

*“Multifidelity Ensemble Kalman Filtering
Using Surrogate Models Defined by
Physics-Informed Autoencoders”*

Andrey A. Popov, Adrian Sandu

Computational Science Laboratory
Report CSL-TR-21-1

Computational Science Laboratory
“Compute the Future!”

Department of Computer Science
Virginia Tech
Blacksburg, VA 24060
Phone: (540) 231-2193
Fax: (540) 231-6075

Email: apopov@vt.edu, sandu@cs.vt.edu
Web: <http://csl.cs.vt.edu>



MULTIFIDELITY ENSEMBLE KALMAN FILTERING USING SURROGATE MODELS DEFINED BY PHYSICS-INFORMED AUTOENCODERS

Andrey A. Popov^{1,} & Adrian Sandu¹*

2202 Kraft Drive, Blacksburg, VA, 24060

*Address all correspondence to: Andrey A. Popov, E-mail: apopov@vt.edu

Original Manuscript Submitted: mm/dd/yyyy; Final Draft Received: mm/dd/yyyy

Data assimilation is a Bayesian inference process that obtains an enhanced understanding of a physical system of interest by fusing information from an inexact physics-based model, and from noisy sparse observations of reality. The multifidelity ensemble Kalman filter (MFEnKF) recently developed by the authors combines a full-order physical model and a hierarchy of reduced order surrogate models in order to increase the computational efficiency of data assimilation. The standard MFEnKF uses linear couplings between models, and is statistically optimal in case of Gaussian probability densities. This work extends MFEnKF to work with non-linear couplings between the models. Optimal nonlinear projection and interpolation operators are obtained by appropriately trained physics-informed autoencoders, and this approach allows to construct reduced order surrogate models with less error than conventional linear methods. Numerical experiments with the canonical Lorenz '96 model illustrate that nonlinear surrogates perform better than linear projection-based ones in the context of multifidelity filtering.

KEY WORDS: *Bayesian inference, control variates, data assimilation, multifidelity ensemble Kalman filter, reduced order modeling, machine learning, surrogate models*

Contents

1	Introduction	1
2	Background	2
2.1	Control variates	2
2.2	Linear control variates with non-linear transformations	4
2.3	The Ensemble Kalman Filter	5
2.4	The Multifidelity Ensemble Kalman Filter	5
2.5	Projection-based reduced order models	8
2.6	Autoencoders	8
3	Non-linear projection-based MFEnKF	9
3.1	NL-MFEnKF Heuristic Corrections	10
3.1.1	Control space unbiased mean adjustment	10
3.1.2	Principal space unbiased mean adjustment	10

3.1.3	Kalman approximate mean adjustment	11
3.2	Telescopic extension	11
4	Models for numerical experiments	11
4.1	Lorenz '96	12
4.1.1	Data for constructing reduced-order models.	12
4.2	Proper Orthogonal decomposition ROM	12
4.3	Autoencoder based ROM	12
4.3.1	Lorenz '96 surrogate	14
5	Numerical experiments	14
5.1	Accuracy of ROM models	14
5.2	Impact of ROM dimension	15
5.3	Ensemble Size and Inflation	16
6	Conclusions	17

1. INTRODUCTION

Data assimilation (Asch et al., 2016; Reich and Cotter, 2015) is a Bayesian inference process that fuses information obtained from an inexact physics-based model, and from noisy sparse observations of reality, in order to enhance our understanding of a physical process of interest. The reliance on physics-based models distinguishes data assimilation from traditional machine learning methodologies, which aim to learn the quantities of interest through purely data-based approaches. From the perspective of machine learning, data assimilation is a learning problem where the quantity of interest is constrained by prior physical assumptions, as captured by the model, and nudged towards the optimum solution by small amounts of data from imperfect observations. Therefore data assimilation can be considered a form of physics-constrained machine learning (Karpatne et al., 2017; Willard et al., 2020). This work improves data assimilation methodologies by combining a mathematically rigorous data assimilation approach and a data rigorous machine learning algorithm through powerful techniques in multilevel inference (Giles, 2008, 2015).

The ensemble Kalman filter (Burgers et al., 1998; Evensen, 1994, 2009) (EnKF) is a family of computational methods that tackle the data assimilation problem using Gaussian assumptions, and a Monte Carlo approach where the underlying probability densities are represented by ensemble of model state realizations. The ensemble size, i.e., the number of physics-based model runs, is typically the main factor that limits the efficiency of EnKF. For increasing the quality of the results when ensembles are small, heuristics correction methods such as covariance shrinkage (Nino-Ruiz and Sandu, 2015, 2018, 2019) and localization (Moosavi et al., 2019; Petrie, 2008; Popov and Sandu, 2019) have been developed. As some form of heuristic correction is required for operation implementations of the ensemble Kalman filter, reducing the need for such heuristic corrections in operational implementations is an important and active area of research.

The dominant cost in operational implementations of EnKF is the large number of expensive high fidelity physics-based model runs, which we refer to as “full order models” (FOM). A natural approach to increase efficiency is to endow the data assimilation algorithms with the ability to make use of inexpensive but less accurate model runs (Cao et al., 2007; Farrell and Ioannou, 2001), which we refer to as “reduced order models” (ROMs). ROMs are constructed to capture the most important aspects of the dynamics of the FOM, at a fraction of the computational cost; typically they use a much smaller number of variables than the corresponding FOM. The idea of leveraging model hierarchies in numerical algorithms for uncertainty quantification (Peherstorfer et al., 2018) is fast gaining traction in both the data assimilation and machine learning communities. Here we focus on two particular types of ROMs: a proper orthogonal decomposition (POD) based ROM, corresponding to a linear projection of the FOM dynamics onto a small linear subspace (Brunton and Kutz, 2019), and a ROM based on autoencoders (Aggarwal et al., 2018), which are neural networks built for optimal dimensionality reduction, corresponding to a non-linear projection of the dynamics onto a small dimensional manifold.

The multifidelity ensemble Kalman filter (MFEnKF) (Popov et al., 2021; Popov and Sandu, 2021) is a correction and extension to the multilevel EnKF (Chada et al., 2020; Chernov et al., 2021; Hoel et al., 2016, 2019) that combines the ensemble Kalman filter with the idea of surrogate modeling. The MFEnKF attempts to utilize both the full-order and a reduced order surrogate model in order to optimally combine the information obtained from model runs with information begotten from the observations. By posing the problem in terms of a mathematically rigorous variance reduction technique—the linear control variate framework—MFEnKF is able to provide robust guarantees about the accuracy of the inference results.

While numerical weather prediction is the dominant driver of innovation in data assimilation literature (Kalnay, 2003), other applications can benefit from our multifidelity approach such as mechanical engineering (Blanchard et al., 2009, 2010a,b) and air quality modeling (Constantinescu et al., 2007a,b,c).

This paper is organized as follows. Section 2 discusses the data assimilation problem, provides background on control variates, the EnKF, and the MFEnKF, as well as ROMs and autoencoders. Section 3 introduces the non-linear extension to the MFEnKF. Section 4 presents the Lorenz '96 model and the corresponding POD-ROM and NN-ROM surrogates. Section 5 provides the results of numerical experiments. Concluding remarks are made in Section 6.

2. BACKGROUND

As the aphorism goes “all models are wrong, but some are useful”, Sequential data assimilation propagates our imperfect knowledge about some physical quantity of interest through an imperfect model which represents a time-evolving dynamical system, typically a chaotic system (Strogatz, 2018). Without an additional influx of external information, our knowledge about the systems would rapidly degrade, therefore acquiring and assimilating imperfect external information about the quantity of interest is the core of the problem that we are interested in.

To be specific, consider a dynamical system of interest whose true state at time t_i is X_i^t . The evolution of the system is captured by the dynamical model

$$X_i = \mathcal{M}_{t_{i-1}, t_i}(X_{i-1}) + \Xi_i, \quad (2.1)$$

where X_i is a random variable whose distribution describes our knowledge of the the state of a physical process at time index i , and Ξ_i is a random variable describing the modeling error. In this paper we assume a perfect model ($\Xi_i \equiv 0$), as the discussion of model error in multifidelity methods is significantly outside the scope of this paper.

Additional independent information about the system is obtained through imperfect measurements of the observable aspects Y_i of the truth X_i^t , i.e., through noisy observations

$$Y_i = \mathcal{H}(X_i^t) + \eta_i, \quad \eta_i \sim \mathcal{N}(0, \Sigma_{\eta_i, \eta_i}), \quad (2.2)$$

where the “observation operator” $\mathcal{H}$ maps the model space onto the observation space (i.e., selects the observable aspects of the state).

