

# Machine Learning Techniques to Construct Patched Analog Ensembles for Data Assimilation

L. Minah Yang\*, Ian Grooms\*

*Department of Applied Mathematics, University of Colorado, Boulder, CO 80309*

---

## Abstract

Using generative models from the machine learning literature to create artificial ensemble members for use within data assimilation schemes has been introduced in [1, Grooms QJRMS, 2020] as constructed analog ensemble optimal interpolation (cAnEnOI). Specifically, we study general and variational autoencoders for the machine learning component of this method, and combine the ideas of constructed analogs and ensemble optimal interpolation in the data assimilation piece. To extend the scalability of cAnEnOI for use in data assimilation on complex dynamical models, we propose using patching schemes to divide the global spatial domain into digestible chunks. Using patches makes training the generative models possible and has the added benefit of being able to exploit parallelism during the generative step. Testing this new algorithm on a 1D toy model, we find that larger patch sizes make it harder to train an accurate generative model (i.e. a model whose reconstruction error is small), while conversely the data assimilation performance improves at larger patch sizes. There is thus a sweet spot where the patch size is large enough to enable good data assimilation performance, but not so large that it becomes difficult to train an accurate generative model. In our tests the new patched cAnEnOI method outperforms the original (unpatched) cAnEnOI, as well as the ensemble square root filter results from [1].

*Keywords:* data assimilation, machine learning, analogs

---

## 1. Introduction

Data assimilation is widely used in geophysical problems as a reliable method to combine theoretical knowledge of the dynamical system at hand with the available real-life observations. The main data assimilation methods are split into those that require an ensemble such as the ensemble Kalman filter (EnKF, see [2]), and variational methods such as 3D-Var and 4D-Var [3]. As ensemble forecasts can be prohibitively expensive, the use of time-independent ensemble perturbations to reduce costs has been introduced as ensemble optimal interpolation (EnOI, see [2, 4]) and implemented in ocean data assimilation (see, e.g., [5]). Naturally, this is less accurate than using an ensemble forecast and to remedy this, it was proposed in [1] to use analog ensembles (AnEnOI) which are time-dependent. Analogs are model states that are close to the current forecast; in AnEnOI analogs of the forecast are found from a library of model states and used to form an analog ensemble. The use of analogs in AnEnOI is a departure from their traditional role in forecasting, where analogs are used to produce a forecast (see [6–8]). Moreover, [1] introduced a new method of using variational autoencoders (VAEs, a machine learning model in [9]) to construct analogs for AnEnOI, rather than finding analogs in a library. This method is referred to as cAnEnOI and it has been shown to produce better results than both EnOI and AnEnOI when tested on a multiscale variant [10] of the Lorenz '96 model [11].

---

\*Corresponding author

*Email address:* `lucia.yang@colorado.edu` (L. Minah Yang)

In this paper, we introduce a patching scheme to the analog construction component of cAnEnOI which we call p-cAnEnOI. As geophysical applications often require high resolution representations of complex dynamical structures, it is likely difficult for a single VAE to capture all of the significant information over the entire spatial domain. For example, consider training a generative model to generate realistic states of an eddy resolving ocean model. A single three-dimensional variable (e.g. temperature) from the global ocean simulation in [12] at nominal 0.1-degree resolution requires approximately 4 GB of memory. There are four three-dimensional variables in a single model state (for this model; others may have more), not to mention the other two-dimensional variables like sea-surface height, and the resolution of ocean models is continually increasing. Not only would training a generative model for something like this require massive computing powers in terms of computations during training, simply moving batches of samples around during training could cause issues in the communication sector.

To avoid the difficulty of dealing with the global model state, we simply divide the domain into equal-sized patches, and train a generative model to reproduce a single patch. A global model state can be constructed out of patches. Some benefits of p-cAnEnOI include more robust training due to simpler complexity of the machine learning model and increase in the number of training samples. Also, when constructing the global model state, patches can be constructed in parallel.

We test p-cAnEnOI in the same multiscale Lorenz '96 model as [1]. We vary the patch sizes in p-cAnEnOI and study how that affects the performance of the machine learning models with respect to reconstruction error (the ability of the VAE to exactly reconstruct the input model state), and the performance of the resulting data assimilation scheme. Using smaller patch sizes makes training quicker and results in VAEs with better (lower) reconstruction errors, but introduces more discontinuities across patch boundaries which can degrade the spatial correlations in the analog ensemble.

While cAnEnOI in [1] used only VAEs to construct analogs, we introduce the use of general (non-variational) autoencoders in constructing analogs. These are quicker to train due to their simpler structure but require additional steps to use as a generative model: the covariance of the training data encoded into the latent space must be computed then factorized using the Cholesky decomposition; but these extra operations are not prohibitive due to the small size of the encoded latent space. We find that p-cAnEnOI with general autoencoders performs just as well as the variant that uses VAEs instead.

Overall, we find that p-cAnEnOI can outperform the analog methods described in [1] as well as an ensemble square root filter, which costs 100 times more for an ensemble size of 100. In the toy model, we found that the largest patch size (equal to 1/4 of the domain) produced the best results, and p-cAnEnOI with patch sizes of 1/8 and 1/16 still performed better than global cAnEnOI but patch size 1/32 performed worse despite having the best reconstruction error. This indicates that the patch size has an effect on the data assimilation quality where it must be large enough to preserve spatial correlations but small enough to allow for efficient training.

## 2. Background

In this section, we summarize the relevant data assimilation and machine learning methods in order to sufficiently describe the new method, p-cAnEnOI, in section 3. This new method is a slight modification of cAnEnOI, which inserts constructed analogs (typically used in forecasting) in ensemble-based data assimilation methods such as ensemble optimal interpolation (EnOI). In addition, variational autoencoders, a machine learning model, is used to construct the analogs. We discuss the data assimilation methods in section 2.1, the machine learning methods in sections 2.2.1 and 2.2.2, cAnEnOI in section 2.3, and the toy model in section 2.4.

### 2.1. Ensemble-based Data Assimilation Methods

In this section, we discuss some relevant ensemble-based data assimilation methods. The Ensemble Kalman Filter (EnKF, see [13–15]) takes the Monte Carlo approach to approximate the background error covariance matrix by forecasting an ensemble of state variables and computing its sample covariance matrix.

This corresponds to forming  $\mathbf{B}$  in eq. (1). The Kalman update formula for the  $i^{th}$  ensemble member is

$$\mathbf{x}_a^{(i)} = \mathbf{x}_b^{(i)} + \mathbf{K}(\mathbf{y}^{(i)} - \mathbf{H}\mathbf{x}_b^{(i)}), \quad \mathbf{K} = \mathbf{B}\mathbf{H}^\top(\mathbf{H}\mathbf{B}\mathbf{H}^\top + \mathbf{R})^{-1}, \quad (1)$$

where  $\mathbf{x}_a^{(i)}$  is the updated ensemble member,  $\mathbf{x}_b^{(i)}$  is the forecast of the  $i^{th}$  ensemble member from the previous time step,  $\mathbf{y}^{(i)}$  is the perturbed observation where the perturbations have zero mean and covariance  $\mathbf{R}$ , and  $\mathbf{H}$  is the observation model. Accurately approximating the background covariance matrix in this approach may require the size of the ensemble to be large, which then increases the number of model forecasts to be made. The cost of the EnKF update is usually small compared to the cost of a forecast, so the total cost of using the EnKF increases linearly with respect to the size of the ensemble,  $N_e$ , and may quickly become impractical. The methods we discuss in this section are motivated by reducing the computational load of forecasting  $N_e$  ensemble members. Since we compare our results in section 4 to the numerical results of [1], the analog methods discussed in the aforementioned paper are included. The quality of the various data assimilation methods is evaluated by the root mean squared error (RMSE) computed across  $T$  assimilation cycles as shown ineq. (2),

$$RMSE = \frac{1}{T} \sum_{t=1}^T \left( \frac{1}{N_x} \sum_{j=1}^{N_x} (RS_j^{(t)} - AM_j^{(t)})^2 \right)^{1/2}, \quad (2)$$

where the difference between the reference simulation  $RS$  and the analysis mean  $AM$  is averaged for a state with  $N_x$  spatial nodes and the subscript and parenthesized superscripts each refer to spatial and temporal indices.

### 2.1.1. Ensemble Square Root Filter

Ensemble square root filters (reviewed in [16]) modify the EnKF to ensure that the posterior covariance exactly satisfies the Kalman filter update formula

$$\mathbf{B}_a = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{B}_f \quad (3)$$

and avoids having to build an ensemble of perturbed observations. In particular, [1] implements the serial ensemble square root filter of [17], (referred to as ESRF hereafter) to serve as a point of reference for the performance of data assimilation methods that use an ensemble forecast at every step. Localization is added to zero out spuriously high covariances across large distances using the following localization function

$$\ell(d_i) = \exp\left(-\frac{1}{2}\left(\frac{d_i}{L}\right)^2\right), \quad (4)$$

where  $d_i$  is the distance from the observation being assimilated to the  $i^{th}$  state variable, and  $L$  is the localization radius, a tunable parameter. This can be efficiently applied with elementwise multiplication within ESRF. In addition, inflation is applied to the posterior as suggested in [18] by multiplying the background covariance matrix by some  $r \geq 1$ . Since the optimal results for ESRF with an ensemble of size  $N_e = 200$  were not much better than that with  $N_e = 100$  in [1], we use  $N_e = 100$  for the remainder of this paper. No new ESRF experiments were performed here; we only quote ESRF results from [1].

### 2.1.2. EnOI and Analogs

Ensemble optimal interpolation (EnOI, see [2, 4]) replaces the background covariance matrix in eq. (1) with the sample covariance computed from a fixed set of states drawn from a catalog of model states. This can be interpreted as approximating the background forecast error with the covariance present in a climatological time scale and is time-independent since the same set of states is used for every assimilation cycle. Since the climatological spread is likely larger than what is expected in the error in a single forecast step, we scale the constructed  $\mathbf{B}$  to have a specified forecast spread  $f$ , which is another tunable parameter, equal to the square root of the mean along the diagonal of  $\mathbf{B}$ .

Analogs refer to model states that are close to the current state and they have been used in forecasting (see [6–8, 19]) following its introduction in [20]. They have primarily been used for forecasting in the context of a catalog of data that contain correlated pairs  $C = \{(\mathbf{x}^{(i)}, \mathbf{x}^{(i+1)})\}_{i=1, \dots}$ . Basic analog forecasting is performed by finding a pair in the catalog such that  $\mathbf{x}^{(i)}$  matches the current state  $\mathbf{x}^*$  as closely as possible, and then using the corresponding component  $\mathbf{x}^{(i+1)}$  as the forecast. Constructed analogs generalize this method by finding many analog pairs in  $C$  whose  $\mathbf{x}^{(i)}$  components are close to  $\mathbf{x}^*$ , and constructing the forecast with some weighted average of the corresponding  $\mathbf{x}^{(i+1)}$  component of the pairs. Analog forecasting is not used in AnEnOI, cAnEnOI, or p-cAnEnOI, though it has been used in data assimilation in [21].

It was proposed in [1] to replace the time-independent ensemble perturbations in EnOI with analog ensemble perturbations. The resulting method (AnEnOI) makes a single forecast  $\mathbf{x}^{(i)}$  (using the dynamical model, not an analog forecast), and then finds  $N_e$  nearest-neighbors in a catalog. The mean of the ensemble of nearest neighbors is then subtracted and replaced by the actual forecast, resulting in an ensemble whose mean is the single forecast and whose perturbations are drawn from a set of model states that are similar to the forecast. The ensemble covariance matrix  $\mathbf{B}$  is then scaled to have forecast spread  $f$ , which is a tunable parameter. AnEnOI exhibited approximately 15% improvement in performance compared to EnOI in the tests of [1].

