

Integrating recurrent neural networks with data assimilation for scalable data-driven state estimation

Stephen Gregory Penny¹, Timothy A Smith¹, Tse-Chun Chen², Jason Alexander Platt³, Hsin-Yi Lin⁴, Michael Goodliff⁵, and Henry Don Isaac Abarbanel⁶

¹Cooperative Institute for Research in Environmental Sciences at the University of Colorado Boulder

²University of Maryland, College Park

³University of California San Diego

⁴The Cooperative Institute for Research in Environmental Sciences (CIRES)

⁵University of Colorado

⁶UC San Diego

November 26, 2022

Abstract

Data assimilation (DA) is integrated with machine learning in order to perform entirely data-driven online state estimation. To achieve this, recurrent neural networks (RNNs) are implemented as surrogate models to replace key components of the DA cycle in numerical weather prediction (NWP), including the conventional numerical forecast model, the forecast error covariance matrix, and the tangent linear and adjoint models. It is shown how these RNNs can be initialized using DA methods to directly update the hidden/reservoir state with observations of the target system. The results indicate that these techniques can be applied to estimate the state of a system for the repeated initialization of short-term forecasts, even in the absence of a traditional numerical forecast model. Further, it is demonstrated how these integrated RNN-DA methods can scale to higher dimensions by applying domain localization and parallelization, providing a path for practical applications in NWP.

Integrating recurrent neural networks with data assimilation for scalable data-driven state estimation

S.G. Penny^{1,2}, T.A. Smith^{1,2}, T.-C. Chen^{1,2}, J. Platt³, H.-Y. Lin^{1,2}, M. Goodliff^{1,2}, H.D.I. Abarbanel^{3,4}

¹Cooperative Institute for Research in Environmental Sciences (CIRES) at the University of Colorado

Boulder

²National Oceanic and Atmospheric Administration (NOAA) Physical Sciences Laboratory (PSL)

³University of California San Diego (UCSD) Department of Physics

⁴Scripps Institution of Oceanography (SIO)

¹216 UCB, University of Colorado Boulder, Boulder, Colorado, 80309, USA

²325 Broadway, Boulder, Colorado, 80305, USA

³UC San Diego 9500 Gilman Dr. La Jolla, CA 92093

⁴8622 Kennel Way, La Jolla, CA 92037

Key Points:

- Recurrent Neural Networks (RNNs) can replace conventional numerical forecast models to produce accurate ensemble forecast statistics and linearized dynamics.
- Data assimilation (DA) is compatible with RNNs by applying state estimation in the hidden/reservoir state space with a small modification to conventional DA algorithms.
- The combined RNN-DA methods can be effectively scaled to higher dimensions by applying domain localization techniques.

Corresponding author: Stephen G. Penny, Steve.Penny@noaa.gov

Abstract

Data assimilation (DA) is integrated with machine learning in order to perform entirely data-driven online state estimation. To achieve this, recurrent neural networks (RNNs) are implemented as surrogate models to replace key components of the DA cycle in numerical weather prediction (NWP), including the conventional numerical forecast model, the forecast error covariance matrix, and the tangent linear and adjoint models. It is shown how these RNNs can be initialized using DA methods to directly update the hidden/reservoir state with observations of the target system. The results indicate that these techniques can be applied to estimate the state of a system for the repeated initialization of short-term forecasts, even in the absence of a traditional numerical forecast model. Further, it is demonstrated how these integrated RNN-DA methods can scale to higher dimensions by applying domain localization and parallelization, providing a path for practical applications in NWP.

Plain Language Summary

Weather forecast models derived from fundamental equations of physics continue to increase in detail and complexity. While this evolution leads to consistently improving daily weather forecasts, it also leads to associated increases in computational costs. In order to make a forecast at any given moment, these models must be initialized with our best guess of the current state of the atmosphere, which typically includes information from a limited set of observations as well as forecasts from the recent past. Modern methods for initializing these computer forecasts typically require running many copies of the model, either simultaneously or in sequence, to compare with observations over the recent past and ensure that our best guess estimate of the current state of the atmosphere agrees closely with those observations before making a new forecast. This repeated execution of the computer forecast model is often a time-consuming and costly bottleneck in the initialization process. Here it is shown that techniques from the fields of artificial intelligence and machine learning (AI/ML) can be used to produce simple surrogate models that provide sufficiently accurate approximations to replace the original costly model in the initialization phase. The resulting process is self-contained, and does not require any further utilization of the original computer model when making daily forecasts.