Our aim is to combine the two sources of information in a Bayesian framework:

$$\pi(X_i|Y_i) \propto \pi(Y_i|X_i) \pi(X_i). \quad (2.3)$$

In the remainder of the paper we use the following notation. Let W and V be random variables. The exact mean of W is denoted by μ_W , and the exact covariance between W and V by $\Sigma_{W,V}$. E_W denotes an ensemble of samples of W , and $\tilde{\mu}_W$ and $\tilde{\Sigma}_{W,V}$ are the empirical ensemble mean of W and empirical ensemble covariance of W and V , respectively.

2.1 Control Variates

Bayesian inference requires that all available information is used in order to obtain correct results (Jaynes, 2003). Variance reduction techniques (Owen, 2013) are methods that provide estimates of some quantity of interest with lower variability. From a Bayesian perspective they represent

a reintroduction of information that was previously ignored. The linear control variate (LCV) approach (Rubinstein and Marcus, 1985) is a variance reduction method that aims to incorporate the new information in an optimal linear sense.

LCV works with two vector spaces, a principal space $\mathbb{X}$ and a control space $\mathbb{U}$, and several random variables, as follows. The principal variate X is a $\mathbb{X}$ -valued random variable, and the control variate $\hat{U}$ is a highly correlated $\mathbb{U}$ -valued random variable. The ancillary variate U is a $\mathbb{U}$ -valued random variable, which is uncorrelated with the preceding two variables, but shares the same mean as $\hat{U}$, meaning $\mu_U = \mu_{\hat{U}}$. The linear control variate framework builds a new $\mathbb{X}$ -valued random variable Z , called the total variate:

$$Z = X - \mathbf{S}(\hat{U} - U), \quad (2.4)$$

where the linear “gain” operator $\mathbf{S}$ is used to minimize the variance in Z by utilizing the information about the distributions of the constituent variates X , $\hat{U}$ and U .

In this work $\mathbb{X}$ and $\mathbb{U}$ are finite dimensional vector spaces. The dimension of $\mathbb{X}$ is n , the dimension of $\mathbb{U}$ is denoted by m when it is the observation space, and by r when it is the reduced order model state space.

The following lemma is a slight generalization of (Rubinstein and Marcus, 1985, Appendix).

Lemma 1 (Optimal gain). *The optimal gain $\mathbf{S}$ that minimizes the trace of the covariance of Z (2.4) is:*

$$\mathbf{S} = \Sigma_{X,\hat{U}}(\Sigma_{\hat{U},\hat{U}} + \Sigma_{U,U})^{-1}. \quad (2.5)$$

Here we focus on the case where the principal and control variates are related through a deterministic projection operator:

$$\hat{U} = \theta(X), \quad (2.6)$$

and where there exists an interpolation operator such that:

$$X \approx \tilde{X} = \phi(\hat{U}), \quad (2.7)$$

which in some optimal sense approximately recovers the information embedded in X . The variable $\tilde{X}$ is known as the reconstruction.

We make the assumption that the projection and interpolation operators are linear:

$$\theta(X) := \Theta X, \quad \phi(\hat{U}) := \Phi \hat{U}. \quad (2.8)$$

We reproduce below the useful result (Popov et al., 2021, Theorem 3.1).

Theorem 1. *Under the assumptions that $\hat{U}$ and U have equal covariances, and that the principal variate residual is uncorrelated with the control variate, the optimal gain of (2.4) is half the interpolation operator:*

$$\Sigma_{\hat{U},\hat{U}} = \Sigma_{U,U} \quad \text{and} \quad \Sigma_{(X-\Phi\hat{U}),\hat{U}} = 0 \quad \Rightarrow \quad \mathbf{S} = \frac{1}{2} \Phi. \quad (2.9)$$

Under the assumptions of Theorem 1 the control variate structure (2.4) is:

$$Z = X - \frac{1}{2} \Phi (\hat{U} - U). \quad (2.10)$$

2.2 Linear Control Variates with Non-Linear Transformations

While working with linear transformations is easy, many applications require reducing the variance of a non-linear transformation of a random variable. We generalize the control variate framework to such an application.

Following Owen (2013), assume that our transformed principal variate is of the form $h(X)$, and the transformed control variate is of the form $g(\hat{U})$, where h and g are smooth non-linear functions. Similarly, assume that the transformed ancillary variate has the form $g(U)$. The total variate Z^h , in the same space as $h(X)$, is defined as:

$$Z^h = h(X) - \mathbf{S} \left(g(\hat{U}) - g(U) \right), \quad (2.11)$$

with the optimal linear gain given by Lemma 1:

$$\mathbf{S} = \Sigma_{h(X), g(\hat{U})} \left(\Sigma_{g(\hat{U}), g(\hat{U})} + \Sigma_{g(U), g(U)} \right)^{-1}. \quad (2.12)$$

Theorem 2. *If $\hat{U}$ and U are Gaussian, and share the same mean and covariance, ($\mu_{\hat{U}} = \mu_U$, $\Sigma_{\hat{U}, \hat{U}} = \Sigma_{U, U}$), then*

$$\mu_{g(\hat{U})} = \mu_{g(U)}, \quad \text{and} \quad \Sigma_{g(\hat{U}), g(\hat{U})} = \Sigma_{g(U), g(U)}, \quad (2.13)$$

the control variate structure (2.11) holds, and the optimal linear gain is:

$$\mathbf{S} = \frac{1}{2} \Sigma_{h(X), g(\hat{U})} \Sigma_{g(\hat{U}), g(\hat{U})}^{-1}. \quad (2.14)$$

Proof. As $\hat{U}$ and U are Gaussian, they are completely defined by their mean and covariance, and any deterministic non-linear transformation of the variates has the same moments, thereby sharing the mean and covariance. In that case $g(\hat{U}) - g(U)$ is unbiased, and the control variate framework can be applied. $\square$

We now consider the case where the principal variate transformation is linear, $h = \mathbf{H}$, and the control and ancillary variate transformation is the interpolation operator, $g = \phi$, such that the underlying control variate is a projection of the principal variate (2.42),

$$\mathbf{H} X \approx \phi(\hat{U}). \quad (2.15)$$

Theorem 3. *Without proof, if the assumption of Theorem 2 hold, and the transformed principal variate residual is uncorrelated with the transformed control variate,*

$$\Sigma_{(\mathbf{H}X - \phi(\hat{U}), \phi(\hat{U}))} \stackrel{\text{assumed}}{=} 0, \quad (2.16)$$

then the optimal gain is

$$\mathbf{S} = \frac{1}{2} \mathbf{I}, \quad (2.17)$$

where $\mathbf{I}$ is the identity operator.

2.3 The Ensemble Kalman Filter

The EnKF is a statistical approximation to the optimal control variate structure (2.4), where the underlying probability density functions are represented by empirical measures using ensembles, i.e., a finite number of realizations of the random variables. The linear control variate framework allows to combine multiple ensembles into one that better represents the desired quantity of interest.

Let $E_{X_i^b} \in \mathbb{R}^{n \times N_X}$ be an ensemble of N_X realizations of the n -dimensional principal variate, which represents our prior uncertainty in the model state at time index i from (2.1). Likewise, let $E_{\mathcal{H}_i(X_i^b)} = \mathcal{H}_i(E_{X_i^b}) \in \mathbb{R}^{m \times N_X}$ be an ensemble of N_X realizations of the m -dimensional control observation state variate, which represents the same model realizations cast into observation space. Let $E_{Y_i} \in \mathbb{R}^{m \times N_X}$ be an ensemble of N_X ‘perturbed observations’, which is a statistical correction required in the ensemble Kalman filter (Burgers et al., 1998).