Next, we introduce autoencoders and variational encoders in section 2.2 before continuing with the descriptions of the data assimilation methods in section 2.3.

## 2.2. Autoencoders and Variational Autoencoders

While [1] exclusively used variational autoencoders as the machine learning component of the constructed analog method, we introduce a method that uses a general autoencoder. We briefly summarize both machine learning methods in sections 2.2.1 and 2.2.2.

### 2.2.1. General Autoencoders

Autoencoders are a type of unsupervised learning in which deep neural networks are trained to learn an encoding of the data. To achieve this, the output of the neural network is set to the input, and the hidden layers of the neural network form two distinct components: the encoder and the decoder. The encoder is comprised of (hidden) layers that code the input variables into the latent space. The structure of the encoder is often designed to result in a latent space with fewer dimensions than that of the input space, performing a dimension reduction of a sort, though dimension reduction is not the goal here. The decoder transforms the encoded variable back to its original input through more layers. The encoder and decoder pair are shown as two mappings in eqs. (5) and (6),

$$e(\mathbf{x}) = \mathbf{z}, \quad \mathbf{x} \in \mathbb{R}^{\text{dsize}}, \quad \mathbf{z} \in \mathbb{R}^{\text{latentDim}} \quad (5)$$

$$d(\mathbf{z}) = \hat{\mathbf{x}}, \quad \hat{\mathbf{x}} \in \mathbb{R}^{\text{dsize}}, \quad (6)$$

$\mathbf{x}$  is the original variable,  $\mathbf{z}$ , its code, and  $\hat{\mathbf{x}}$ , the reconstructed variable, and  $\text{latentDim} \leq \text{dsize}$ . In practice, an autoencoder can be thought of and trained as a normal deep neural network with a special hidden layer that has lower dimension than the size of the input variable. Figure 1 illustrates the architecture of an autoencoder.

Autoencoders can be used for specialized lossy compression (meaning exact recovery is not guaranteed) or dimension reduction tasks for specific datasets they were trained on (see [22, 23]). Their use in anomaly detection ([24] and image processing (super resolution [25]; low light enhancement [26]) has also been studied. We use autoencoders with convolutional layers for the encoder and the decoder halves.

Many types of autoencoders have been developed for use in data science applications, i.e. sparse autoencoders for classification and variational autoencoders (VAEs) as a generative model. We describe VAEs in the next section, and refer to autoencoders that are not VAEs as general autoencoders.

### 2.2.2. VAE

A variational autoencoder is a type of an autoencoder that imposes extra conditions to regularize the latent space (see [9, 27]) to resemble a Gaussian distribution. These extra conditions are imposed via the

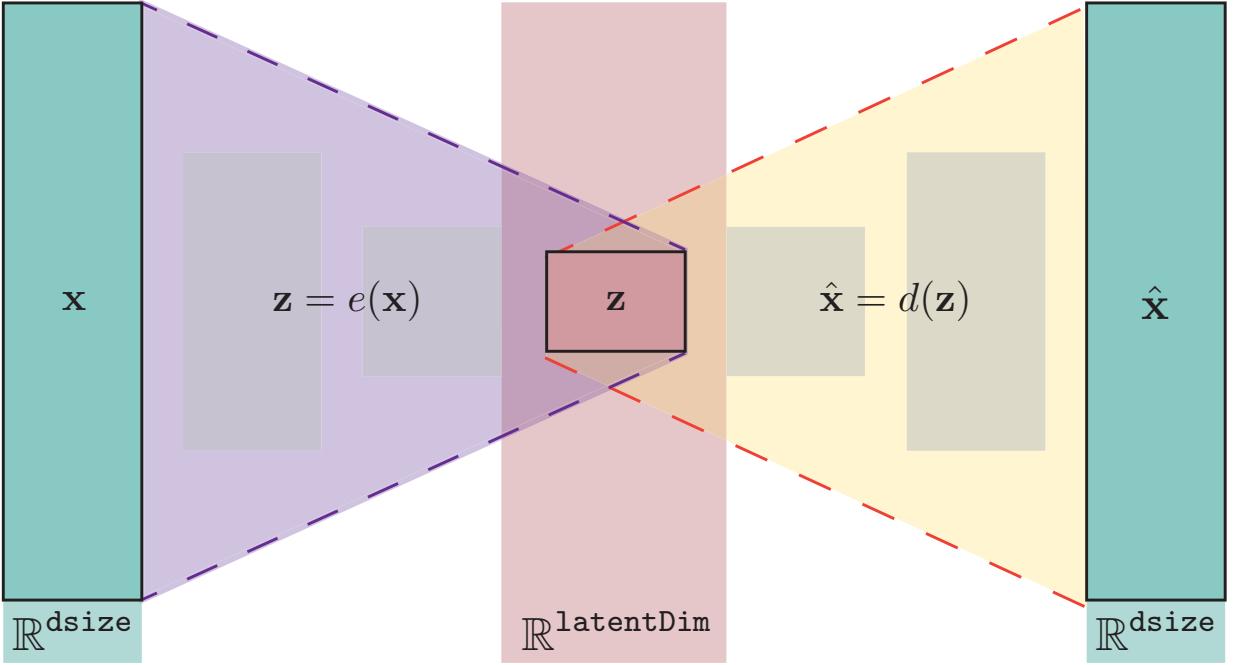


Figure 1: Example architecture of a general autoencoder. Purple region represents the encoder, and the yellow region represents the decoder.

reparametrization trick (see [9]), described in eqs. (8) and (9) and by adding an extra loss term called the Kullback-Leibler (KL) divergence, which measures the distance from one probabilistic distribution to another. The full loss function is shown in eq. (7), where the KL divergence represents the distance from the multivariate standard normal distribution.

$$\text{loss}(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{2} \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 - \frac{1}{2} \sum_{k=1}^{\text{latentDim}} (1 + (\log(\mathbf{z}_\Sigma^2))_k - (\mathbf{z}_\mu^2)_k - (\mathbf{z}_\Sigma)_k) \quad (7)$$

In contrast to eq. (5), the encoder in a VAE outputs two variables of size `latentDim`. Equation (8) shows the VAE variant of an encoder that encodes the input variable  $\mathbf{x}$  into two variables,  $\mathbf{z}_\mu$  and  $\mathbf{z}_{\log(\Sigma)}$ , which are then combined via eq. (9) to mimic sampling from a normal distribution.

$$e(\mathbf{x}) = [\mathbf{z}_\mu \ \mathbf{z}_\Sigma] \in \mathbb{R}^{\text{latentDim} \times 2} \quad (8)$$

$$\mathbf{z} = \mathbf{z}_\mu + \epsilon \circ \mathbf{z}_\Sigma, \quad \epsilon \sim N(\mathbf{0}, \mathbf{I}), \quad (9)$$

where  $\circ$  represents an elementwise product. (We also denote  $e(\mathbf{x})_\mu := \mathbf{z}_\mu$  and  $e(\mathbf{x})_{\log(\Sigma)} := \mathbf{z}_{\log(\Sigma)}$ .) This encoded state,  $\mathbf{z}$  is then fed through the decoder, which is constructed exactly as in eq. (6).

The regularized structure of the latent space guarantees that a small perturbation added in the latent space should result in small differences once decoded. That is,  $d(e(\mathbf{x}) + \epsilon) - d(e(\mathbf{x}))$  should be reasonably small for a small  $\epsilon$ . Figure 2 provides a visualization of the structure of VAEs as a basis of comparison against fig. 1.

### 2.3. cAnEnOI

In [1], the author suggested using VAEs as a generative model for constructing analogs in the following way. To generate  $N_e$  ensemble members close to the current forecast state,  $\mathbf{x}$ ,  $N_e$  distinct, independent, samples of random noise with amplitude  $r_z$  would be added to the mean encoded state ( $e(\mathbf{x})_\mu \equiv \mathbf{z}_\mu$ ) in

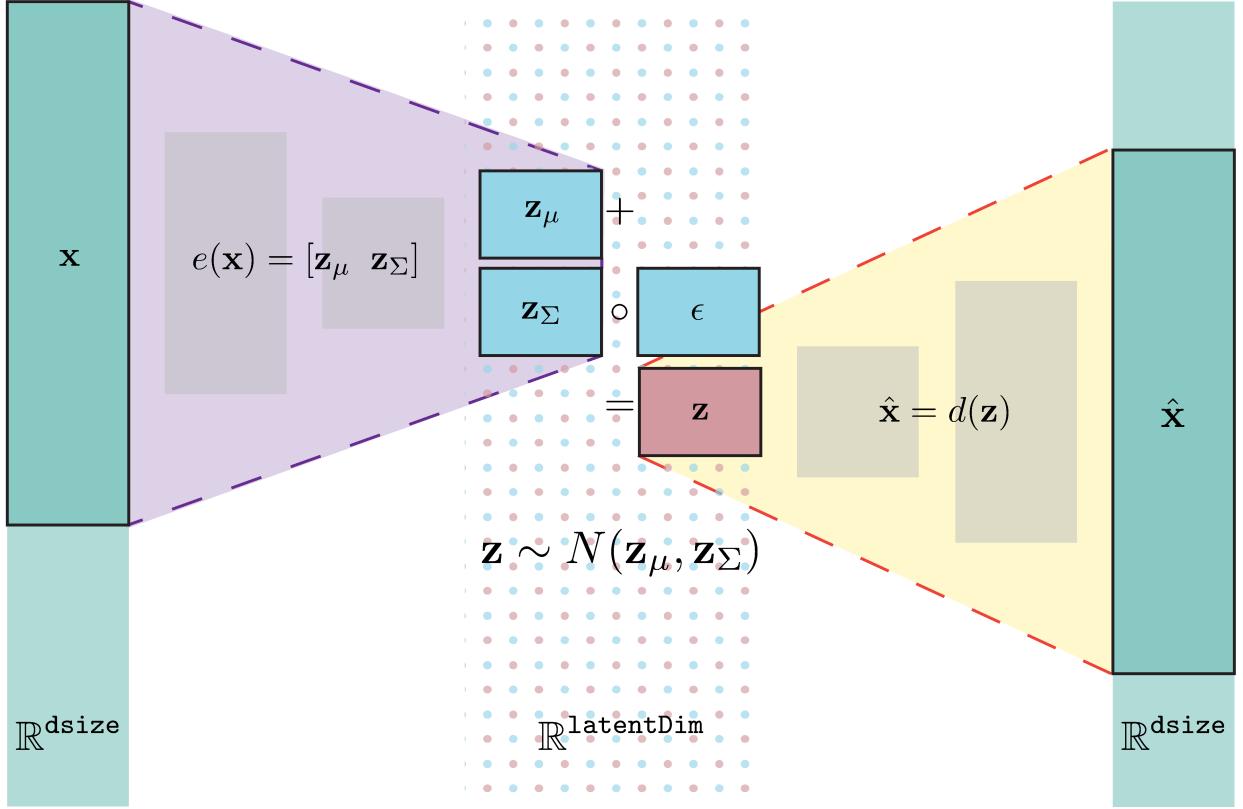


Figure 2: Example architecture of a variational autoencoder (VAE).

latent space and then decoded to return to the original dimension, as shown in eq. (10). Each ensemble member is perturbed by a unique random noise, denoted by  $\epsilon^{(i)}$  for  $i = 1, \dots, N_e$ .

$$\mathbf{x}^{(i)} = d\left(e(\mathbf{x})_\mu + r_z \epsilon^{(i)}\right), \quad \epsilon^{(i)} \sim N(\mathbf{0}, \mathbf{I}), \quad i \in \{1, \dots, N_e\} \quad (10)$$

The sample covariance matrix of these constructed analogs approximates the background covariance matrix in eq. (1). The amplitude  $r_z$  of the noise in latent space is a tunable parameter. In section 3.2 we propose a slightly different version of eq. (10) to construct an analog ensemble using a general autoencoder.

