approximate data assimilation schemes for stable and unstable dynamics. *Journal of the Meteorological Society of Japan. Ser. II*, 74(1):63–75, 1996.
- [18] A. Fillion, M. Bocquet, S. Gratton, S. Gürol, and P. Sakov. An iterative ensemble kalman smoother in presence of additive model error. *SIAM/ASA Journal on Uncertainty Quantification*, 8(1):198–228, 2020.
- [19] L. Fulton, V. Modi, D. Duvenaud, D. I. W. Levin, and A. Jacobson. Latent-space dynamics for reduced deformable simulation. *Computer Graphics Forum*, 2019.
- [20] A. Garcia-Garcia, S. Orts-Escolano, S. Oprea, V. Villena-Martinez, P. Martinez-Gonzalez, and J. Garcia-Rodriguez. A survey on deep learning techniques for image and video semantic segmentation. *Applied Soft Computing*, 70:41–65, 2018.

- [21] M. Gardner, J. Grus, M. Neumann, O. Tafjord, P. Dasigi, N. F. Liu, M. Peters, M. Schmitz, and L. Zettlemoyer. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia, July 2018. Association for Computational Linguistics.
- [22] A. J. Geer. Learning earth system models from observations: machine learning or data assimilation? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194):20200089, 2021.
- [23] E. Haber, K. Lensink, E. Treister, and L. Ruthotto. Imxnet: A forward stable deep neural network, 2019.
- [24] E. Haber and L. Ruthotto. Stable architectures for deep neural networks. *Inverse Problems*, 34(1):014004, dec 2017.
- [25] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, Aug 2018.
- [26] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [27] Y. Khoo, J. Lu, and L. Ying. Solving parametric pde problems with artificial neural networks. *European Journal of Applied Mathematics*, page 1–15, 2020.
- [28] M. A. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2):233–243, 1991.
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 1097–1105. Curran Associates, Inc., 2012.
- [30] C. Kuenzer, M. Ottinger, M. Wegmann, H. Guo, C. Wang, J. Zhang, S. Dech, and M. Wikelski. Earth observation satellite sensors for biodiversity monitoring: potentials and bottlenecks. *International Journal of Remote Sensing*, 35(18):6599–6647, 2014.
- [31] Q. V. Le, M. Ranzato, R. Monga, M. Devin, K. Chen, G. S. Corrado, J. Dean, and A. Y. Ng. Building high-level features using large scale unsupervised learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML’12*, page 507–514, Madison, WI, USA, 2012. Omnipress.
- [32] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar. Fourier Neural Operator for Parametric Partial Differential Equations. *arXiv e-prints*, page [arXiv:2010.08895](https://arxiv.org/abs/2010.08895), Oct. 2020.
- [33] E. Lorenz. Predictability: a problem partly solved. In *Seminar on Predictability, 4-8 September 1995*, volume 1, pages 1–18, Shinfield Park, Reading, 1995. ECMWF, ECMWF.
- [34] Z. Lu, T. K. Leen, R. van der Merwe, S. Frolov, and A. M. Baptista. Sequential data assimilation with sigma-point kalman filter on low-dimensional manifold. *submitted to Journal of Inverse Problems*, 2007.
- [35] J. Mack, R. Arcucci, M. Molina-Solana, and Y. Guo. Attention-based convolutional autoencoders for 3d-variational data assimilation. *Computer Methods in Applied Mechanics and Engineering*, 372:113291, 2020.
- [36] J. Mandel, E. Bergou, S. Gürol, S. Gratton, and I. Kusanický. Hybrid levenberg–marquardt and weak-constraint ensemble kalman smoother method. *Nonlinear Processes in Geophysics*, 23(2):59–73, 2016.
- [37] R. Maulik, B. Lusch, and P. Balaprakash. Reduced-order modeling of advection-dominated systems with recurrent neural networks and convolutional autoencoders. *Physics of Fluids*, 33(3):037106, 2021.
- [38] R. N. Miller and M. A. Cane. A kalman filter analysis of sea level height in the tropical pacific. *Journal of Physical Oceanography*, 19(6):773 – 790, 01 Jun. 1989.
- [39] L. Mitchell and A. Carrassi. Accounting for model error due to unresolved scales within ensemble kalman filtering. *Quarterly Journal of the Royal Meteorological Society*, 141(689):1417–1428, 2015.