Remark 4 (EnKF Perturbed Observations). *Each ensemble member of the perturbed observations is sampled from a Gaussian distribution with mean the measured value, and the known observation covariance from (2.2):*

$$[E_{Y_i}]_{:,e} \sim \mathcal{N}(\mu_{Y_i}, \Sigma_{\eta_i, \eta_i}). \quad (2.18)$$

The prior ensemble at time step i is obtained by propagating the posterior ensemble at time $i - 1$ through the model equations,

$$E_{X_i^b} = \mathcal{M}_{t_{i-1}, t_i}(E_{X_{i-1}^a}), \quad (2.19)$$

where the slight abuse of notation indicates an independent model propagation of each ensemble member. Application of the Kalman filter formula constructs an ensemble $E_{X_i^a}$ describing the posterior uncertainty:

$$E_{X_i^a} = E_{X_i^b} - \tilde{\mathbf{K}}_i (E_{\mathcal{H}_i(X_i^b)} - E_{Y_i}), \quad (2.20)$$

where the statistical Kalman gain is an ensemble-based approximation to the exact gain in Lemma 1:

$$\tilde{\mathbf{K}}_i = \tilde{\Sigma}_{X_i^b, \mathcal{H}_i(X_i^b)} \left(\tilde{\Sigma}_{\mathcal{H}_i(X_i^b), \mathcal{H}_i(X_i^b)} + \Sigma_{\eta_i, \eta_i} \right)^{-1}. \quad (2.21)$$

Remark 5 (Inflation). *Inflation is a probabilistic correction necessary to account for the Kalman gain being correlated to the ensemble (Popov and Sandu, 2020). In inflation the ensemble anomalies (deviations from the statistical mean) are multiplied by a constant $\alpha > 1$, thereby increasing the covariance of the distribution described by the ensemble:*

$$E_{X_{i+1}^b} \leftarrow \tilde{\mu}_{X_{i+1}^b} + \alpha (E_{X_{i+1}^b} - \tilde{\mu}_{X_{i+1}^b}). \quad (2.22)$$

2.4 The Multifidelity Ensemble Kalman Filter

The Multifidelity Ensemble Kalman Filter (MFEnKF) (Popov et al., 2021) seeks to merge the information from a hierarchy of models and the corresponding observations into a coherent representation that can be propagated forward in time. This necessitates that the information from the models is decoupled, but implicitly preserving of some underlying structure; we take the linear control variate structure for convenience.

Without loss of generality, we discuss here a bi-fidelity approach, with a telescopic extension to multiple fidelities provided at the end of the section. Instead of having access to one model $\mathcal{M}$, assume that we have access to a hierarchy of models. In the bi-fidelity case, the principal space model is denoted by $\mathcal{M}^X$ and the control space model is denoted by $\mathcal{M}^U$.

We now consider the total variate

$$Z_i^b = X_i^b - \frac{1}{2} \Phi(\hat{U}_i^b - U_i^b), \quad (2.23)$$

that describes the prior total information from a model that evolves in principal space (X_i^b) and a model that evolves in ancillary space ($\hat{U}_i^b$ and U_i^b).

Assume that our prior total variate is represented by the three ensembles $E_{X_i^b} \in \mathbb{R}^{n \times N_X}$ consisting of N_X realizations of the n -dimensional principal model state variate, $E_{\hat{U}_i^b} \in \mathbb{R}^{r \times N_X}$ consisting of N_X realizations of the r -dimensional control model state variate, and $E_{U_i^b} \in \mathbb{R}^{r \times N_U}$ consisting of N_U realizations of the r -dimensional ancillary model state variate. Each of these ensembles has a corresponding ensemble of m -dimensional control observation space realizations.

MFEnKF performs sequential data assimilation using the above constituent ensembles, without having to explicitly calculate the ensemble of the total variates. The MFEnKF forecast step propagates the three ensembles from the previous step:

$$\begin{aligned} E_{X_i^b} &= \mathcal{M}_{t_{i-1}, t_i}^X(E_{X_{i-1}^a}), \\ E_{\hat{U}_i^b} &= \mathcal{M}_{t_{i-1}, t_i}^U(E_{\hat{U}_{i-1}^a}), \\ E_{U_i^b} &= \mathcal{M}_{t_{i-1}, t_i}^U(E_{U_{i-1}^a}). \end{aligned} \quad (2.24)$$

Two observation operators $\mathcal{H}_i^X$ and $\mathcal{H}_i^U$ cast the principal model and control model spaces, respectively, into the control observation space. In this paper we assume that the principal model space observation operator is the canonical observation operator (2.2):

$$\mathcal{H}_i^X(X_i) := \mathcal{H}_i(X_i), \quad (2.25)$$

and that the control model space observation operator is the canonical observation operator (2.2) applied to the linear interpolated reconstruction (2.7) of a variable in control model space:

$$\mathcal{H}_i^U(U_i) := \mathcal{H}_i(\Phi U_i). \quad (2.26)$$

Additionally, we define an (approximate) observation operator for the total model variate :

$$\mathcal{H}_i^Z(Z_i) := \mathcal{H}_i^X(X_i) - \frac{1}{2} \left(\mathcal{H}_i^U(\hat{U}_i) - \mathcal{H}_i^U(U_i) \right), \quad (2.27)$$

which, under the linearity assumptions on $\mathcal{H}_i^X$ of Theorem 3 and the underlying Gaussian assumptions on $\hat{U}_i$ and U_i of Theorem 2, begets that $\mathcal{H}_i^Z = \mathcal{H}_i^X$. Even without the linearity assumption the definition (2.27) is operationally useful.

The MFEnKF analysis updates each constituent ensemble as follows:

$$\begin{aligned} E_{X_i^a} &= E_{X_i^b} - \tilde{\mathbf{K}}_i \left(E_{\mathcal{H}_i^X(X_i^b)} - E_{Y_i^X} \right), \\ E_{\hat{U}_i^a} &= E_{\hat{U}_i^b} - \Theta \tilde{\mathbf{K}}_i \left((E_{\mathcal{H}_i^U(\hat{U}_i^b)} - E_{Y_i^X}) \right), \\ E_{U_i^a} &= E_{U_i^b} - \Theta \tilde{\mathbf{K}}_i \left((E_{\mathcal{H}_i^U(U_i^b)} - E_{Y_i^U}) \right), \end{aligned} \quad (2.28)$$

with the heuristic correction to the means

$$\tilde{\mu}_{X_i^a} \leftarrow \tilde{\mu}_{Z_i^a}, \quad \tilde{\mu}_{\hat{U}_i^a} \leftarrow \Theta \tilde{\mu}_{Z_i^a}, \quad \tilde{\mu}_{U_i^a} \leftarrow \Theta \tilde{\mu}_{Z_i^a}, \quad (2.29)$$

applied in order to fulfill the unbiasedness requirement of the control variate structure:

$$\tilde{\mu}_{Z_i^a} = \tilde{\mu}_{Z_i^b} - \tilde{\mathbf{K}}_i \left(\tilde{\mu}_{\mathcal{H}_i^X(Z_i^b)} - \mu_{Y_i} \right). \quad (2.30)$$

The Kalman gain and the covariances are defined by the semi-linearization:

$$\tilde{\mathbf{K}}_i = \tilde{\Sigma}_{Z_i^b, \mathcal{H}_i^Z(Z_i^b)} \left(\tilde{\Sigma}_{\mathcal{H}_i^Z(Z_i^b), \mathcal{H}_i^Z(Z_i^b)} + \Sigma_{\eta_i, \eta_i} \right)^{-1} \quad (2.31)$$

$$\begin{aligned} \tilde{\Sigma}_{Z_i^b, \mathcal{H}_i^Z(Z_i^b)} &= \tilde{\Sigma}_{X_i^b, \mathcal{H}_i^X(X_i^b)} + \frac{1}{4} \tilde{\Sigma}_{\Phi \hat{U}_i^b, \mathcal{H}_i^U(\hat{U}_i^b)} + \frac{1}{4} \tilde{\Sigma}_{\Phi U_i^b, \mathcal{H}_i^U(U_i^b)} \\ &\quad - \frac{1}{2} \tilde{\Sigma}_{X_i^b, \mathcal{H}_i^U(\hat{U}_i^b)} - \frac{1}{2} \tilde{\Sigma}_{\Phi \hat{U}_i^b, \mathcal{H}_i^X(X_i^b)}, \end{aligned} \quad (2.32)$$

$$\begin{aligned} \tilde{\Sigma}_{\mathcal{H}_i^Z(Z_i^b), \mathcal{H}_i^Z(Z_i^b)} &= \tilde{\Sigma}_{\mathcal{H}_i^X(X_i^b), \mathcal{H}_i^X(X_i^b)} + \frac{1}{4} \tilde{\Sigma}_{\mathcal{H}_i^U(\hat{U}_i^b), \mathcal{H}_i^U(\hat{U}_i^b)} + \frac{1}{4} \tilde{\Sigma}_{\mathcal{H}_i^U(U_i^b), \mathcal{H}_i^U(U_i^b)} \\ &\quad - \frac{1}{2} \tilde{\Sigma}_{\mathcal{H}_i^X(X_i^b), \mathcal{H}_i^U(\hat{U}_i^b)} - \frac{1}{2} \tilde{\Sigma}_{\mathcal{H}_i^U(\hat{U}_i^b), \mathcal{H}_i^X(X_i^b)}. \end{aligned} \quad (2.33)$$