#### 2.4. 1D Toy Model: Multiscale Lorenz '96

The multiscale Lorenz '96 model introduced in [10] is described by a single system of ODEs in eq. (11)

$$\dot{\mathbf{x}} = h \mathbf{N}_S(\mathbf{x}) + J \mathbf{T}^\top \mathbf{N}_L(\mathbf{T}\mathbf{x}) - \mathbf{x} + F \mathbb{1}, \quad (11)$$

where  $h, F \in \mathbb{R}$  are coupling and forcing parameters,  $J \in \mathbb{N}$  is the ratio between the two scales, and  $\mathbb{1}$  is a vector of ones. The small and large scale nonlinearities are described by eqs. (12) and (13), and  $\mathbf{T}$  is a mapping that projects  $\mathbf{x}$  onto the largest  $K$  Fourier modes and interpolates the result to  $K$  equally-spaced points. The nonlinearities are defined by

$$\mathbf{N}_S(\mathbf{x})_i = -\mathbf{x}_{i+1}(\mathbf{x}_{i+2} - \mathbf{x}_{i-1}), \quad i \in \{1, \dots, JK\}, \quad (12)$$

$$\mathbf{N}_L(\mathbf{X})_k = -\mathbf{X}_{k-1}(\mathbf{X}_{k-2} - \mathbf{X}_{k+1}), \quad k \in \{1, \dots, K\} \quad (13)$$

where the indices are periodic, i.e.  $\mathbf{x}_{i+JK} = \mathbf{x}_i$  and  $\mathbf{X}_{k+K} = \mathbf{X}_k$ . We use this model exclusively to train the VAEs and to test our new method, patched-cAnEnOI (p-cAnEnOI), in section 4. Figure 3 shows the

evolution of the model eq. (11) configured with  $h = 0.5$ ,  $F = 8$ ,  $J = 64$ , and  $K = 41$ , and initialized from a standard normal distribution at time 0. The multiscale model typically displays about 8 large-scale waves in the domain, similar to the standard Lorenz ‘96 model, with additional small-scale instabilities appearing intermittently in time and space.

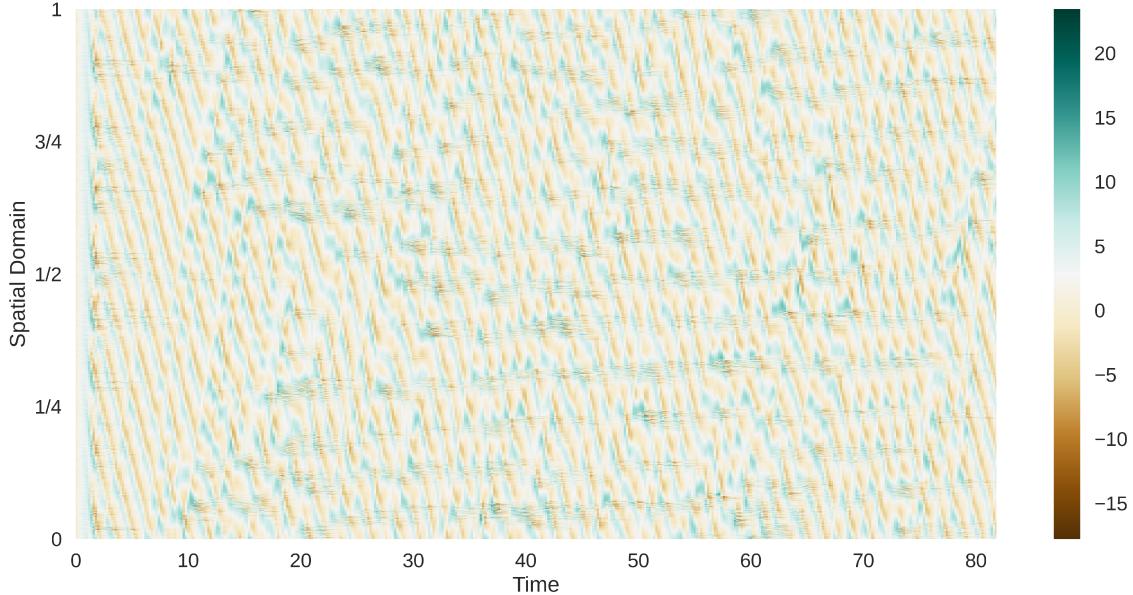


Figure 3: Evolution of the multiscale Lorenz-96 model initialized by a sample from a standard normal distribution at time  $T = 0$  and run until final time  $T = 81.8$ .

### 3. Method: Patched Analogs (p-cAnEnOI)

Training a VAE that learns the dynamics over the entire spatial domain for use in cAnEnOI is feasible for small problems such as the multiscale Lorenz ‘96 from [10] as is demonstrated in [1]. However, a direct implementation of cAnEnOI on more complex, high resolution models poses several problems. First, the number of parameters being trained for convolutional neural networks grows large as the size of the input variable increases, making training computationally expensive. Second, training a VAE to learn the complex dynamics of the global spatial domain may require extremely large libraries of training samples that may not be available. For example, a GCM may have several dynamically significant locations that each deserve careful analysis but the global variable as a whole is too specific of a state for a VAE to be trained to learn. Our new method, p-cAnEnOI, attempts to address these issues. We describe the method in detail in section 3.1 and discuss its impact on the covariance matrix within eq. (1) in section 3.3.

#### 3.1. Constructing Patched Analogs with VAEs

This new method modifies how analogs are generated at each assimilation cycle. Recall that to generate  $N_e$  analogs, cAnEnOI uses a VAE to encode the state variable at time of assimilation, add  $N_e$  distinct random noise, and decode those perturbations. Instead, p-cAnEnOI first partitions the state variable of size `dsize` into  $p$  equal sized patches of size `psize=dsize/p` given that  $p$  is a divisor of `dsize`. These  $p$  patches are then individually encoded with a VAE whose input and output variables correspond to `psize`. To generate  $N_e$  ensemble members,  $N_e$  distinct random noises are added to each encoded patch, then decoded. The decoded patch perturbations are rearranged to cover the original spatial domain.

Consider a trivial partitioning scheme which places patch boundaries to the left of locations  $\{k \times \text{psize}\}_{k=1}^{p-1}$ . This forces discontinuities only at those  $(p-1)$  locations for all constructed analogs and unevenly distributes the errors associated with these discontinuities across the domain. To remedy this, we shift the locations of the patch boundaries. Since the toy model is periodic, any shift  $s$  in  $\{0, \dots, \text{psize} - 1\}$  creates a unique set of  $p$  patch boundaries at  $\{s + k \times \text{psize}\}_{k=0}^{p-1}$ . That is, the  $p$  patches of a state  $\mathbf{x}$  with a right shift  $s$  range over the nodes as indicated in eq. (14), where  $\mathbf{x}^{(j,s)}$  is the  $j^{th}$  patch created with shift  $s$  and  $\mathbf{x}_k$  is the  $k^{th}$  element of  $\mathbf{x}$ .

$$\mathbf{x}^{(j,s)} = \begin{cases} [\mathbf{x}_{s+j \times \text{psize}+1}, \dots, \mathbf{x}_{s+(j+1) \times \text{psize}}], & j = 0, \dots, p-2 \\ [\mathbf{x}_{s+(j-1) \times \text{psize}+1}, \dots, \mathbf{x}_{\text{dsiz}}], \mathbf{x}_1, \dots, \mathbf{x}_s], & j = p-1 \end{cases} \quad (14)$$

The patching is undone by simply concatenating all  $p$  patches in order and by shifting to left by  $s$ . We discuss the impact of shifted partitioning on the covariance matrix in section 3.3, patching in non-periodic domains in section 5, and alg. 1 formalizes this new method of constructing analogs.

---

**Algorithm 1:**  $\{\mathbf{a}^{[i]}\}_{i=1}^{N_e} = \text{gen\_patch\_ensemble}(\mathbf{x}, N_e, p, r_z, S)$ . This algorithm constructs  $N_e$  analogs by using a vae on  $pN_e$  patches with shifts in  $S$ .

---

<p><b>Input:</b> <math>\mathbf{x}, N_e, p, r_z, S</math></p> <p>1 <b>for</b> <math>i = 1 : N_e</math> <b>do</b></p> <p style="padding-left: 20px;">/* Partition <math>\mathbf{x}</math> into <math>p</math> patches using shift <math>s_i \in S</math> (See eq. (14)). */</p> <p style="padding-left: 20px;">2 <math>[\mathbf{x}^{(1,s_i)}, \dots, \mathbf{x}^{(p,s_i)}] \leftarrow \text{patch}(\mathbf{x}, s_i)</math></p> <p style="padding-left: 20px;">3 <b>for</b> <math>j = 1 : p</math> <b>do</b></p> <p style="padding-left: 40px;">4 <math>[z_\mu, z_{\log(\Sigma)}] \leftarrow e(\mathbf{x}^{(j,s_i)})</math></p> <p style="padding-left: 40px;">5 <math>\mathbf{x}^{(j,s_i)} \leftarrow d(z_\mu + r_z \epsilon^{(j,i)})</math>, <span style="float: right;">// <math>\epsilon^{(j,i)} \sim N(\mathbf{0}, \mathbf{I})</math></span></p> <p style="padding-left: 20px;">6 <math>\mathbf{a}^{[i]} \leftarrow \text{undo\_patching}(\{\mathbf{x}^{(j,s_i)}\}_{j=1}^p)</math></p> <p>7 <b>return</b> <math>\{\mathbf{a}^{[i]}\}_{i=1}^{N_e}</math></p>	<p><b>Output:</b> <math>\{\mathbf{a}^{[i]}\}_{i=1}^{N_e}</math></p>
---	---

---

Algorithm 1 assumes that we have a trained VAE ready to be used on the patches. This adds only few extra steps to the training portion of cAnEnOI: a new library must be generated from the original library of analogs. If  $1/p$  is the ratio between the patch size we wish to use and the size of the domain, then each model state in the original training library can generate at least  $p$  patches, creating a library that has at minimum  $p \times$  the number of elements of the original library. This library can now be split into training and validation sets, and the appropriate VAE can now be trained. Note that whatever patching scheme is used to generate this new library should be used within alg. 1 as well. Furthermore, there are many opportunities for parallelization in alg. 1. The inner loop can be computed simultaneously with  $p$  nodes, and even the outer loop can be computed in parallel if copies of model state  $\mathbf{x}$  can be quickly distributed.

As an example, consider implementing this to a state variable of the model eq. (11) with  $J = 64$  and  $K = 41$ , which has dimension 2624. Figure 4 shows four distinct ways of splitting the state variable into four equal-sized patches of size 656. Each color represents a single patch of size 1/4 of the original domain size in each of the four patching variants, and the gray vertical lines mark the boundaries between two adjacent patches. The patching variants use shifts of 0, 164, 328, and 492 units to the right. To retrieve the original variable, the patches are concatenated from orange to green to red to purple from left to right, then shifted by 0, 164, 328, and 492 units to the left. Note that 656 distinct patching variants exist for this particular example. Lastly, fig. 5 shows four distinct analog ensemble members generated from the same forecast. Each ensemble member was created from the corresponding patch shift, as well as having had different perturbations of amplitude  $r_z = 0.75$  added while in their encoded states. The different colors again represent the patches, the vertical lines represent the patch boundaries, and the grey waves in the background show the original forecast. Next, we discuss how a general autoencoder can be used instead of a VAE within p-cAnEnOI.