1 Introduction

Numerical weather prediction (NWP) requires two primary components: a computational forecast model and an initialization method, both of which have been demonstrated to contribute approximately equally to the steady improvement in forecast skill over the past 40 years. We seek to replace the computational forecast model with a data-driven surrogate model and integrate these two critical components. Weather forecast models typically push the boundaries of computational feasibility, even on the largest supercomputers in the world, with a drive towards increased grid resolutions and better resolved physical processes. The most sophisticated initialization methods require executing the forecast model many times, using iterative loops and ensembles of forecasts initialized from perturbed initial conditions. This creates a competing paradigm where computational resources must be balanced between model fidelity and initialization accuracy. As a result, the models serve as a major limiting factor in the development of new initialization methods.

The application of artificial intelligence and machine learning (AI/ML) methods in weather and climate is a rapidly growing activity. Boukabara et al. (2021) described multiple instances in which AI/ML is being developed for target applications in operational NWP. An area of interest noted by Boukabara et al. (2021) is the synergy between AI/ML and data assimilation (DA). Abarbanel et al. (2018) noted deep connec-




Figure 1. Data assimilation is applied to update the hidden/reservoir state space $\mathbf{s}(t)$. Observations \mathbf{y}^o are sampled from the nature state space $\mathbf{x}^{(N)}(t)$, while the composition $H \circ W_{out}$ is used as an observation operator to map the hidden/reservoir state space to an equivalent representation that can be used to form the innovations $\mathbf{d}(t) = \mathbf{y}^o(t) - H \circ W_{out}(\mathbf{s}(t))$ in the observation space.

tions between ML and DA, and in special cases mathematical equivalences (further details will be provided in an upcoming work (Abarbanel, 2022)). Recent approaches that have been applied to combine ML with DA include the application of a neural network design combined with a DA operation to train the network on noisy data (Brajard et al., 2020), the application of artificial neural networks for correcting errors in numerical model forecasts in the DA cycle (Bonavita & Laloyaux, 2020), the use of a convolution neural network (CNN) to enforce conservation of mass in a DA procedure (Ruckstuhl et al., 2021), the development of a NN-based tangent linear and adjoint model for use in variational DA methods (Hatfield et al., 2021), and an end-to-end application of a combined DA and model bias correction (Arcucci et al., 2021).

Here we will focus on a simplified form of recurrent neural network (RNN), based on the reservoir computing (RC) paradigm, that can be used to replace the numerical model in the DA process. Integrating a data-driven model with DA techniques requires accurate characterization of dynamical error growth. We will demonstrate that the RNN architecture can produce sufficiently accurate representations of such error growth to the degree that the RNN-based models can replace key components of foundational DA algorithms such as the ensemble Kalman filter (EnKF) (Evensen, 1994) and the 4D variational method (4D-Var) (Talagrand & Courtier, 1987; Courtier & Talagrand, 1987; Courtier et al., 1994). Such key components include forecast error covariance statistics derived from ensemble forecasts, and the tangent linear model and its adjoint.

We introduce a method to achieve this in a direct manner by applying DA to update the ‘hidden’ or ‘reservoir’ space of the RNN/RC dynamics (Figure 1). Results are shown assimilating both fully observed and sparsely observed dynamics, with a range of observational noise levels, using an RNN-based ensemble transform Kalman filter (ETKF) (Bishop et al., 2001; Hunt et al., 2007) and an RNN-based strong constraint incremental 4D-Var (Courtier et al., 1994).

2 Methods

The weather and climate communities regularly produce three primary categories of data: (1) observations of the Earth system, either drawn directly from in situ measurements or from satellite-based sensors, (2) simulations produced on regular gridded fields, generated from numerical models derived from basic principles of physics, and (3) retrospective analyses (reanalyses), that attempt to optimally combine the previous two categories to produce historical reconstructions of the atmosphere or other components of the Earth system. All of these datasets have inherent weaknesses. Observations are generally sparse and irregular, and can contain unpredictable errors. Numerical models have systematic errors that result from abstractions, approximations, and unresolved processes. Reanalysis products attempt to mitigate these weaknesses, but still inherit them to some degree. However, as these datasets have already been produced and archived, they are valuable resources that can be leveraged to develop data-driven methods.