In order to ensure that the control variate $\hat{U}$ remains highly correlated to the principal variate X , at the end of each analysis step we replace the analysis control variate ensemble with the corresponding projection of the principal variate ensemble:

$$\mathbf{E}_{\hat{U}_i^a} \leftarrow \Theta \mathbf{E}_{X_i^a}. \quad (2.34)$$

Remark 6 (MFEnKF Perturbed observations). *There is no unique way to perform perturbed observations (remark 4) in the MFEnKF. We will present one way in this paper. As Theorem 2 requires both the control and ancillary variates to share the same covariance, we utilize here the ‘control space uncertainty consistency’ approach. The perturbed observations ensembles in (2.28) is defined by:*

$$[\mathbf{E}_{Y_i^X}]_{:,e} \sim \mathcal{N}(\mu_{Y_i}, \Sigma_{\eta_i, \eta_i}), \quad (2.35)$$

$$[\mathbf{E}_{Y_i^U}]_{:,e} \sim \mathcal{N}(\mu_{Y_i}, s \Sigma_{\eta_i, \eta_i}), \quad (2.36)$$

where the scaling factor is $s = 1$. See (Popov et al., 2021, Section 4.2) for a more detailed discussion about perturbed observations.

Remark 7 (MFEnKF Inflation). *Similarly to the EnKF (see Remark 5), the MFEnKF also requires inflation in order to account for the statistical Kalman gain being correlated to its constituent ensembles. For a bi-fidelity MFEnKF, two inflation factors are required: α_X which acts on the anomalies of the principal and control variates (as they must remain highly correlated) and α_U which acts on the ensemble anomalies of the ancillary variate:*

$$\begin{aligned} \mathbf{E}_{X_{i+1}^b} &\leftarrow \tilde{\mu}_{X_{i+1}^b} + \alpha_X \left(\mathbf{E}_{X_{i+1}^b} - \tilde{\mu}_{X_{i+1}^b} \right), \\ \mathbf{E}_{\hat{U}_{i+1}^b} &\leftarrow \tilde{\mu}_{\hat{U}_{i+1}^b} + \alpha_X \left(\mathbf{E}_{\hat{U}_{i+1}^b} - \tilde{\mu}_{\hat{U}_{i+1}^b} \right), \\ \mathbf{E}_{U_{i+1}^b} &\leftarrow \tilde{\mu}_{U_{i+1}^b} + \alpha_U \left(\mathbf{E}_{U_{i+1}^b} - \tilde{\mu}_{U_{i+1}^b} \right). \end{aligned} \quad (2.37)$$

2.5 Projection-Based Reduced Order Models

It is well known that the important information of many dynamical systems can be expressed with significantly fewer dimensions than the discretization dimension n (Sell and You, 2013). For many infinite dimensional equations it is possible to construct a finite-dimensional inertial manifold that represents the dynamics of the system (including the global attractor). The Hausdorff dimension of the global attractor of some dynamical system is a useful lower bound for the minimal representation of the dynamics, though a representation of just the attractor is likely not sufficient to fully capture all the ‘useful’ aspects of the data. For data-based reduced order models an important aspect is the intrinsic dimension (Lee and Verleysen, 2007) of the data. The authors’ are not aware of any formal statements relating the dimension of an inertial manifold and the intrinsic dimension of some finite discretization of the dynamics. For some reduced dimension r that is sufficient to represent either the dynamics or the data, or both, it is possible to build a ‘useful’ surrogate model.

We will now discuss the construction of reduced order models for problems posed as ordinary differential equations. The following derivations are similar to those found in Lee and Carlberg (2020), but assume vector spaces and no re-centering.

Just like in the control variate framework in section 2.1, our full order model will reside in the principal space $\mathbb{X}$ and our reduced order model will be defined in the space $\mathbb{U}$, which is related to $\mathbb{X}$ through some coupling.

Given an initial value problem in $\mathbb{X}$:

$$\frac{dX}{dt} = f(X), \quad X(t_0) = X_0, \quad t \in [t_0, t_f], \quad (2.38)$$

and the projection operator (2.42), the induced reduced order model initial value problem in $\mathbb{U}$ is defined by simple differentiation of $U = \theta(X)$, by dynamics in the space $\mathbb{U}$,

$$\frac{dU}{dt} = \theta'(X)f(X), \quad X(t_0) = X_0, \quad t \in [t_0, t_f]. \quad (2.39)$$

As is common, the full order trajectory is not available during integration, as there is no bijection from $\mathbb{X}$ to $\mathbb{U}$, thus an approximation using the interpolation operator (2.7) that fully resides in $\mathbb{U}$ is used instead:

$$\frac{dU}{dt} = \theta'(\phi(U))f(\phi(U)), \quad U(t_0) = \phi(X_0), \quad t \in [t_0, t_f]. \quad (2.40)$$

Note that this is not the only way to obtain a reduced order model by using arbitrary projection and interpolation operators. It is however the simplest extension of the POD-based ROM framework.

Remark 8 (Linear ROM). *In the linear case (2.8) the reduced order model (2.40) takes the form*

$$\frac{dU}{dt} = \Theta f(\Phi U), \quad U(t_0) = \phi(X_0), \quad t \in [t_0, t_f]. \quad (2.41)$$

2.6 Autoencoders

An autoencoder (Aggarwal et al., 2018) is an artificial neural network consisting of two components, an encoder θ and a decoder ϕ , such that given a variable X in the principal space, the variable

$$\hat{U} = \theta(X), \quad (2.42)$$

resides in the control space of the encoder. Conversely the reconstruction,

$$\tilde{X} = \phi(\hat{U}), \quad (2.43)$$

is an approximation to X in the principal space. In our context we think of the encoder as a non-linear projection operator (2.42), and of the decoder as a nonlinear interpolation operator (2.7). While the relative dimension n of the principal space is relatively high, the arbitrary structure of an artificial neural networks allows the autoencoder to learn the optimal r -dimensional (small) representation of the data.

3. NON-LINEAR PROJECTION-BASED MFENKF

We extend MFEnKF to work with non-linear projection and interpolation operators. Theoretical extensions of the linear control variate framework to the non-linear case (Nelson, 1987) are not as-of-yet satisfactory. As is traditional in data assimilation literature, we resort to heuristic assumptions, which in practice may work significantly better than the theory predicts. The new algorithm is named NL-MFEnKF.

The main idea is to replace the optimal control variate structure for linear projection and interpolation operators (2.10) with one that works with their non-linear counterparts (2.11):

$$Z_i^b = X_i^b - \frac{1}{2} \left(\phi(\hat{U}_i^b) - \phi(U_i^b) \right). \quad (3.1)$$

We assume that $\hat{U}$ and U are Gaussian, and have the same mean and covariance, such that they obey the assumptions made in Theorem 2 and in Theorem 3 for the optimal gain.

Similar to MFEnKF (2.26), the control model space observation operator is the application of the canonical observation operator (2.2) to the reconstruction

$$\mathcal{H}_i^U(U_i) := \mathcal{H}_i(\phi(U_i)), \quad (3.2)$$

with the other observation operators (2.25, 2.27) defined as in the MFEnKF.