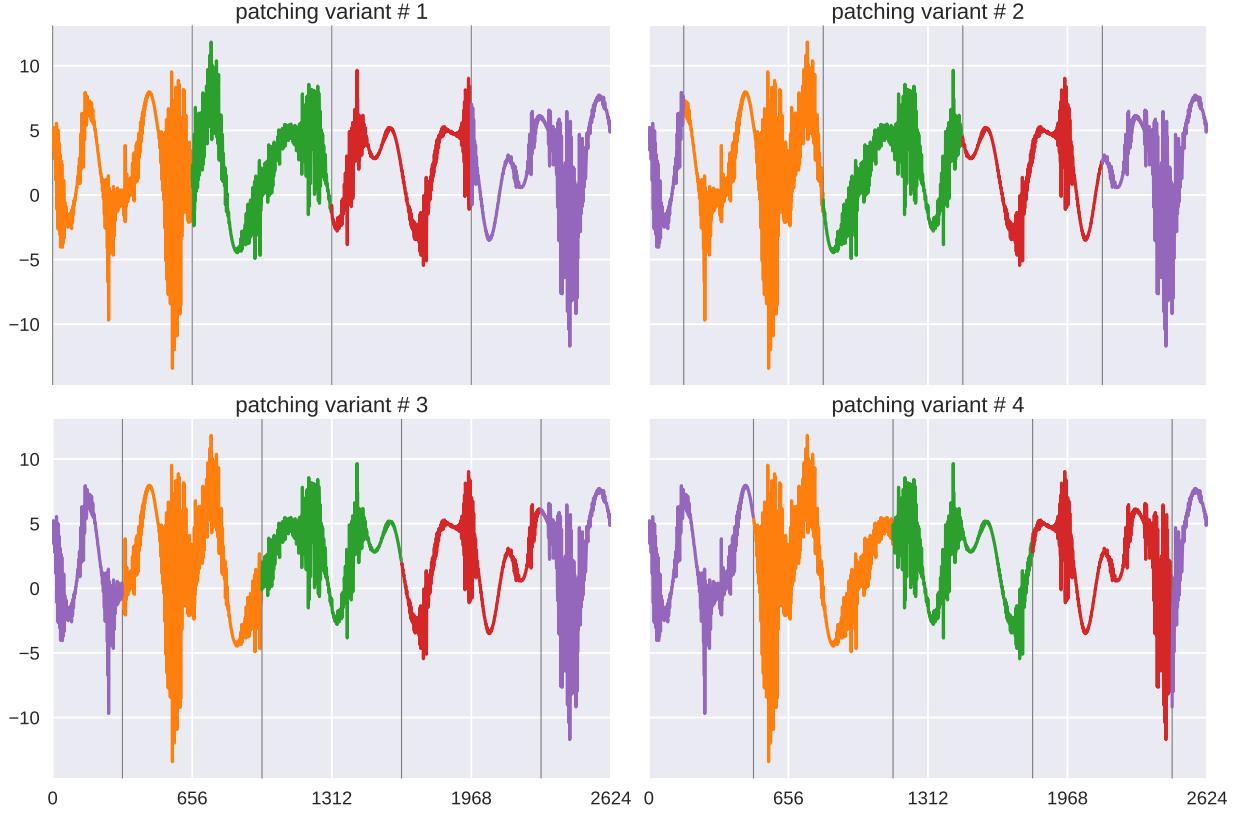


Figure 4: Four patching variants of one instance of the 1D toy model in eq. (11) with  $J = 64$  and  $K = 41$ .

### 3.2. Constructing Analogs with General Autoencoders

Recall that VAEs are autoencoders with extra conditions that force regularity in the latent space by adding an extra term in the loss function and require sampling via eq. (9). In theory, the optimal encoder of a VAE should map from the space of the dataset at hand to the standard normal distribution, whereas the optimal encoder of a general autoencoder does not guarantee any special structure of the latent space. This lack of regularity in the latent space makes general autoencoders a subpar candidate for generative modeling. For example, in alg. 1, we add noise sampled from  $N(\mathbf{0}, r_z^2 \mathbf{I})$  to the encoded state, since we can assume that the encoded state is a sample from  $N(\mathbf{0}, \mathbf{I})$ . For a general autoencoder there is no a priori way to predict the structure of the data distribution in latent space, and therefore no a priori way of knowing how much noise to add in latent space when constructing analogs. One possible way to set the structure of the noise in latent space is to mimic the approach for a variational autoencoder. Let  $\mathbf{Z}$  be a random variable corresponding to the latent space distribution of a general autoencoder, and let the mean and covariance of  $\mathbf{Z}$  be  $\mu$  and  $\Sigma$ . If  $\mathbf{L}$  is the Cholesky factor of  $\Sigma$ , then

$$\mathbf{Z}' = \mathbf{L}^{-1}(\mathbf{Z} - \mu) \quad (15)$$

has mean  $\mathbf{0}$  and covariance  $\mathbf{I}$ , analogous to the standard normal distribution in latent space that is targeted by a VAE. Since we generate analogs using the VAE by adding noise sampled from  $N(\mathbf{0}, r_z^2 \mathbf{I})$  to the encoded forecast state, it is natural for the general autoencoder to add noise sampled from  $N(\mathbf{0}, r_z^2 \mathbf{I})$  to  $\mathbf{Z}'$ . Let us call the encoded forecast state  $\mathbf{z} \in \mathbb{R}^{\text{latentDim}}$ , which should be a sample from the distribution  $\mathbf{Z}$ . Analogs in latent space  $\mathbf{z}^{(i)}$  are constructed by transforming  $\mathbf{z}$  to  $\mathbf{z}'$ , then adding noise, then converting back:

$$\mathbf{z}^{(i)} = \mathbf{L} \left( \mathbf{L}^{-1}(\mathbf{z} - \mu) + r_z \epsilon^{(i)} \right) + \mu = \mathbf{z} + r_z \mathbf{L} \epsilon^{(i)}, \quad \epsilon^{(i)} \sim N(\mathbf{0}, \mathbf{I}), \quad i \in \{1, \dots, N_e\}. \quad (16)$$

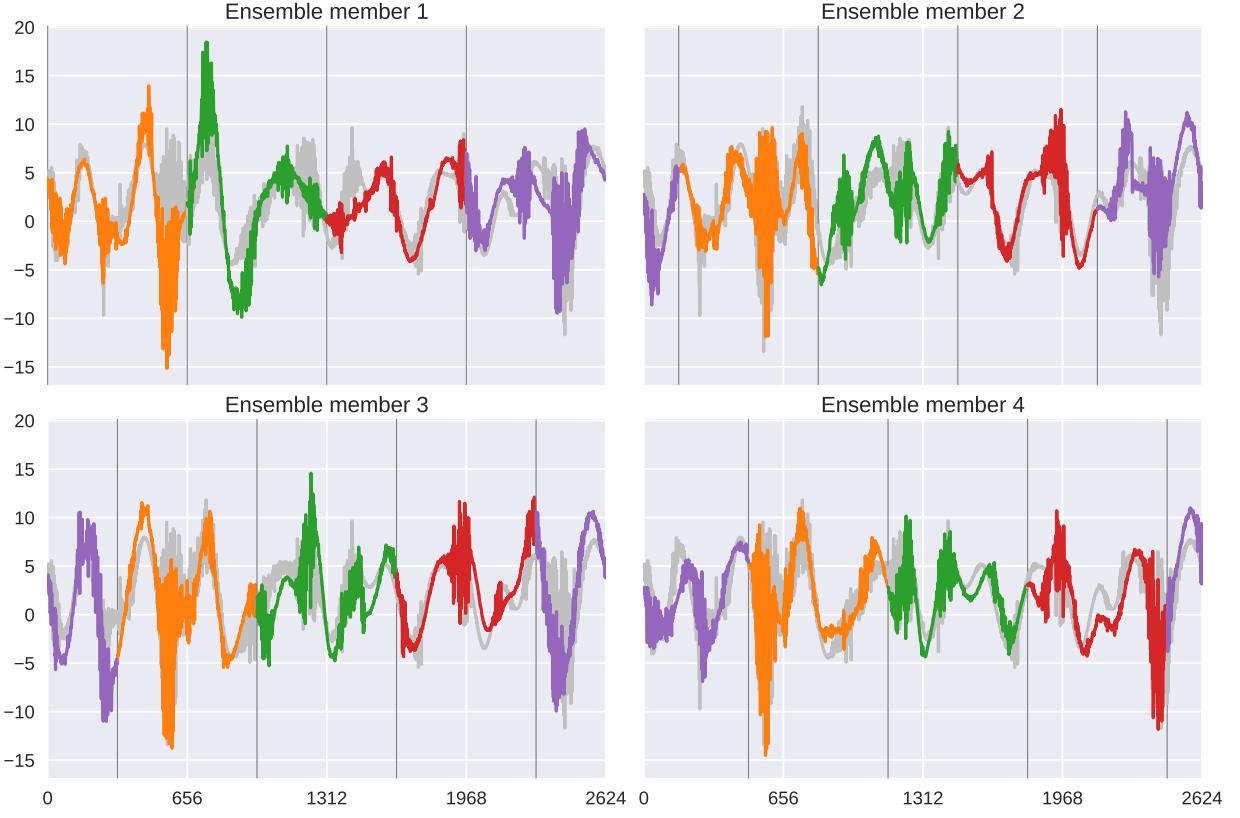


Figure 5: Example of four ensemble members generated by alg. 1. The grey line is the forecast from which the patched analogs are derived.

Since the covariance  $\Sigma$  is not known a priori, it can be approximated by encoding many samples from the training set into the latent space and computing the mean vector  $\hat{\mu}$ , the covariance matrix,  $\hat{\Sigma}$ , and the Cholesky factor such that  $\hat{\Sigma} \approx \hat{\mathbf{L}}\hat{\mathbf{L}}^\top$ . Finally, we can replace line 5 of alg. 1 with

$$\mathbf{x}^{(j,i)} \leftarrow d(\mathbf{z}_\mu + r_z \hat{\mathbf{L}}\epsilon^{(j,i)}) \quad (17)$$

to add regular perturbations even when using a general autoencoder. Outside the costs of training, this adds  $N_e$  matrix-vector multiplications, each of size `latentDim`, at every assimilation cycle as well as the initial costs of computing  $\hat{\Sigma}$ , and its Cholesky factorization. However, this will not significantly affect the computational cost since `latentDim` is likely sufficiently small and these extra operations will not significantly inflate the leading order costs.

### 3.3. Analysis: Impact of Patching on Covariance Matrix

There are two different effects at play in the structure of the background covariance matrix in p-cAnEnOI: sampling errors and the effect of patching. To isolate the effect of patching we develop in this section a patched stochastic process model, and then derive an explicit formula for its covariance function. Consider a stochastic process  $Z(s)$  with finite second-order moments. Without loss of generality we assume the mean is zero, and for simplicity of exposition we take  $s \in \mathbb{R}$ ; extensions to higher spatial dimension are discussed at the end of the section. Denote the covariance function of  $Z$  by

$$\text{Cov}[Z(s), Z(t)] = C(s, t). \quad (18)$$

Next construct a ‘patched’ stochastic process  $Z_p(s)$  as follows. Without loss of generality, let the patch width be 1. A single sample of  $Z_p$  is constructed by the following two-step procedure.

1. Draw a shift  $s_*$  from the uniform distribution on  $[0, 1)$  and define intervals  $I_k = [s_* + k - 1, s_* + k]$  for  $k = -\infty, \dots, \infty$ .
2. For each  $k$  draw an independent sample of  $Z(s)$  and then let  $Z_p(s) = Z(s)$  for  $s \in I_k$ .