- [40] A. Mohan, N. Lubbers, D. Livescu, and M. Chertkov. Embedding hard physical constraints in neural network coarse-graining of 3d turbulence. *arXiv: Computational Physics*, 2020.
- [41] D. W. Otter, J. R. Medina, and J. K. Kalita. A survey of the usages of deep learning for natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [42] S. M. Pawar and O. San. Data assimilation empowered neural network parameterizations for subgrid processes in geophysical flows. *arXiv: Computational Physics*, 2020.
- [43] C. Qi, H. Su, K. Mo, and L. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017.
- [44] M. Raissi, P. Perdikaris, and G. Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *ArXiv*, abs/1711.10561, 2017.
- [45] M. Raissi, P. Perdikaris, and G. Karniadakis. Physics informed deep learning (part ii): Data-driven discovery of nonlinear partial differential equations. *ArXiv*, abs/1711.10566, 2017.
- [46] M. Raissi, Z. Wang, M. S. Triantafyllou, and G. E. Karniadakis. Deep learning of vortex-induced vibrations. *Journal of Fluid Mechanics*, 861:119–137, Dec 2018.

- [47] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat. Deep learning and process understanding for data-driven earth system science. *Nature*, 566:195–204, 2019.
- [48] M. Rottmann, K. Maag, M. Peyron, N. Krejic, and H. Gottschalk. Detection of iterative adversarial attacks via counter attack, 2021.
- [49] P. Sakov and M. Bocquet. Asynchronous data assimilation with the enfk in presence of additive model error. *Tellus A: Dynamic Meteorology and Oceanography*, 70(1):1–7, 2018.
- [50] P. Sakov, J.-M. Haussaire, and M. Bocquet. An iterative ensemble kalman filter in the presence of additive model error. *Quarterly Journal of the Royal Meteorological Society*, 144(713):1297–1309, 2018.
- [51] S. Seo and Y. Liu. Differentiable physics-informed graph networks. *ArXiv*, abs/1902.02950, 2019.
- [52] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [53] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [54] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.
- [55] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14, page 3104–3112, Cambridge, MA, USA, 2014. MIT Press.
- [56] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013.
- [57] A. Tangborn and S. Q. Zhang. Wavelet transform adapted to an approximate kalman filter system. *Applied Numerical Mathematics*, 33(1):307–316, 2000.
- [58] A. Trevisan, M. D’Isidoro, and O. Talagrand. Four-dimensional variational assimilation in the unstable subspace and the optimal subspace dimension. *Quarterly Journal of the Royal Meteorological Society*, 136(647):487–496, 2010.
- [59] A. Trevisan and L. Palatella. On the kalman filter error covariance collapse into the unstable subspace. *Nonlinear Processes in Geophysics*, 18(2):243–250, 2011.
- [60] A. Trevisan and F. Uboldi. Assimilation of standard and targeted observations within the unstable subspace of the observation–analysis–forecast cycle system. *Journal of the Atmospheric Sciences*, 61(1):103–113, 01 Jan. 2004.
- [61] D. Tuan Pham, J. Verron, and M. Christine Roubaud. A singular evolutive extended kalman filter for data assimilation in oceanography. *Journal of Marine Systems*, 16(3):323–340, 1998.
- [62] F. Uboldi and A. Trevisan. Detecting unstable structures and controlling error growth by assimilation of standard and adaptive observations in a primitive equation ocean model. *Nonlinear Processes in Geophysics*, 13(1):67–81, 2006.
- [63] P. R. Vlachas, W. Byeon, Z. Y. Wan, T. P. Sapsis, and P. Koumoutsakos. Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2213):20170844, May 2018.
- [64] L. von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, J. Pfrommer, A. Pick, R. Ramamurthy, M. Walczak, J. Garcke, C. Baukhage, and J. Schücker. Informed machine learning – a taxonomy and survey of integrating knowledge into learning systems. *arXiv: Machine Learning*, 2019.
- [65] S. Wiewel, M. Becher, and N. Thuerey. Latent-space physics: Towards learning the temporal evolution of fluid flow. *CoRR*, abs/1802.10123, 2018.
- [66] S. Wiewel, B. Kim, V. C. Azevedo, B. Solenthaler, and N. Thuerey. Latent space subdivision: Stable and controllable time predictions for fluid flow. *Computer Graphics Forum*, 39(8):15–25, 2020.
- [67] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144, 2016.