Remark 9. *It is of independent interest to explore control model space observation operators that are not of the form (3.2). For example, if the interpolation operator ϕ is created through an autoencoder, the control model space observation operator $\mathcal{H}^U$ could similarly be a different decoder of the same latent space.*

The MFEnKF equations (2.28) are replaced by their non-linear counterparts in a manner similar to what is done with non-linear observation operators,

$$\begin{aligned} E_{X_i^a} &= E_{X_i^b} - \tilde{\mathbf{K}}_i \left(E_{\mathcal{H}_i^X(X_i^b)} - E_{Y_i^X} \right), \\ E_{\hat{U}_i^a} &= E_{\hat{U}_i^b} - \tilde{\mathbf{K}}_i^\theta \left(E_{\mathcal{H}_i^U(\hat{U}_i^b)} - E_{Y_i^X} \right), \\ E_{U_i^a} &= E_{U_i^b} - \tilde{\mathbf{K}}_i^\theta \left(E_{\mathcal{H}_i^U(U_i^b)} - E_{Y_i^U} \right), \end{aligned} \quad (3.3)$$

where, as opposed to (2.28), there are now two Kalman gains, defined by,

$$\tilde{\mathbf{K}}_i = \tilde{\Sigma}_{Z_i^b, \mathcal{H}_i^Z(Z_i^b)} \left(\tilde{\Sigma}_{\mathcal{H}_i^Z(Z_i^b), \mathcal{H}_i^Z(Z_i^b)} + \Sigma_{\eta_i, \eta_i} \right)^{-1}, \quad (3.4)$$

$$\tilde{\mathbf{K}}_i^\theta = \tilde{\Sigma}_{\theta(Z_i^b), \mathcal{H}_i^Z(Z_i^b)} \left(\tilde{\Sigma}_{\mathcal{H}_i^Z(Z_i^b), \mathcal{H}_i^Z(Z_i^b)} + \Sigma_{\eta_i, \eta_i} \right)^{-1}, \quad (3.5)$$

for our purposes the covariances are semi-linear approximations, similar to (2.32) and (2.33), and the perturbed observations are defined just like in the MFEnKF (Remark 6).

3.1 NL-MFEnKF Heuristic Corrections

For linear operators the projection of the mean is the mean of the projection. This is however not true for general non-linear operators. Thus in order to correct the means like in the MFEnKF (2.29), additional assumptions have to be made.

The empirical mean of the total analysis variate (similar to (3.1) and (2.30)) is

$$\tilde{\mu}_{Z_i^a} = \tilde{\mu}_{X_i^a} - \frac{1}{2} \left(\tilde{\mu}_{\Phi(\hat{U}_i^a)} - \tilde{\mu}_{\Phi(U_i^a)} \right). \quad (3.6)$$

We use it to find the optimal mean adjustments in reduced space. Specifically, we set the mean of the principal variate to be the mean of the total variate (3.6),

$$\tilde{\mu}_{X_i^a} \leftarrow \tilde{\mu}_{Z_i^a}, \quad (3.7)$$

enforce the recorrelation of the principal and control variates (2.34) via

$$E_{\hat{U}_i^a} \leftarrow \theta(E_{X_i^a}), \quad (3.8)$$

and define the control variate mean adjustment as,

$$\tilde{\mu}_{\hat{U}_i^a} \leftarrow \tilde{\mu}_{\theta(X_i^a)}. \quad (3.9)$$

Unlike the linear control variate framework of the MFEnKF (2.23), the non-linear framework of the NL-MFEnKF (3.1) does not induce a unique way to impose unbiasedness on the control-space variates. There are many possible non-linear formulations to the MFEnKF, and many possible heuristic corrections of the mean the ancillary variate. Here we discuss three approaches based on:

1. control space unbiased mean adjustment,
2. principal space unbiased mean adjustment, and
3. Kalman approximate mean adjustment,

each stemming from a different assumption on the relationship between the ancillary variate and the other variates.

3.1.1 Control Space Unbiased Mean Adjustment

By assuming that the control variate $\hat{U}_i^a$ and the ancillary variate U_i^a are unbiased in the control space, the mean adjustment of $\hat{U}_i^a$ in (3.9) directly defines the mean adjustment of the ancillary variate:

$$\tilde{\mu}_{U_i^a} \leftarrow \tilde{\mu}_{\hat{U}_i^a}. \quad (3.10)$$

The authors will choose this method of correction in the numerical experiments for both its properties and ease of implementation.

3.1.2 Principal Space Unbiased Mean Adjustment

If instead we assume that the control variate $\Phi(\hat{U}_i)$ and the ancillary variate $\Phi(U_i)$ are unbiased in the principal space, then the mean of the total variate Z_i (3.1) equals the mean of the principal variate X_i , a desirable property.

Finding a mean adjustment for U_i^a in the control space such that the unbiasedness is satisfied in the principal space is a non-trivial problem. Explicitly, we seek a vector $\mathbf{v}$ such that:

$$\tilde{\boldsymbol{\mu}}_{\phi(\tilde{U}_i^a)} = \tilde{\boldsymbol{\mu}}_{\phi(U_i^a - \tilde{\boldsymbol{\mu}}_{U_i^a} + \mathbf{v})}. \quad (3.11)$$

The solution to (3.11) requires the solution to an expensive inverse problem. Note that (3.11) is equivalent to (3.10) under the assumptions of Theorem 2, in the limit of ensemble size.

3.1.3 Kalman Approximate Mean Adjustment

Instead of assuming that the control and ancillary variates are unbiased, we can consider directly the mean of the control-space total variate:

$$\tilde{\boldsymbol{\mu}}_{\theta(Z_i^a)} = \tilde{\boldsymbol{\mu}}_{\theta(X_i^a)} - \frac{1}{2} \left(\tilde{\boldsymbol{\mu}}_{\tilde{U}_i^a} - \tilde{\boldsymbol{\mu}}_{U_i^a} \right), \quad (3.12)$$

defined with the mean values in NL-MFEnKF formulas (3.3). The following adjustment to the mean of the ancillary variate:

$$\tilde{\boldsymbol{\mu}}_{U_i^a} \leftarrow \tilde{\boldsymbol{\mu}}_{\theta(Z_i^a)}, \quad (3.13)$$

is not unbiased with respect to the control variate in any space, but provides a heuristic approximation of the total variate mean in control space.

3.2 Telescopic Extension

As in (Popov et al., 2021, Section 4.5) one can telescopically extend the bi-fidelity NL-MFEnKF algorithm to $\mathcal{L}+1$ fidelities. Assume that the nonlinear operator ϕ_ℓ interpolates from the space of fidelity ℓ to the space of fidelity $\ell-1$, where ϕ_1 interpolates to the principal space. A telescopic extension of (3.1) is

$$Z = X - \sum_{\ell=1}^{\mathcal{L}} 2^{-\ell} \left(\bar{\Phi}(\hat{U}_\ell) - \bar{\Phi}(U_\ell) \right), \quad (3.14)$$

where the projection operator at each fidelity is defined as,

$$\bar{\Phi}_\ell = \phi_1 \circ \dots \circ \phi_\ell, \quad (3.15)$$

projecting from the space of fidelity ℓ to the principal space. The telescopic extension of the NL-MFEnKF is not analyzed further in this work.

4. MODELS FOR NUMERICAL EXPERIMENTS

For numerical experiments we use the Lorenz '96 model, and two surrogate models that approximate it's dynamics:

1. a principal orthogonal decomposition-based quadratic reduced order model (POD-ROM), and
2. an autoencoder neural network-based reduced order model (NN-ROM).

For each case we construct reduced order models (ROMs) for reduced dimension sizes of $r = 7, 14, 21, 28$, and 35 . The Lorenz '96 and the POD-ROM models are implemented in the ODE-test-problems suite (Computational Science Laboratory, 2020; Roberts et al., 2019).

4.1 Lorenz '96

The Lorenz '96 model (Lorenz, 1996) is conjured from the PDE (Reich and Cotter, 2015),

$$\frac{dy}{dt} = -yy_x - y + F, \quad (4.1)$$

where the forcing parameter is set to $F = 8$. In the semi-discrete version $\mathbf{y} \in \mathbb{R}^n$, and the nonlinear term is approximated by a numerically unstable finite difference approximation,

$$[\mathbf{y}\mathbf{y}_x]_k = (\hat{\mathbf{I}}\mathbf{y}) \cdot (\hat{\mathbf{D}}\mathbf{y}) = ([\mathbf{y}]_{k-1}) \cdot ([\mathbf{y}]_{k-2} - [\mathbf{y}]_{k+1}), \quad (4.2)$$

where $\hat{\mathbf{I}}$ is a (linear) shift operator, and the linear operator $\hat{\mathbf{D}}$ is a first order approximation to the first spatial derivative. The canonical $n = 40$ variable discretization with cyclic boundary conditions is used. The canonical fourth order Runge-Kutta method is used to discretize the time dimension.

For the given discrete formulation of the Lorenz '96 system, 14 represents the number of non-negative Lyapunov exponents, 28 represents the rounded-up Kaplan-Yorke dimension of 27.1, and 35 represents an approximation of the intrinsic dimension of the system (calculated by the method provided by Bahadur and Paffenroth (2019)). To the authors' knowledge, the inertial manifold of the system, if it exists, is not known. The relatively high ratio between the intrinsic dimension of the system and the spatial dimension of the system makes constructing a reduced order model particularly challenging.