In practice when constructing the ensemble members the patch boundaries are not drawn randomly; however, if a single ensemble member is drawn at random, the patch boundaries for that ensemble member can be considered as random variables.

To find the covariance function for the patched process  $Z_p$  we will first find the covariance conditioned on the shift  $s_*$ , and then average over the uniform distribution on  $s_*$ . If  $s$  and  $t$  are not in the same interval  $I_k$  then the conditional covariance  $\text{Cov}[Z_p(s), Z_p(t)]$  is zero. If they are in the same interval, then the conditional covariance is the same as for  $Z$ , i.e.  $C(s, t)$ . So the unconditional covariance function of  $Z_p$  is  $C(s, t)$  multiplied by the probability that  $s$  and  $t$  are both in the same interval, then averaged over the distribution of  $s_*$ .

Let  $t = s + d$ . The probability that  $s$  and  $t$  are both in  $I_k$  is zero if  $|d| \geq 1$ . If  $|d| < 1$ , the probability that  $s$  and  $t$  are both in  $I_k$  is 1 minus the probability that  $s_*$  is in between them. So

$$P[s, t \in I_k] = 1 - \min\{1, |s - t|\}.$$

The covariance function for the patched process is simply

$$\text{Cov}[Z_p(s), Z_p(t)] = C(s, t) [1 - \min\{1, |s - t|\}]. \quad (19)$$

In higher dimensions  $\mathbf{s} \in \mathbb{R}^d$  the form of the function changes slightly to reflect the general principle that the covariance function of the patched process is the covariance function of the original process multiplied by one minus the probability that there is at least one patch boundary between  $\mathbf{s}$  and  $\mathbf{t}$ .

This is equivalent to using a distance-based localization (called ‘tapering’ in the statistics literature) where the localization function is a ‘tent’ function (the shape of the localization function changes in higher dimensions). In order for the covariance function of the original and patched processes to remain similar, it is necessary to choose a patch size that is larger than the correlation length scale of the true process. How much larger is, of course, situation dependent. It is worth noting that the covariance matrix estimated from a patched ensemble can still exhibit spurious correlations at long ranges because of sampling errors, so localization methods still need to be used in patched cAnEnOI. In a future work we plan to explore overlapping (or layering) the patches with smooth transitions across the finite overlap, rather than a complete independence across patch boundaries.

#### 4. Numerical Results

We test our method on the 1D toy model described in section 2.4, a multiscale modification introduced in [10] to the Lorenz ’96 model of [11], and recently used as a data assimilation test model in [1, 28]. First, we generated a library of snapshots by numerically solving eq. (11) via `ode45` in Matlab until final time  $t = 200$  where each snapshot covered the full spatial domain. Since 0.2 time units is about one day of atmospheric dynamics according to [11], the simulation is run almost to 3 years. This library contains 270,576 samples, 25% of which were reserved for validation. After separating this library into a training set and a validation set, every element in the library was partitioned into smaller patches as our goal was to test the patch version of cAnEnOI. The patch sizes ranged from  $1/32$  to  $1/4$  of the original domain, and thus we created in total five libraries of the same data but with different data sizes. We elaborate on the details of training the VAEs and using them in p-cAnEnOI in section 4.1 and show the results of using general autoencoders in section 4.2.

#### 4.1. *p-cAnEnOI* with VAE

A main component of our method outlined in section 3 is the use of a machine learning model as a generative tool to create artificial ensemble members. In this section, we share the architectures of the specific machine learning models we used, how these models were trained, and the performance of these models with respect to standard machine learning evaluation metrics. Ultimately, we are most interested in how the artificial ensemble members impact the performance of the data assimilation task and we discuss the relationship between the quality of the machine learning models and the accuracy of the data assimilation in section 4.1.2. Recently, there have been many approaches of utilizing machine learning for data assimilation, some of which include [29–32]. We investigated autoencoders and variational autoencoders (see [27]) and trained several models using the library of patches of snapshots of the 1D toy model discussed in section 2.4.

##### 4.1.1. VAE Training

We used the variational autoencoder structure used in [1], Encoder 1,Decoder 1, for patches with sizes 1/4, 1/8, and 1/16 of the original variable. Encoder 1 reduces the dimension of the data to 1/4 of the input size with two max pooling layers (each pooling layer halves the dimension), and the fully connected layer reduces it further to two outputs that are each approximately 1/6 of the input size. These two outputs  $z_\mu$  and  $z_{\log(\Sigma)}$  are then combined via eq. (9), completing the encoding into the latent space. Finally, the original dimensionality is recovered by Decoder 1. We used Encoder 2 for the 1/32 patches because each sample was not divisible by 4 and therefore could not undergo the quarter reduction of the dimensions with the two max pooling layers. Therefore, Encoder 2 only has one max pooling layer and uses the fully connected layer to reduce the dimension to approximately 1/6 of the input size. Similarly, Decoder 2 only has one transposed convolutional layer with stride (2, 1) as opposed to the two layers in Decoder 1.

The VAEs were trained using the loss function as defined in section 2.2. This loss term includes the reconstruction error and the KL-divergence of the training set. Due to risk of overfitting to the training data, a validation set is used to compute an unbiased measure of the quality of a model while being trained. We examine only the reconstruction error portion of the loss on the validation set. Suppose that there are  $N$  elements of a dataset where each element is of `dszie` dimensions. Then, the reconstruction error (RE) is defined via eq. (20),

$$\text{RE} = \left( \frac{1}{N} \sum_{i=1}^N \frac{1}{\text{dszie}} \sum_{j=1}^{\text{dszie}} (\mathbf{x}_i^{(j)} - d(e(\mathbf{x}_i^{(j)})))^2 \right)^{1/2}, \quad (20)$$

where  $e : \mathbb{R}^{\text{dszie}} \rightarrow \mathbb{R}^{\text{latentDim}}$  and  $d : \mathbb{R}^{\text{latentDim}} \rightarrow \mathbb{R}^{\text{dszie}}$  are some encoder and decoder, and  $\mathbf{x}_i^{(j)}$  is the  $j^{th}$  element of the  $i^{th}$  sample. Recall that we used one library of snapshots to generate four more libraries of the same snapshots partitioned into the four different patch sizes. We evaluate the quality of our VAEs via eq. (20), which averages over a dataset with  $N$  elements of dimension `dszie` and applies the square root at the very end to ensure uniformity across the four different libraries. Note that  $N \times \text{dszie}$  for the validation set for all five libraries remains at the same value, 44,374,464.

In fig. 6, we show the averaged RE for the validation set every 20 epochs until the 1,000<sup>th</sup> for the patch sizes 1/4, 1/8, 1/16, and 1/32 while training the VAE architectures [Encoder 1, Decoder 1] and [Encoder 2,Decoder 2]. We can see that the VAEs for all patch sizes experience a significant drop in the reconstruction error in the first 200 epochs. The drop happens sooner at smaller patch sizes, to the extent that at the smallest patch size it happens before the first point on the plot, at 20 epochs. If we exclude the smallest patch size, then we observe that the smaller patch sizes are able to achieve smaller reconstruction errors. Furthermore, the VAE reconstruction errors for the 1/4 and 1/32 patch sizes seem to have reached their minima within the first 500 epochs, whereas the errors for the other two patch sizes could potentially decrease further.

Next, fig. 7 shows how the quality of the 1/4 patch VAE improves with more training for a single sample. The original sample is shown in blue, and the other colors show the reconstruction of that sample with the VAE after 20, 160, 300, and 440 epochs of training. Corresponding to the large drop in the reconstruction error in the 1/4 patch between the 140<sup>th</sup> and the 180<sup>th</sup> epoch in fig. 6, the top right panel shows the

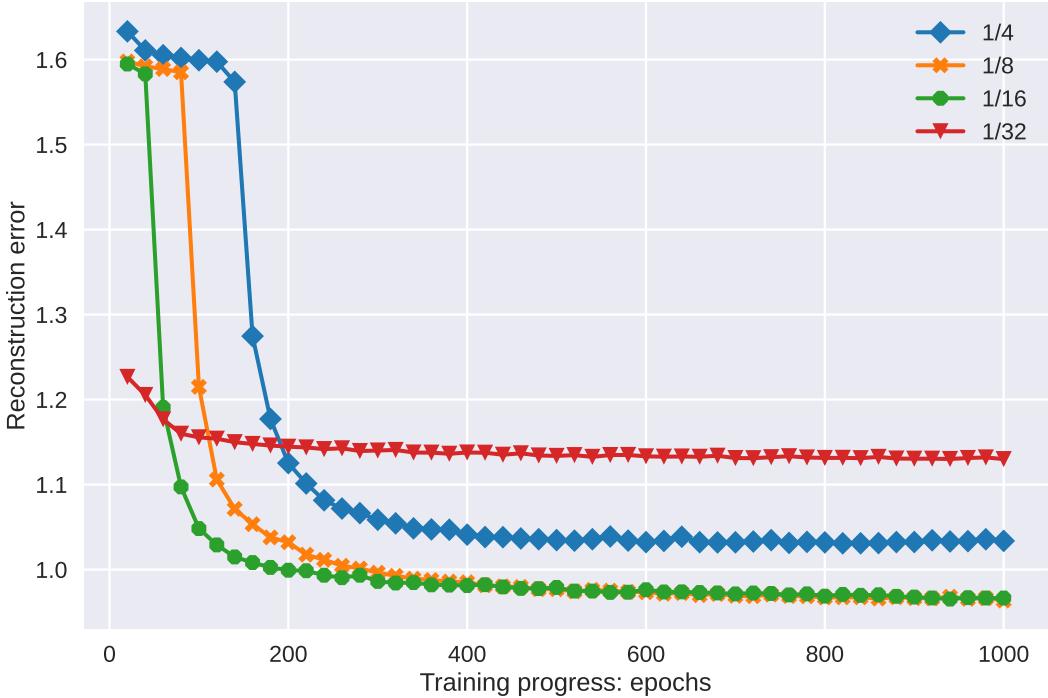


Figure 6: Averaged reconstruction error computed via eq. (20) on the validation set shown every 20 epochs during training.

most visibly recognizable improvement at epoch 160, while the bottom two plots certainly show continued progress. Combining fig. 6 and fig. 7 we conclude that the VAE rapidly learns to reproduce the large scales, and then after sufficient training it begins to learn the small scales; the rapid drop in reconstruction error evident in fig. 6 is evidently associated with learning the small scales.

Finally, fig. 8 shows the relative errors of the spectra for VAEs for the four patch sizes for 500 epochs on a sample from the validation set. There is a sharp increase in the error exactly at the scale boundary, which suggests that the large scale dynamics are well emulated by the all four VAEs. The smaller patch VAEs actually show lower errors in the large scales, but the errors are indistinguishable for the large wave numbers. In conclusion, the VAEs for the larger patch sizes incur more loss in reconstruction in comparison to the VAEs for smaller patch sizes when trained for the same number of epochs.

#### 4.1.2. Data Assimilation Results

The reference simulation is initialized with standard normal noise and run with MATLAB's `ode45` (which has adaptive time-stepping) until  $t = 9$ , by which the statistical equilibrium has been reached. The data assimilation starts from the state at  $t = 9$ , and the reference is simulated for 73 additional time units with outputted at a time interval of 0.2, which corresponds to one 'day' according to [11]. The observations are taken at every 4th spatial node and assimilated at every 0.2 time units, and the observation errors are sampled from  $N(0, 1/2)$ . There are a total of  $73 \div 0.2 = 365$  assimilation cycles, of which the first 73 are discarded as the burn-in period of the data assimilation. The root mean squared error (RMSE, see eq. (2)) between the reference simulation and the analysis mean for the latter 292 assimilation cycles quantify the performance of p-cAnEnOI for the specific set of parameters used. The relevant tunable parameters include:  $r_z$ , the amplitude of noise added in the latent space when generating artificial ensemble members via eq. (10);

$f$ , the forecast spread of the generated ensemble;  $L$ , the localization radius to zero out spuriously high correlations across far distances. Each combination of these three parameters undergoes 8 trials to account for statistical variability and we summarize the performance with the mean and standard error across these 8 trials. Note that several experiments from [1] were continued to 1,000 assimilation cycles with no change in the performance, indicating that 8 independent trials with 365 cycles each is sufficient to assess the performance of the different methods.