We assume that for a realistic application either a long numerical model simulation or retrospective historical analysis is available as training data. Thus, we train the RNN to generate accurate predictions using a dedicated training dataset that resolves all components of the target dynamics. For this study, we use the output from the model described in section 2.4 to train our RNN models. We focus our attention on developing the capabilities to integrate the data-driven RNN model with an online DA process that repeatedly ingests new noisy and sparse observations, updates the state estimate of the system, and makes new short-term forecasts. This has applications ranging from operational NWP to the efficient reconstruction of historical Earth system states.

2.1 Network Design

Reservoir computing is a category of machine learning methods. It originated in the works of Jaeger (2001), who introduced Echo State Networks (ESN), and Maass et al. (2002) who introduced Liquid State Machines (LSM). Both methods assume that an input signal can be mapped to a fixed nonlinear system of higher dimension than the input. Such systems can be trained using a readout layer to map the state of the reservoir to the desired output. The result is a simple model that can reproduce the potentially complex dynamics of the original system. A survey is provided by Lukoševičius & Jaeger (2009), while Konkoli (2017) provides further discussion on the generality of RC. Successful applications of RC have been demonstrated for the prediction of spatiotemporally chaotic dynamics (Chattopadhyay et al., 2020; Pathak et al., 2018; Platt et al., 2021), and in particular of geophysical fluid dynamics (Arcomano et al., 2020; Lin & Penny, 2021-submitted).

RNNs have long been used as a preferred ML method for cases in which temporal considerations are necessary. We use a basic RNN (Elman, 1990) with a simplified structure to produce a variant of the RC network. The RNN/RC design is useful for prediction because the ‘hidden’ or ‘reservoir’ state $\mathbf{s}(t_i)$ provides a short-term memory of the target system trajectory $\mathbf{x}(t)$ up to time t_i . Assuming an accurate forecast can be made of the hidden/reservoir state $\mathbf{s}(t_{i+1})$ by a well-trained RNN, then that forecast can be mapped to the target system space to produce a forecast $\mathbf{x}'(t_{i+1})$ of the true system state $\mathbf{x}(t_{i+1})$.

As highlighted by Schrauwen et al. (2007), Steil (2004) showed that the state-of-the-art learning rule for RNNs at the time had the same weight dynamics as the methods proposed by Jaeger (2001) and Maass et al. (2002). The Atiya & Parlos (2000) recurrent learning rule trains the output weights while the internal weights are only globally scaled up or down. Similarly, we classify the RNN model parameters as either ‘macro-scale’ or ‘micro-scale’. We note that what we call ‘macro-scale’ parameters were called ‘global parameter’ by Lukoševičius (2012). In what follows, all matrix elements are classified as micro-scale parameters while all scalars are classified as macro-scale parameters.

ters. The general form of our RNN is given as,

$$\mathbf{s}(t_{i+1}) = F(\mathbf{s}(t_i), \mathbf{x}(t_i)) = l * f(\rho \mathbf{W}_{res} \mathbf{s}(t_i) + \sigma \mathbf{W}_{in} \mathbf{x}(t_i)) + (1 - l) * \mathbf{s}(t_i), \quad (1)$$

$$\mathbf{x}'(t_{i+1}) = G(\mathbf{s}(t_{i+1})) = \mathbf{W}_{out}(\mathbf{s}(t_{i+1})), \quad (2)$$

where $\mathbf{x}(t_i)$ is the system state at time t_i , provided from data, $\mathbf{s}(t_i)$ is the hidden/reservoir state, and $\mathbf{x}'(t_{i+1})$ is the predicted system state at the next time t_{i+1} . The parameter ρ determines the spectral radius of the reservoir adjacency matrix \mathbf{W}_{res} , σ scales the input signal mapped to the reservoir space by \mathbf{W}_{in} , and l is the ‘leak’ parameter that gates new information into the system. We will use $f = \tanh$.

After training the RNN model parameters, we expect the predicted state $\mathbf{x}'(t_{i+1})$ to be close to the true system state $\mathbf{x}(t_{i+1})$. In a typical RNN, all parameters of the system described by equations (1) and (2) are trained, usually by a gradient descent type optimization method. For RC, all model parameters in equation (1) are assumed fixed. Equation (1) is then iterated to generate a time series $\mathbf{s}(t)$ of hidden/reservoir states that corresponds to the training data $\mathbf{x}(t)$. The ‘readout’ operator \mathbf{W}_{out} in equation (2) is typically assumed to be linear.