4.1.1 Data for Constructing Reduced-Order Models.

The data to construct the reduced order models is taken to be $T = 5000$ state snapshots from a representative model run. The snapshots are spaced 36 time units apart, equivalent to six months in model time.

4.2 Proper Orthogonal Decomposition ROM

Using the method of snapshots (Sirovich, 1987), we construct optimal linear operators, $\Phi^T = \Theta \in \mathbb{R}^{r \times n}$, such that the projection captures the dominant orthogonal modes of the system dynamics. The reduced order model approximation with linear projection and interpolation operators (2.41) is quadratic (similar to Mou et al. (2021); San and Iliescu (2015))

$$\frac{d\mathbf{u}}{dt} = \mathbf{a} + \mathbf{B}\mathbf{u} + \mathbf{u}^T \mathcal{C} \mathbf{u}, \quad (4.3)$$

where the corresponding vector $\mathbf{a}$, matrix $\mathbf{B}$, and 3-tensor $\mathcal{C}$ are defined by:

$$\mathbf{a} = F\Theta 1_n, \quad (4.4a)$$

$$\mathbf{B} = -\Theta\Phi, \quad (4.4b)$$

$$[\mathcal{C}]_{pqr} = -(\hat{\mathbf{I}}\Phi_{:,p})^T (\hat{\mathbf{D}}\Phi_{:,q}) \Phi_{:,r}. \quad (4.4c)$$

4.3 Autoencoder based ROM

We now discuss building the neural network-based reduced order model (NN-ROM). Given the principal space variable X , consider an encoder that computes its projection $\hat{U}$ onto the control space (2.42), and a decoder that computes the reconstruction $\tilde{X}$ (2.43).

Canonical autoencoders simply aim to minimize the reconstruction error:

$$X \approx \tilde{X}, \quad (4.5)$$

which attempts to capture the dominant modes of the intrinsic manifold of the data. However, for our purposes, we seek to preserve two additional properties.

The first property that is especially important to our application is the stable preservation of the latent space through multiple projections and interpolations:

$$\theta(\phi(\hat{U})) \approx \hat{U}, \quad (4.6)$$

which we will call the *left-inverse property*. For POD, this property is automatically preserved by construction and the linearity of the methods. For non-linear operators, the authors have not explicitly seen this property preserved, however, as the MFEnKF requires sequential applications of projections and interpolations, we believe that for the use-case outlined in this paper, the property is especially important.

The autoencoder structure automatically induces a reduced order model $\mathcal{M}^{AE}$, corresponding to the equation (2.40), that captures the dynamics of the system in the r -dimensional control space. The second property is that the r -dimensional ROM dynamics follows accurately the n -dimensional dynamics of the full order model.

Remark 10. *Note that unlike the POD-ROM whose linear structure induces a purely r -dimensional initial value problem, the NN-ROM (2.40) still involves n -dimensional function evaluations. In a practical method it would be necessary to reduce the internal dimension of the ROM, however that is significantly outside the scope of this paper.*

We wish to construct a proof-of-concept neural network surrogate. We seek to ensure that the induced dynamics (2.40) makes accurate predictions, by not only capturing the intrinsic manifold of the data, but also attempting to capture the inertial manifold of the system. Explicitly, we wish to ensure that

$$\|\mathcal{M}(X) - \phi(\mathcal{M}^{AE}(\theta(X)))\| \quad (4.7)$$

is minimized. In this sense (2.40) would represent an approximation of the dynamics along a submanifold of the inertial manifold. In practice we replace (4.7) by the difference between a short trajectory in the full space started from a certain initial value, and the corresponding short trajectory in the latent space started from the projected initial value.

Combining the canonical autoencoder reconstruction error term (4.5), the latent space reconstruction error term (preserving property (4.6)), and the temporal reconstruction error (4.7) (computed along K small steps), we arrive at the following loss function for each snapshot:

$$\begin{aligned} \ell_j(X_j) = & \frac{1}{n} \|X_j - \phi(\theta(X_j))\|_2^2 + \frac{\lambda_1}{r} \|\theta(X_j) - \theta(\phi(\theta(X_j)))\|_2^2 \\ & + \sum_{k=1}^K \frac{\lambda_2}{n} \left\| \mathcal{M}_{t_j, (t_j+k\Delta t)}(X_j) - \phi(\mathcal{M}_{t_j, (t_j+k\Delta t)}^{AE}(\theta(X_j))) \right\|_2^2, \end{aligned} \quad (4.8)$$

where the hyper-parameters λ_1 and λ_2 represent the inverse relative variance of the mismatch.

The full loss function, combining the cost functions for all T snapshots,

$$L(X) = \sum_{j=1}^T \ell_j(X_j), \quad (4.9)$$

can be minimized through typical stochastic optimization methods.

4.3.1 Lorenz '96 Surrogate

Similar to the POD model, we construct r -dimensional NN-based surrogate ROMs. To this end, we use one hidden layer networks with the tanh activation function for the encoder (2.42) and decoder (2.43):

$$\theta(X) = \mathbf{W}_2^\theta \tanh(\mathbf{W}_1^\theta X + \mathbf{b}_1^\theta) + \mathbf{b}_2^\theta, \quad (4.10)$$

$$\mathbf{W}_1^\theta \in \mathbb{R}^{h \times n}, \mathbf{W}_2^\theta \in \mathbb{R}^{r \times h}, \mathbf{b}_1^\theta \in \mathbb{R}^h, \mathbf{b}_2^\theta \in \mathbb{R}^r,$$

$$\phi(U) = \mathbf{W}_2^\phi \tanh(\mathbf{W}_1^\phi U + \mathbf{b}_1^\phi) + \mathbf{b}_2^\phi, \quad (4.11)$$

$$\mathbf{W}_1^\phi \in \mathbb{R}^{h \times r}, \mathbf{W}_2^\phi \in \mathbb{R}^{n \times h}, \mathbf{b}_1^\phi \in \mathbb{R}^h, \mathbf{b}_2^\phi \in \mathbb{R}^n,$$

where the hidden layer dimension on both the encoder and decoder is set to $h = 200$ in order to approximate well most non-linear projection and interpolation functions.

The hyperparameter values used in the numerical experiments are $\lambda_1 = 10^3$, $\lambda_2 = 1$, and $K = 5$. We employ the ADAM (Kingma and Ba, 2014) optimization procedure to train the NN and produce the various ROMs needed. Gradients of the loss function are computed through automatic differentiation.

5. NUMERICAL EXPERIMENTS

The numerical experiments compare the following four methodologies :

1. Standard EnKF with the Lorenz '96 full order model;
2. MFEnKF with the POD surrogate model, an approach named MFEnKF(POD);
3. NL-MFEnKF with the autoencoder surrogate model, named NL-MFEnKF(NN); and
4. MFEnKF with the autoencoder surrogate model, named MFEnKF(NN).

Since MFEnKF does not support non-linear projections and interpolations, in MFEnKF(NN) the ensembles are interpolated into the principal space, and assimilated under the assumption that $\Theta = \Phi = \mathbf{I}$.

To measure the accuracy of the analysis mean with respect to the truth we use the spatio-temporal root mean square error (RMSE):

$$\text{RMSE} = \sqrt{\sum_{i=1}^N \sum_{k=1}^n \left(\left[\tilde{\mu}_{X_i^s} \right]_k - [X_i^t]_k \right)^2}, \quad (5.1)$$

where N is the number of data assimilation steps for which we wish to measure the error.

For sequential data assimilation experiments we observe all 40 variables of the Lorenz '96 system, with an observation error covariance of $\Sigma_{\eta_i, \eta_i} = \mathbf{I}$. Observations are performed every 0.05 time units corresponding to 6 hours in model time. We run 20 independent realizations (independent ensemble initializations) for 1100 time steps, but discard the first 100 steps for spinup.

5.1 Accuracy of ROM Models

Our first experiment is concerned with the preservation of energy by our ROMs, and seeks to compare the accuracy of NN-ROM against that of POD-ROM. For the Lorenz '96 model, we

r	POD-ROM	NN-ROM
7	0.52552	0.56753
14	0.70200	0.76303
21	0.82222	0.88997
28	0.90161	0.95494
35	0.96251	0.98483

TABLE 1: Relative kinetic energies preserved by the reconstructions of the POD-ROM and the NN-ROM solutions of the Lorenz '96 system. Various reduced-order model dimensions r are considered.

use the following equation (Karimi and Paul, 2010) to model the spatio-temporal kinetic energy,

$$\text{KE} = \sum_{i=1}^T \sum_{k=1}^n ([y_i]_k)^2, \quad (5.2)$$

where T is the number of temporal snapshots that were used to construct the ROMs. Table 1 shows the relative kinetic energies of the POD-ROM and the NN-ROM reconstructed solutions (2.7) (the energies of the reconstructed ROM solutions are by the kinetic energy of the full order solution).