These settings are set identical to the experiments in [1] so the patched constructed analog method can be compared to those. The data assimilation performance was measured using RMSE eq. (2), which measures the error between reference simulation and the ensemble mean after assimilating observations. It was found in [1] that a 200 sized ensemble showed little improvement over an ensemble size of 100, and therefore  $N_e = 100$  was used for all methods, ESRF, EnOI, AnEnOI, and cAnEnOI. EnOI provides a lower bound as expected, and AnEnOI and cAnEnOI result in lower RMSE than EnOI. The mean RMSE for cAnEnOI is actually lower than the mean RMSE from ESRF but the 8 independent trials were used to show that this difference is not statistically significant. We summarize the relevant results from [1] in table 1.

	ESRF	EnOI	AnEnOI	cAnEnOI
DA RMSE	1.35	2.27	2.01	1.30

Table 1: Summary of various analog data assimilation methods from [1] with optimal data assimilation parameters and  $r_z$  for cAnEnOI.

Our new method, p-cAnEnOI, uses a generative machine learning model to construct patched-analog

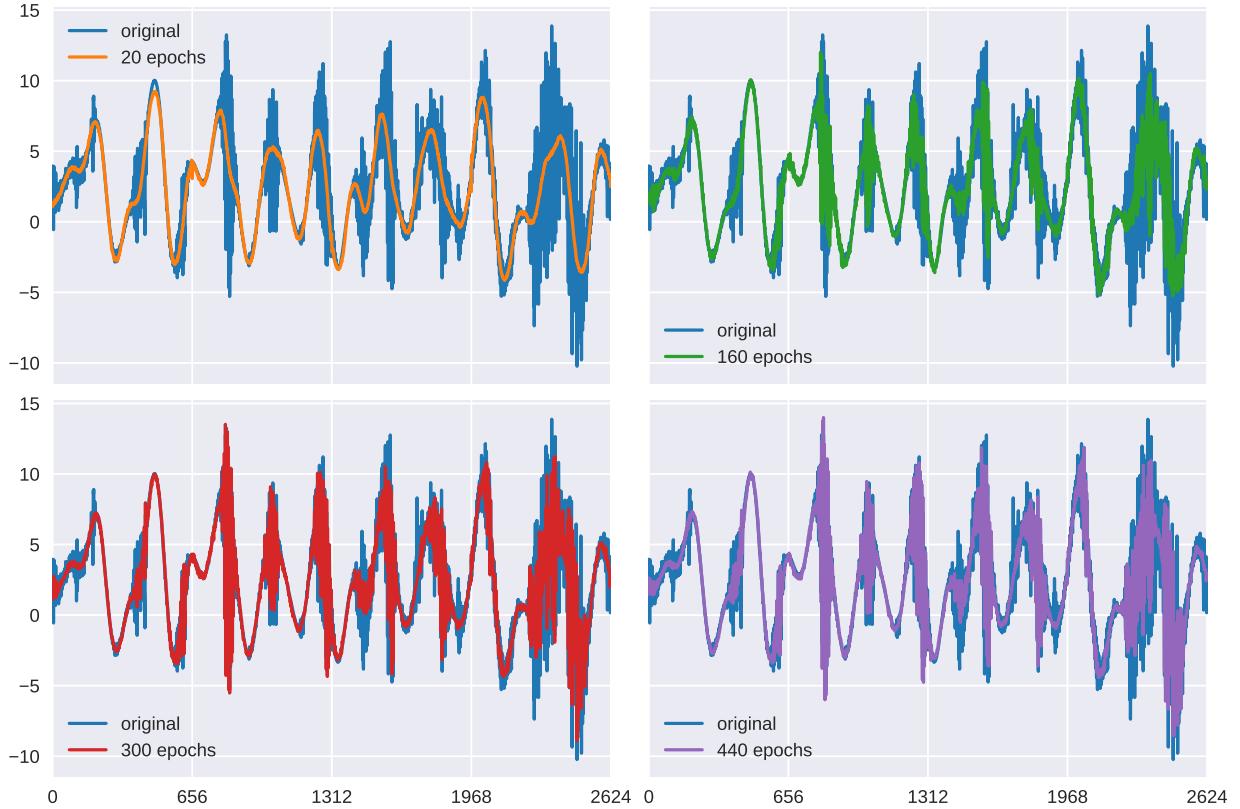


Figure 7: A sample in the validation set is shown with its reconstruction after 20, 160, 300, and 440 epochs of training. The patch size is 1/4 of the domain size.

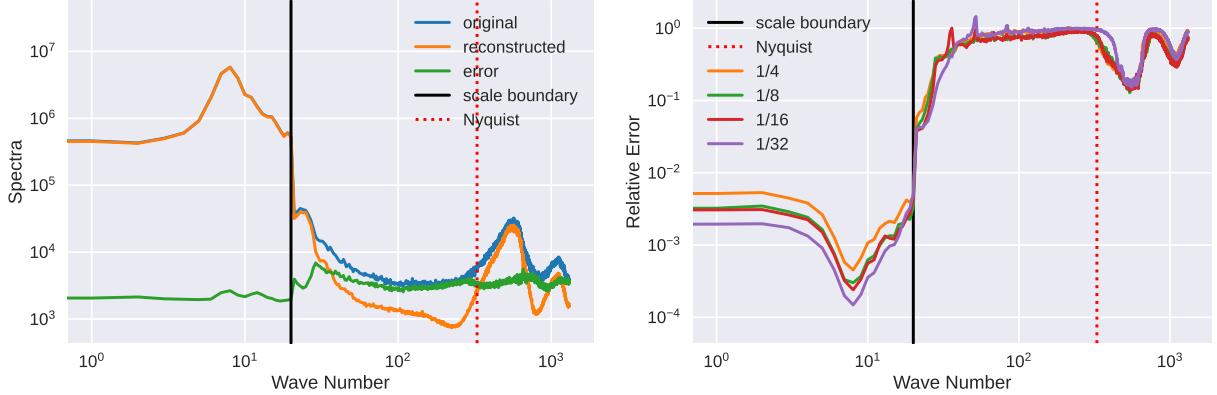


Figure 8: Left: The Fourier spectra of a sample from the validation set, its reconstruction using a full domain VAE, and the absolute error. Right: The relative error in the Fourier spectra of the same sample VAEs with various patch sizes. The solid black line indicates the cutoff scale between the 41 large-scale modes and the remaining small-scale modes, while the dashed red line indicates the Nyquist wavenumber of the observing system.

ensemble members for use within EnOI. We searched through the parameter spaces of  $r_z$ ,  $f$ , and  $L$  for p-cAnEnOI with VAEs for each of the patch sizes 1/4, 1/8, 1/16, and 1/32, with 8 trials for each combination. The VAEs used the architecture [Encoder 1, Decoder 3] for patch sizes 1/4, 1/8, 1/16 and [Encoder 2, Decoder 4] for patch size 1/32.

Figure 9 shows the results of the parameter sweep for the p-cAnEnOI that uses patch size 1/4. The color of each point represents the average RMSE value over 8 trials for a specific combination of the three parameters  $r_z$ ,  $f$ , and  $L$ , and we see that the parameters significantly influence the results. When  $r_z = 0.10$ ,  $f = 0.8$ , and  $L = 16$ , the optimal RMSE of 0.89 is achieved for the p-cAnEnOI scheme with 100 patched-analog ensemble members, which is significantly lower than that achieved by the global cAnEnOI and the ESRF. This improvement could suggest that the localization associated with patching (section 3.3) is more effective than the traditional localization described by eq. (4).

Table 2 shows the optimal set of parameters for each of the patch sizes used, the corresponding mean and standard errors of the RMSE, and the mean reconstruction error of the validation set. If we exclude the column with the global cAnEnOI, there is a clear trend that shows larger patch sizes achieving the best optimal RMSE. While the optimal  $f$  value stays within [0.70, 0.90] for all variants, the localization radius  $L$  and amplitude of noise in latent space,  $r_z$ , vary quite a bit. The global cAnEnOI from [1] bucks this trend, but it might be that the VAE used in the global patch has not yet been trained to its optimal performance level, because it is significantly more expensive to train than the VAEs used for the smaller patches.

The variation in optimal localization radius may be explained by reference to the results in section 3.3. Patched analogs have an effect similar to localizing the covariances. At large patch sizes the amount of localization needed is comparable to an un-patched method; at small patch sizes the effect of patching already significantly reduces short-range correlations, but some localization is still needed to remove spurious correlations at long ranges.

The larger patch sizes (1/4 and 1/8) benefit from smaller  $r_z$ , whereas the smaller two patch sizes (1/16 and 1/32) prefer larger  $r_z$ 's. This sensitivity to  $r_z$  might be explained as follows. Recall that VAEs are designed so that the encoder is a mapping from the space of instances of input variables to the latent space, which is forced towards the multivariate standard normal distribution,  $N(\mathbf{0}, \mathbf{I})$ . It is possible that the VAEs for the two larger patch sizes encode into latent spaces having enough dimensions to allow for preserving deeper complexities than the VAEs for the smaller patch sizes. If so, even a small amplitude noise added in those more complex latent spaces could generate artificial ensemble members that exhibit enough diversity, whereas a larger amplitude is necessary to generate distinct enough ensemble members for the VAEs with smaller latent space dimensions.

The above explanation is inconsistent with the large value of  $r_z$  for the global cAnEnOI method. The

discrepancy can be explained as follows. It was noted in [1] that in the  $r_z \rightarrow 0$  limit of cAnEnOI the analog ensemble covariance matrix is rank-deficient, with rank limited by the dimension of the latent space, whereas at larger  $r_z$  the nonlinearity of the decoder comes into play and enables the analog ensemble to span a subspace of higher rank. Specifically, in the limit  $r_z \rightarrow 0$  the range of the cAnEnOI covariance must be a subspace of the Jacobian of the decoder. This limitation does not apply to the p-cAnEnOI method, so the range of the p-cAnEnOI covariance can access a larger subspace, even in the limit  $r_z \rightarrow 0$ .