The results lead to several observations. First, the NN-ROM always preserves more energy than the POD-ROM. We have achieved our goal to build an NN-ROM that is more accurate than the POD-ROM. Second, the NN-ROMs with dimensions $r = 21$ and $r = 28$ preserve almost as much energy as the POD-ROMs with dimension $r = 28$ and $r = 35$, respectively. Intuitively this tells us that they should be just as accurate.

5.2 Impact of ROM Dimension

The second set of experiments seeks to learn how the ROM dimension affects the accuracy of the various multifidelity data assimilation algorithms.

We take the principal ensemble size to be $N_X = 32$, and the surrogate ensemble sizes equal to $N_U = r - 3$, in order to always work in the undersampled regime. All multifidelity algorithms are run with inflation factors $\alpha_X = 1.05$ and $\alpha_U = 1.01$. The traditional EnKF using the full order model is run with an ensemble size of $N = N_X$ and an inflation factor $\alpha = 1.06$ to ensure stability.

The results are shown in Figure 1. For the ‘interesting’ dimensions $r = 28$, $r = 35$, and an underrepresented $r = 21$, the NL-MFEnKF(NN) performs significantly better than the MFEnKF(POD). For a severely underrepresented ROM dimension of $r = 7$ and $r = 14$ the MFEnKF(POD) outperforms or ties the NL-MFEnKF(NN). The authors suspect that this is either due to the NN-ROMs sharing the same hyperparameters, meaning that if the hyperparameters were tailored to each r dimensional NN-ROM individually, the authors suspect that it would beat the POD-ROM in all cases.

Of note is that excluding the case of $r = 35$, the MFEnKF(NN), using the standard MFEnKF method in the principal space, is the least accurate among all algorithms, indicating that the non-linear method presented in this paper is truly needed for models involving non-linear model reduction.




FIG. 1: Analysis RMSE versus ROM dimension for three multifidelity data assimilation algorithms and the classical EnKF. Ensemble sizes are $N_X = 32$ and $N_U = r - 3$. Error bars show two standard deviations. The inflation factor for the surrogate ROMs is fixed at $\alpha_U = 1.01$, the inflation of $\alpha = 1.07$ is used for the EnKF, and $\alpha_X = 1.05$ is used for all other algorithms.

We note that for $r = 35$, the suspected intrinsic dimension of the data, the NL-MFEnKF(NN) outperforms the EnKF, both in terms of RMSE and variability within runs. This is additionally strengthened by the results of the MFEnKF(NN) assimilated in the principal space, as it implies that there is little-to-no loss of information in the projection into the control space.

5.3 Ensemble Size and Inflation

Our last set of experiments focuses on the particular ROM dimension $r = 28$, as we believe that it is representative of an operationally viable dimension reduction, covering the dimensionality of the global attractor.

For each of the four algorithms we vary the principal ensemble size $N_X = N$, and the principal inflation factor $\alpha_X = \alpha$. As before, we set the control ensemble size to $N_U = r - 3 = 25$ and the control-space inflation factor to $\alpha_U = 1.01$.

Figure 2 shows the spatio-temporal RMSE for various choices of ensemble sizes and inflation factors. The results show compelling evidence that NL-MFEnKF(NN) is competitive when compared to MFEnKF(POD); the two methods have similar stability properties for a wide range of principal ensemble size N_X and principal inflation α_X , but NL-MFEnKF(NN) yields smaller analysis errors for almost all scenarios for which the two methods are stable.

For a few points with low values of principal inflation α_X , the NL-MFEnK(NN) is not as stable as the MFEnKF(POD). This could be due to either an instability in the NN-ROM itself, in the NL-MFEnKF itself, or in the projection and interpolation operators.

An interesting observation is that the MFEnKF(NN), which is assimilated naively in the principal space, becomes less stable for larger ensemble sizes N_X . One possible explanation for this is that the ensemble mean estimates become more accurate, thus the bias between the ancillary and control variates is amplified in (2.4), and more error is introduced from the surrogate model. This is in contrast to most other ensemble based methods, including all others in this paper, whose error is lowered by increasing ensemble size.




FIG. 2: Analysis RMSE for four data assimilation algorithms and various values of the principal ensemble size N_X and inflation factor α_X . The surrogate ROM size is fixed to $r = 28$ with ensemble size of $N_U = 25$ and inflation of $\alpha_U = 1.01$. The top left panel represents the MFEEnKF with a POD-ROM surrogate, the top right panel represents the NL-MFEEnKF with a NN-ROM, the bottom left panel represents the MFEEnKF with a NN-ROM surrogate assimilated in the principal space, and the bottom right panel represents the classical EnKF.

6. CONCLUSIONS

This work extends the multifidelity ensemble Kalman filter framework to work with nonlinear surrogate models. The control variate framework is generalized to incorporate optimal non-linear projection and interpolation operators implemented using autoencoders.

The results obtained in this paper indicate that reduced order models based on non-linear projections that fully capture the intrinsic dimension of the data provide excellent surrogates for use in multifidelity sequential data assimilation. Moreover, nonlinear generalizations of the control variate framework result in small approximation errors, and thus the assimilation can be carried out efficiently in the space of a nonlinear reduced model.

Future directions to be pursued include more robust constructions of projection and interpolation operators, the use of recurrent neural network models, and applications to operational test problems.

ACKNOWLEDGMENTS

This work was supported by awards NSF ACI-1709727, NSF CDS&E-MSS-1953113, DOE ASCR DE-SC0021313, and by the Computational Science Laboratory at Virginia Tech.

REFERENCES

- Aggarwal, C.C. , *Neural Networks and Deep Learning*, Springer, 2018.
- Asch, M., Bocquet, M., and Nodet, M., *Data Assimilation: Methods, Algorithms, and Applications*, SIAM, 2016.
- Bahadur, N. and Paffenroth, R., Dimension Estimation using Autoencoders, *arXiv preprint arXiv:1909.10702*, 2019.
- Blanchard, E., Sandu, A., and Sandu, C., Parameter Estimation for Mechanical Systems via an Explicit Representation of Uncertainty, *Engineering Computations. International Journal for Computer-Aided Engineering and Software*, vol. **26**, no. 5, pp. 541–569, 2009.
URL <http://dx.doi.org/10.1108/02644400910970185>
- Blanchard, E., Sandu, A., and Sandu, C., A Polynomial Chaos-Based Kalman Filter Approach for Parameter Estimation of Mechanical Systems, *Journal of Dynamic Systems, Measurement, and Control*, vol. **132**, no. 6, p. 18, 2010a.
URL <http://dx.doi.org/10.1115/1.4002481>
- Blanchard, E., Sandu, A., and Sandu, C., Polynomial Chaos based Parameter Estimation Methods for Vehicle Systems, *Journal of Multi-body Dynamics*, vol. **224**, pp. 59–81, 2010b.
URL <http://dx.doi.org/10.1243/14644193JMBD204>
- Brunton, S.L. and Kutz, J.N., *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, Cambridge University Press, 2019.
- Burgers, G., van Leeuwen, P.J., and Evensen, G., Analysis Scheme in the Ensemble Kalman Filter, *Monthly Weather Review*, vol. **126**, pp. 1719–1724, 1998.
- Cao, Y., Zhu, J., Navon, I.M., and Luo, Z., A Reduced-Order Approach to Four-Dimensional Variational Data Assimilation using Proper Orthogonal Decomposition, *Int. J. Numer. Meth. Fluids*, vol. **53**, no. 10, pp. 1571–1583, 2007.
- Chada, N.K., Jasra, A., and Yu, F., Multilevel Ensemble Kalman-Bucy Filters, *arXiv preprint arXiv:2011.04342*, 2020.
- Chernov, A., Hoel, H., Law, K.J.H., Nobile, F., and Tempone, R., Multilevel Ensemble Kalman Filtering for Spatio-Temporal Processes, *Numerische Mathematik*, vol. **147**, no. 1, pp. 71–125, 2021.
URL <https://doi.org/10.1007/s00211-020-01159-3>
- Computational Science Laboratory, ODE test problems, , 2020.
URL <https://github.com/ComputationalScienceLaboratory/ODE-Test-Problems>
- Constantinescu, E., Sandu, A., Chai, T., and Carmichael, G., Assessment of ensemble-Based chemical data assimilation in an idealized setting, *Atmospheric Environment*, vol. **41**, no. 1, pp. 18–36, 2007a.
URL <http://dx.doi.org/10.1016/j.atmosenv.2006.08.006>
- Constantinescu, E., Sandu, A., Chai, T., and Carmichael, G., Ensemble-Based chemical data assimilation. I: General approach, *Quarterly Journal of the Royal Meteorological Society*, vol. **133**, no. 626, pp. 1229–1243, 2007b.
URL <http://dx.doi.org/10.1002/qj.76>