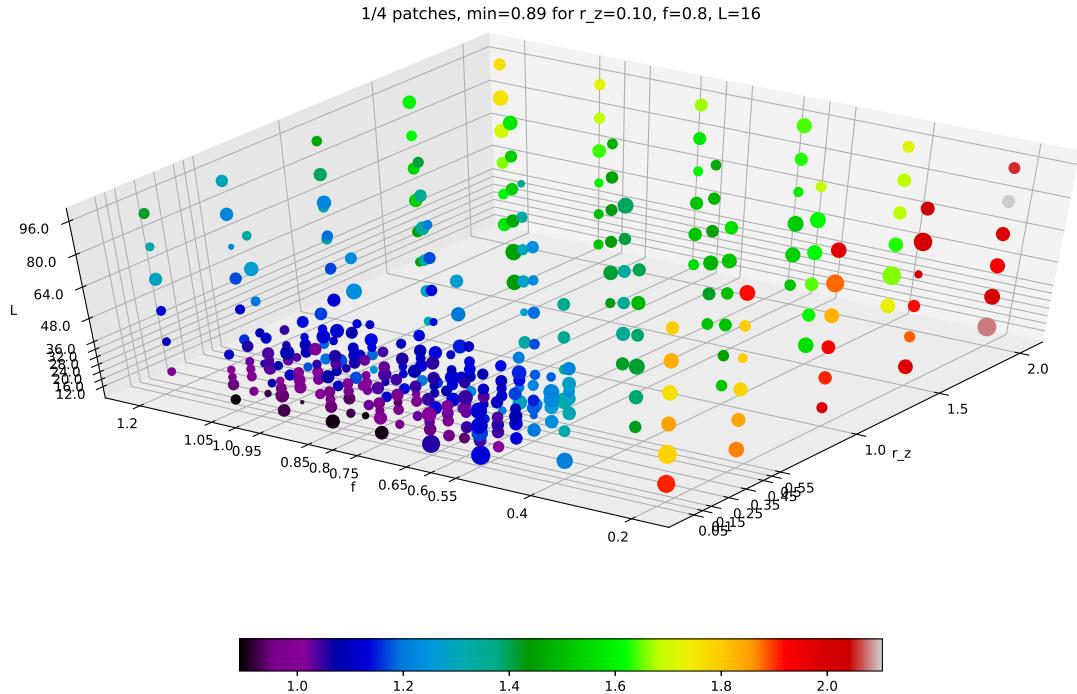


Figure 9: The color and the size of the points denote the computed mean and the standard error of the RMSE across 8 trials and their locations mark the combination of the three relevant parameters:  $r_z$ , the amplitude of noise added in latent space to generate artificial ensemble members;  $f$ , the forecast spread; and  $L$ , the localization radius. Larger points imply larger RMSE, and smaller points, smaller RMSE. This plot shows the data for the 1/4 patch size p-cAnEnOI.

To explore the relationship between the accuracy of the VAEs as measured by the reconstruction error eq. (20) and the performance of the corresponding p-cAnEnOI, we return to the set of VAEs discussed in section 4.1.1 which used [Encoder 1, Decoder 1] for patch sizes 1/4, 1/8, 1/16 and [Encoder 2, Decoder 2] for patch size 1/32. The VAEs were saved every 20 epochs during training to track the progress of the quality of the VAEs with respect to the reconstruction error (see eq. (20)) over the validation set. Using the set of optimal parameters  $(L, f, r_z)$  found using Decoders 3 and 4 (and shown in table 2) for each patch size, we computed the data assimilation mean RMSE over 8 trials for each of the VAEs saved throughout the first 500 epochs of training. Figure 10 shows the trend between the reconstruction error eq. (20) and the mean RMSE eq. (2) for each patch size, and the optimal RMSE out of the 50 saved VAEs for each patch size are reported in table 2.

As the reconstruction error decreases with more training, the quality of the data assimilation also improves as is shown by lower mean RMSE values. In fact, there is little distinction between p-cAnEnOI with patch size 1/4, 1/8, and 1/16 for VAEs that yielded reconstruction error between 1.2 to 1.5, other than that the 1/8 patch size has a slightly steeper slope. The 1/16 variant splits off from the other three and a higher mean RMSE is computed where the reconstruction error ranges approximately from 0.95 to 1.2. The 1/8 variant produces mean RMSE close to 1 even when the reconstruction error continues to decrease from about 1.3 to 0.95, and the 1/4 variant produces mean RMSE that dips slightly below 1 despite the

Patch size	1	1/4	1/8	1/16	1/32
$L$	40	16	12	24	36
$f$	0.70	0.80	0.70	0.80	0.90
$r_z$	0.70	0.10	0.05	0.25	0.45
DA RMSE: Decoder 3, 4	n/a	$0.89 \pm 0.01$	$0.95 \pm 0.01$	$1.13 \pm 0.02$	$1.18 \pm 0.03$
DA RMSE: Decoder 1, 2	$1.30 \pm 0.05$	$0.88 \pm 0.01$	$0.93 \pm 0.02$	$1.17 \pm 0.02$	$1.39 \pm 0.03$
RE: Decoder 3, 4	n/a	1.03	0.97	0.96	0.78
RE: Decoder 1, 2	1.26	1.03	0.99	0.97	1.13

Table 2: Summary of optimal parameters for p-cAnEnOI for patch sizes 1/4, 1/8, 1/16, and 1/32, as well as the global cAnEnOI, and the corresponding mean and standard error of the data assimilation RMSE and reconstruction error (RE).

reconstruction error stagnating around 1.3. The reconstruction error for 1/32 cannot improve after 1.12 and the RMSE similarly is restricted to range from 1.35 to 1.5, which indicates that the small latent space of this VAE cannot represent the full complexity of the dynamics of this model. Overall, there is a clear linear relationship between the reconstruction error and the mean RMSE of the data assimilation performance. The slight disparities between these p-cAnEnOI variants may be attributed to the varying patch sizes and the second loss term, the KL-divergence. For example, it is possible that the 1/16 p-cAnEnOI yields a higher RMSE than p-cAnEnOI with patch sizes 1/4 and 1/8 despite having the same reconstruction error of 1.03 because that latent space of the 1/16 patch vae is less regular in comparison to the other two. Or, it could be that that 1/16 patch VAE fails to capture the large scale motions accurately enough since the smallest large scale wave has length approximately 1/8 of the domain.

In table 2, we see that the two sets of VAEs (with different decoders) for each patch size 1/4, 1/8, and 1/16 produce RMSE and RE that are very similar, but the two VAEs for patch size 1/32 are quite different from each other. The mean RE for the VAE with Decoder 4 for patch size 1/32 yields the smallest reconstruction error out of all 9 VAEs, but still yields the second worst RMSE. This strongly suggests that too small of a patch size loses too much information and cannot act as a good generative model for constructing analogs.

Most importantly, we observe that the p-cAnEnOI schemes with patch sizes (strictly) larger than 1/32 and the global cAnEnOI outperform the optimal ESRF although ESRF could probably do better with more sophisticated localization and inflation. While we did not measure the exact computation cost for these variants, implementation of ESRF is far more expensive in the forecasting step, which often accounts for the leading order computational expense.

#### 4.2. *p-cAnEnOI with a general autoencoder*

A main component of VAEs is the effort to regularize the latent space by encouraging it to look like the standard normal distribution, and is done by eq. (9) and the inclusion of the KL-divergence in the loss term as is discussed in section 2.2.2. To provide a basic comparison, we used Encoder 1 that only outputs one variable of size `latentDim` at the last layer, and Decoder 3 exactly as is to form a simple autoencoder for patch size 1/8 and trained it only with the reconstruction error as the loss term. After a week of training, the reconstruction error computed via eq. (20) using the autoencoder on the validation set yielded 0.92, a significantly lower value than the two corresponding VAEs (a week of training with Decoder 3 and 1000 epochs with Decoder 1). The computed RMSEs had mean 0.96 and standard error 0.02, which shows that a general encoder can produce similar results as VAE within p-cAnEnOI with the additional step described in section 3.2.

## 5. Conclusion

Using analogs and constructed analogs in EnOI (see [1] for details) is a recent development in the growing field of using machine learning techniques for data assimilation. AnEnOI and cAnEnOI alleviate the

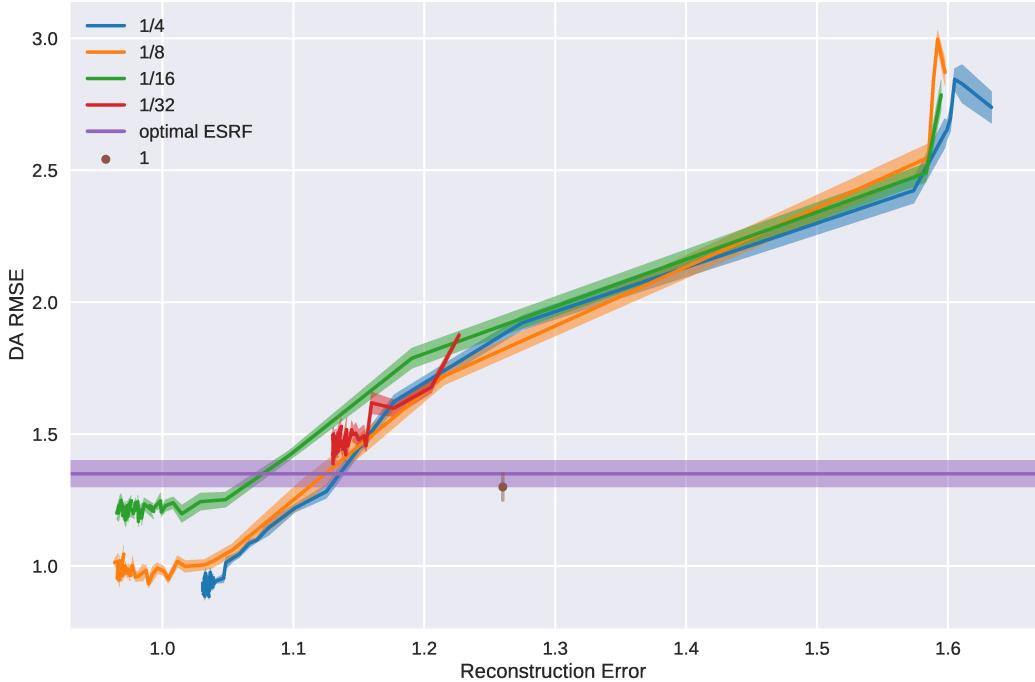


Figure 10: Averaged reconstruction error of VAEs in training over validation data set compared against the resulting mean and standard error of DA RMSE. The shaded regions show the range of the standard error around the mean values.

computational challenge of having to forecast large ensembles in EnKF and its variants, but improves upon EnOI by inserting time dependency to the ensemble such that the background covariance depends on the current forecast. A limitation of cAnEnOI from [1] is the need to train a generative machine learning model to construct samples of the entire model state, which can include billions of variables for geophysical applications. We have introduced p-cAnEnOI which uses VAEs and general autoencoders to generate patches of analogs that form whole analogs when assembled. This new scheme addresses the scalability of implementing cAnEnOI in geophysical applications in several ways: 1) Dividing the spatial domain into smaller subsets reduces the complexity of the machine learning model to be trained and encourages robust training with fewer parameters learned with more training samples available; 2) Multiple analog patches can be simultaneously generated in a parallel process then combined together, allowing for speed-ups.

Our numerical experiments were designed to test if and how the patch size affects the data assimilation performance, as well as how the accuracy of the generative model affects the data assimilation performance. First, we trained VAEs that learned patches of instances of the toy model (defined in eq. (11)) with sizes ranging from 1/4, 1/8, 1/16, and 1/32 of the spatial domain. Then, we solved for the optimal set of three tunable parameters of the data assimilation system ( $f$ , forecast spread;  $r_z$ , the amplitude of noise added in latent space; and  $L$ , the localization radius) by running 8 trials of p-cAnEnOI for ranges of those parameters for each of the four patch sizes. We used the root mean squared error between the reference simulation and the analysis mean as defined in eq. (2) to measure the quality of data assimilation, and found that using larger patches led to better results. While the optimal forecast spread was similar across the different patch sizes, the localization radius and  $r_z$  varied quite a bit. Notably, the larger patches required smaller localization radii unlike smaller patches that naturally perform some localization. However, patch sizes that

are large enough to capture the dominant large-scale structures in the data have similar optimal localization radii and small  $r_z$ , and variations in these parameters appear only for patch sizes that are sub-optimal.

All of the p-cAnEnOI variants produced better data assimilation RMSE than cAnEnOI, AnEnOI, EnOI, and ESRF. Given that the large patches produce better results, it may be surprising that p-cAnEnOI performed better than cAnEnOI. Since we found that the normalized reconstruction errors (see eq. (20)) of VAEs for the small patch sizes were lower, we sought to find a relationship between the reconstruction error and the data assimilation RMSE. In this second experiment, we saved the VAEs at various stages during training and tested them within p-cAnEnOI with the optimal parameters found from the previous experiment. As expected, more training produced better reconstruction errors, which then led to better data assimilation results. However, when comparing p-cAnEnOI with similar reconstruction errors, we found that the variants that use larger patches still performed marginally better. The lower than expected performance of global cAnEnOI from [1] might thus be due to its use of an incompletely trained VAE. Training the global VAE to its optimal performance is significantly more expensive than training a VAE for local patches.

Lastly, we tested the use of general (non-variational) autoencoders in p-cAnEnOI. We implemented the method outlined in section 3.2 for the 1/8 patch with the parameters found for that patch size in the first experiment. The data assimilation RMSE averaged over 8 trials for this general autoencoder version produced results not statistically different from the VAE version. This is a promising result, since it suggests that autoencoders with simpler architecture than VAEs, and which are therefore easier to train, might perform as well as VAEs for this task.

In conclusion, constructing patched analogs shows encouraging results that prompt future studies of applying cAnEnOI and p-cAnEnOI to more realistic models. Both the patching scheme and the use of general autoencoders reduce the cost of training the autoencoder compared to global cAnEnOI. In future work, we will address what patching scheme to use in non-periodic domains with boundaries. We will also investigate overlapping the patches with a smooth transition region rather than a sharp patch boundary in order to mitigate the across-patch discontinuities. Since convolutional neural networks can be used for categorization, we believe they may be able to learn diverse sets of samples that can encompass the dynamics near the boundaries as well as away from the boundaries. As there is an abundance of historical datasets of geophysical systems, our study of utilizing these in an efficient way to improve data assimilation methods is meaningful.

## Acknowledgements

Funding: This work was funded by the US National Science Foundation under grant no. DMS-1821074. This work used the Extreme Science and Engineering Discovery Environment (XSEDE; Towns et al., 2014) Bridges (Nystrom et al., 2015) at the Pittsburgh Super Computing Center through allocation TG-DMS190025.

- [1] I. Grooms, Analog ensemble data assimilation and a method for constructing analogs with variational autoencoders, *Quarterly Journal of the Royal Meteorological Society* 147 (2021) 139–149.
- [2] G. Evensen, The Ensemble Kalman Filter: Theoretical formulation and practical implementation, *Ocean Dynamics* 53 (2003) 343–367.
- [3] W. Lahoz, B. Khattatov, R. Menard (Eds.), *Data Assimilation*, Springer Berlin Heidelberg, 2010. URL: <https://doi.org/10.1007/978-3-540-74703-1>. doi:10.1007/978-3-540-74703-1.
- [4] P. R. Oke, J. S. Allen, R. N. Miller, G. D. Egbert, P. M. Kosro, Assimilation of surface velocity data into a primitive equation coastal ocean model, *Journal of Geophysical Research: Oceans* 107 (2002) 1–25.
- [5] P. R. Oke, G. B. Brassington, D. A. Griffin, A. Schiller, Ocean data assimilation: A case for ensemble optimal interpolation, *Australian Meteorological and Oceanographic Journal* 59 (2010) 67–76.
- [6] L. Delle Monache, F. A. Eckel, D. L. Rife, B. Nagarajan, K. Searight, Probabilistic weather prediction with an analog ensemble, *Monthly Weather Review* 141 (2013) 3498–3516.
- [7] F. A. Eckel, L. Delle Monache, A hybrid nwp-analog ensemble, *Monthly Weather Review* 144 (2016) 897–911.
- [8] Z. Zhao, D. Giannakis, Analog forecasting with dynamics-adapted kernels, *Nonlinearity* 29 (2016) 2888.
- [9] D. P. Kingma, M. Welling, Auto-encoding variational bayes, *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings* (2014) 1–14.
- [10] I. Grooms, Y. Lee, A framework for variational data assimilation with superparameterization, *Nonlinear Processes in Geophysics* 22 (2015) 601–611.
- [11] E. N. Lorenz, Predictability: A problem partly solved, in: *Proc. Seminar on predictability*, volume 1, 1996.

- [12] B. K. Johnson, F. O. Bryan, S. A. Grodsky, J. A. Carton, Climatological annual cycle of the salinity budgets of the subtropical maxima, *Journal of Physical Oceanography* 46 (2016) 2981–2994.
- [13] G. Evensen, Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *Journal of Geophysical Research* 99 (1994).
- [14] P. L. Houtekamer, H. L. Mitchell, Data assimilation using an ensemble kalman filter technique, *Monthly Weather Review* 126 (1998) 796–811.
- [15] G. Burgers, P. Jan van Leeuwen, G. Evensen, Analysis scheme in the ensemble kalman filter, *Monthly weather review* 126 (1998) 1719–1724.
- [16] M. K. Tippett, J. L. Anderson, C. H. Bishop, T. M. Hamill, J. S. Whitaker, Ensemble square root filters, *Monthly Weather Review* 131 (2003) 1485–1490.
- [17] J. S. Whitaker, T. M. Hamill, Ensemble data assimilation without perturbed observations, *Monthly weather review* 130 (2002) 1913–1924.
- [18] M. El Gharamti, K. Raeder, J. Anderson, X. Wang, Comparing adaptive prior and posterior inflation for ensemble filters using an atmospheric general circulation model, *Monthly Weather Review* 147 (2019) 2535–2553.
- [19] H. Van den Dool, J. Huang, Y. Fan, Performance and analysis of the constructed analogue method applied to us soil moisture over 1981–2001, *Journal of Geophysical Research: Atmospheres* 108 (2003).
- [20] E. N. Lorenz, Atmospheric predictability as revealed by naturally occurring analogues, *Journal of the Atmospheric sciences* 26 (1969) 636–646.
- [21] R. Lguensat, P. Tandeo, P. Ailliot, M. Pulido, R. Fablet, The analog data assimilation, *Monthly Weather Review* 145 (2017) 4093–4107.
- [22] J. Ballé, V. Laparra, E. P. Simoncelli, End-to-end optimized image compression, *arXiv preprint arXiv:1611.01704* (2016).
- [23] L. Theis, W. Shi, A. Cunningham, F. Huszár, Lossy image compression with compressive autoencoders, *arXiv preprint arXiv:1703.00395* (2017).
- [24] J. An, S. Cho, Variational autoencoder based anomaly detection using reconstruction probability, *Special Lecture on IE* 2 (2015) 1–18.
- [25] K. Zeng, J. Yu, R. Wang, C. Li, D. Tao, Coupled deep autoencoder for single image super-resolution, *IEEE transactions on cybernetics* 47 (2015) 27–37.
- [26] K. G. Lore, A. Akintayo, S. Sarkar, Llnet: A deep autoencoder approach to natural low-light image enhancement, *Pattern Recognition* 61 (2017) 650–662.
- [27] D. P. Kingma, M. Welling, An introduction to variational autoencoders, *Foundations and Trends in Machine Learning* 12 (2019) 307–392.
- [28] G. Robinson, I. Grooms, A hybrid particle-ensemble kalman filter for problems with medium nonlinearity, *arXiv preprint arXiv:2006.04699* (2020).
- [29] R. S. Cintra, H. F. de Campos Velho, Data assimilation by artificial neural networks for an atmospheric general circulation model, *Advanced Applications for Artificial Neural Networks* 265 (2018).
- [30] T. Arcomano, I. Szunyogh, J. Pathak, A. Wikner, B. R. Hunt, E. Ott, A machine learning-based global atmospheric forecast model, *Geophysical Research Letters* 47 (2020) e2020GL087776.
- [31] A. Wikner, J. Pathak, B. Hunt, M. Girvan, T. Arcomano, I. Szunyogh, A. Pomerance, E. Ott, Combining machine learning with knowledge-based modeling for scalable forecasting and subgrid-scale closure of large, complex, spatiotemporal systems, *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30 (2020) 053111.
- [32] J. Brajard, A. Carrassi, M. Bocquet, L. Bertino, Combining data assimilation and machine learning to infer unresolved scale parametrization, *Philosophical Transactions of the Royal Society A* 379 (2021) 20200086.

## Appendix A. Model Architectures

### Encoder 1 (Used for global and 1/4, 1/8, 1/16 patches)

1. A convolutional layer with 3 filters of size (3,1) and an elu activation.
2. A convolutional layer with 9 filters of size (3,1) and an elu activation.
3. A convolutional layer with 27 filters of size (3,1) and an elu activation.
4. A max pooling layer with pool size (2,1).
5. A convolutional layer with 27 filters of size (3,1) and an elu activation.
6. A convolutional layer with 27 filters of size (3,1) and an elu activation.
7. A max pooling layer with pool size (2,1).
8. A convolutional layer with 27 filters of size (3,1) and an elu activation.
9. A convolutional layer with 27 filters of size (3,1) and an elu activation.

10. A max pooling layer with pool size (2,1).
11. A fully connected layer with two outputs  $z_\mu$  and  $z_{\log(\Sigma)}$ , each of size `latentDim`.

**Encoder 2 (Used for 1/32 patches)**

1. A convolutional layer with 3 filters of size (3,1) and an elu activation.
2. A convolutional layer with 9 filters of size (3,1) and an elu activation.
3. A convolutional layer with 27 filters of size (3,1) and an elu activation.
4. A max pooling layer with pool size (2,1).
5. A convolutional layer with 27 filters of size (3,1) and an elu activation.
6. A convolutional layer with 27 filters of size (3,1) and an elu activation.
7. A max pooling layer with pool size (2,1).
8. A fully connected layer with two outputs  $z_\mu$  and  $z_{\log(\Sigma)}$ , each of size `latentDim`.

**Decoder 1 (Used for global and 1/4, 1/8, 1/16 patches)**

1. A fully connected layer, which is then reshaped to 27 channels, followed by an elu activation.
2. A transposed convolutional layer with 27 filters of size (3,1), stride of (2,1), and an elu activation.
3. A convolutional layer with 27 filters of size (3,1), unit stride, and an elu activation.
4. A transposed convolutional layer with 9 filters of size (3,1), stride of (2,1), and an elu activation.
5. A convolutional layer with 9 filters of size (3,1), unit stride, and an elu activation.
6. A convolutional layer with 1 filter of size (3,1), unit stride, and an elu activation.

**Decoder 2 (Used for 1/32 patches)**

1. A fully connected layer, which is then reshaped to 27 channels, followed by an elu activation.
2. A transposed convolutional layer with 9 filters of size (3,1), stride of (2,1), and an elu activation.
3. A convolutional layer with 9 filters of size (3,1), unit stride, and an elu activation.
4. A convolutional layer with 1 filter of size (3,1), unit stride, and an elu activation.

**Decoder 3 (Used for global and 1/4, 1/8, 1/16 patches)**

1. A fully connected layer, which is then reshaped to 27 channels, followed by an elu activation.
2. A transposed convolutional layer with 27 filters of size (3,1), stride of (2,1), and an elu activation.
3. An elu activation.
4. A transposed convolutional layer with 9 filters of size (3,1), stride of (2,1), and an elu activation.
5. A transposed convolutional layer with 9 filters of size (3,1), unit stride, and an elu activation.
6. A transposed convolutional layer with 1 filter of size (3,1), unit stride, and an elu activation.

**Decoder 4 (Used for 1/32 patches)**

1. A fully connected layer, which is then reshaped to 27 channels, followed by an elu activation.
2. An elu activation layer.
3. A transposed convolutional layer with 9 filters of size (3,1), stride of (2,1), and an elu activation.
4. A transposed convolutional layer with 9 filters of size (3,1), unit stride, and an elu activation.
5. A transposed convolutional layer with 1 filter of size (3,1), unit stride, and an elu activation.