- Constantinescu, E., Sandu, A., Chai, T., and Carmichael, G., Ensemble-Based chemical data assimilation. II: Covariance localization, *Quarterly Journal of the Royal Meteorological Society*, vol. **133**, no. 626, pp. 1245–1256, 2007c.
URL <http://dx.doi.org/10.1002/qj.77>
- Evensen, G., Sequential Data Assimilation with a Nonlinear Quasi-Geostrophic Model using Monte Carlo Methods to Forecast Error Statistics, *Journal of Geophysical Research: Oceans*, vol. **99**, no. C5, pp. 10143–10162, 1994.
- Evensen, G., *Data Assimilation: the Ensemble Kalman Filter*, Springer Science & Business Media, 2009.
- Farrell, B.F. and Ioannou, P.J., State Estimation using a Reduced-Order Kalman Filter, *Journal of the Atmospheric Sciences*, vol. **58**, no. 23, pp. 3666–3680, 2001.
- Giles, M.B., Multilevel Monte Carlo Path Simulation, *Operations Research*, vol. **56**, no. 3, pp. 607–617, 2008.
- Giles, M.B., Multilevel Monte Carlo Methods, *Acta Numerica*, vol. **24**, pp. 259–328, 2015.
- Hoel, H., Law, K.J.H., and Tempone, R., Multilevel Ensemble Kalman Filtering, *SIAM Journal on Numerical Analysis*, vol. **54**, no. 3, 2016.
- Hoel, H., Shaimerdenova, G., and Tempone, R., Multilevel Ensemble Kalman Filtering based on a Sample Average of Independent EnKF Estimators, *Foundations of Data Science*, p. 0, 2019.
- Jaynes, E.T., *Probability Theory: The Logic of Science*, Cambridge university press, 2003.
- Kalnay, E., *Atmospheric Modeling, Data Assimilation, and Predictability*, Cambridge Univ Pr, 2003.
- Karimi, A. and Paul, M.R., Extensive Chaos in the Lorenz-96 Model, *Chaos: An interdisciplinary journal of nonlinear science*, vol. **20**, no. 4, p. 043105, 2010.
- Karpatne, A., Atluri, G., Faghmous, J.H., Steinbach, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N., and Kumar, V., Theory-Guided Data Science: A New Paradigm for Scientific Discovery from Data, *IEEE Transactions on knowledge and data engineering*, vol. **29**, no. 10, pp. 2318–2331, 2017.
- Kingma, D.P. and Ba, J., Adam: A Method for Stochastic Optimization, *arXiv preprint arXiv:1412.6980*, 2014.
- Lee, J.A. and Verleysen, M., *Nonlinear Dimensionality Reduction*, Springer Science & Business Media, 2007.
- Lee, K. and Carlberg, K.T., Model Reduction of Dynamical Systems on Nonlinear Manifolds using Deep Convolutional Autoencoders, *Journal of Computational Physics*, vol. **404**, p. 108973, 2020.
- Lorenz, E.N., Predictability: A Problem Partly Solved, *Proc. Seminar on predictability*, Vol. 1, 1996.
- Moosavi, A., Attia, A., and Sandu, A., Tuning Covariance Localization using Machine Learning, *Machine Learning and Data Assimilation for Dynamical Systems track, International Conference on Computational Science ICCS 2019*, R.J. et al., Ed., Vol. 11539 of *Lecture Notes in Computer Science*, Faro, Algarve, Portugal, pp. 199–212, 2019.
URL https://doi.org/10.1007/978-3-030-22747-0_16
- Mou, C., Wang, Z., Wells, D.R., Xie, X., and Iliescu, T., Reduced Order Models for the Quasi-Geostrophic Equations: A Brief Survey, *Fluids*, vol. **6**, no. 1, p. 16, 2021.
- Nelson, B.L., On Control Variate Estimators, *Computers & Operations Research*, vol. **14**, no. 3, pp. 219–225, 1987.
- Nino-Ruiz, E. and Sandu, A., Ensemble Kalman Filter Implementations based on Shrinkage Covariance Matrix Estimation, *Ocean Dynamics*, vol. **65**, no. 11, pp. 1423–1439, 2015.
URL <http://dx.doi.org/10.1007/s10236-015-0888-9>
- Nino-Ruiz, E. and Sandu, A., An Ensemble Kalman Filter Implementation based on Modified Cholesky Decomposition for Inverse Covariance Matrix Estimation, *SIAM Journal on Scientific Computing*, vol. **40**, no. 2, pp. A867–A886, 2018.

- Nino-Ruiz, E. and Sandu, A., Efficient Parallel Implementation of DDDAS Inference using an Ensemble Kalman Filter with Shrinkage Covariance Matrix Estimation, *Cluster Computing*, vol. **22**, pp. 2211–2221, 2019.
URL <https://doi.org/10.1007/s10586-017-1407-1>
- Owen, A.B., *Monte Carlo Theory, Methods and Examples*, 2013.
- Peherstorfer, B., Willcox, K., and Gunzburger, M., Survey of Multifidelity Methods in Uncertainty Propagation, Inference, and Optimization, *Siam Review*, vol. **60**, no. 3, pp. 550–591, 2018.
- Petrie, R., Localization in the Ensemble Kalman Filter, *MSc Atmosphere, Ocean and Climate University of Reading*, 2008.
- Popov, A.A., Mou, C., Sandu, A., and Iliescu, T., A Multifidelity Ensemble Kalman Filter with Reduced Order Control Variates, *SIAM Journal on Scientific Computing*, vol. **accepted**, no. TBD, p. in press, 2021.
- Popov, A.A. and Sandu, A., A Bayesian Approach to Multivariate Adaptive Localization in Ensemble-Based Data Assimilation with Time-Dependent Extensions, *Nonlinear Processes in Geophysics*, vol. **26**, pp. 109–122, 2019.
URL <https://doi.org/10.5194/npg-26-109-2019>
- Popov, A.A. and Sandu, A., An explicit probabilistic derivation of inflation in a scalar ensemble Kalman filter for finite step, finite ensemble convergence, , 2020.
- Popov, A.A. and Sandu, A., 2021. Multifidelity data assimilation for physical systems, *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications (Vol. IV)*. S.K. Park and L. Xu, Eds. Springer, p. in press.
- Reich, S. and Cotter, C., *Probabilistic Forecasting and Bayesian Data Assimilation*, Cambridge University Press, 2015.
- Roberts, S., Popov, A.A., and Sandu, A., ODE test problems: a MATLAB suite of initial value problems, , 2019.
- Rubinstein, R.Y. and Marcus, R., Efficiency of Multivariate Control Variates in Monte Carlo Simulation, *Operations Research*, vol. **33**, no. 3, pp. 661–677, 1985.
- San, O. and Iliescu, T., A Stabilized Proper Orthogonal Decomposition Reduced-Order Model for Large Scale Quasigeostrophic Ocean Circulation, *Adv. Comput. Math.*, pp. 1289–1319, 2015.
- Sell, G.R. and You, Y., *Dynamics of Evolutionary Equations*, Vol. 143, Springer Science & Business Media, 2013.
- Sirovich, L., Turbulence and the Dynamics of Coherent Structures. I. Coherent Structures, *Quarterly of applied mathematics*, vol. **45**, no. 3, pp. 561–571, 1987.
- Strogatz, S.H., *Nonlinear Dynamics and Chaos: with Applications to Physics, Biology, Chemistry, and Engineering*, CRC Press, 2018.
- Willard, J., Jia, X., Xu, S., Steinbach, M., and Kumar, V., Integrating Physics-Based Modeling with Machine Learning: A Survey, *arXiv preprint arXiv:2003.04919*, 2020.