The micro-scale parameters in the \mathbf{W}_{out} are trained (while the other parameters of our RNN remain fixed) by solving a regularized least squares equation using a loss function that targets forecasts that are one time step in the future (Jaeger, 2001). The corresponding loss function is,

$$\mathcal{L}_{micro}(\mathbf{W}_{out}) = \|\mathbf{W}_{out} \mathbf{S}_{data} - \mathbf{X}_{data}\|^2 + \beta \|\mathbf{W}_{out}\|^2, \quad (3)$$

where \mathbf{S}_{data} and \mathbf{X}_{data} are matrices that comprise the vector-valued states $\mathbf{s}(t)$ and $\mathbf{x}(t)$ for the entire training dataset, ordered columnwise, and β is a Tikhonov regularization parameter.

To move closer towards the approach of the general RNN, in which all parameters are trained, we regard the scalars l , ρ , σ , and β as macro-scale parameters subject to training. The matrix \mathbf{W}_{res} is initialized with a spectral radius of 1 prior to training, but otherwise the values are assigned randomly using a uniform distribution centered at 0. For computational efficiency, \mathbf{W}_{res} is assumed to be sparse (i.e. only 1% of entries are nonzero). The matrix \mathbf{W}_{in} is initialized using a uniform random distribution with values ranging from -1 to 1.

An extended forecast $\mathbf{x}^f(t)$ is made with the RNN from time t_0 to t_i , for $i > 0$, by recursively replacing the input state with the RNN prediction initialized from the previous time. Defined inductively, commencing with $\mathbf{x}(t_0)$,

$$\mathbf{x}^f(t_1) = \mathbf{W}_{out} \circ F(\mathbf{s}(t_0), \mathbf{x}(t_0)), \quad (4)$$

$$\mathbf{x}^f(t_i) = \mathbf{W}_{out} \circ F(\mathbf{s}(t_{i-1}), \mathbf{x}^f(t_{i-1})). \quad (5)$$

We note that due to the nature of chaotic dynamical systems (with leading Lyapunov exponent greater than 0), any error in one step of this recursion will accumulate and lead to exponential error growth over time.

The macro-scale parameters are trained using a nonlinear Bayesian optimization method, which uses a surrogate model representation of the loss function (Jones et al., 1998; Ginsbourger et al., 2010). For this macro-scale optimization we use a loss function that targets longer-range prediction,

$$\mathcal{L}_{macro}(\mathbf{x}^f(t)) = \sum_{i=1}^M \sum_{t=t_i}^{t_{i+N}} \|\mathbf{x}^f(t) - \mathbf{x}(t)\|^2 \exp\left(-\frac{t-t_i}{t_{i+N}-t_i}\right), \quad (6)$$

where M represents the number of separate initial times $[t_1, t_2, t_3, \dots, t_M]$ used to make independent forecasts, selected randomly without replacement from the training dataset, and N represents the number of time steps used for each forecast. We apply an exponential scaling term in order to account for the exponential growth of errors that is typical of chaotic dynamics. This term gives the forecast errors in the earlier portion of the forecast more weight, as this period is more relevant for cycled DA applications.

Recall that by definition a hyperparameter is any design decision that is set before the learning process begins, is generally tunable, and can directly affect how well a model trains. The hyperparameters of this optimization are provided in the appendix in Table A1. We note that Griffith et al. (2019) similarly applied long forecasts at multiple initial times in a reservoir computing application to identify model parameters that resulted in the reconstruction of the full system attractor.

By computing the Jacobian of the forecast model defined by equations (1), (2), (4), and (5) with respect to the hidden/reservoir state, we can determine the linear propagator \mathbf{M} of the reservoir dynamics from time t_i to t_{i+1} as,

$$\mathbf{W} = \rho \mathbf{W}_{res} + \sigma \mathbf{W}_{in} \mathbf{W}_{out}, \quad (7)$$

$$\mathbf{M}_{[t_{i+1}, t_i]}(\mathbf{s}(t_i)) = \frac{dF}{d\mathbf{s}} = l * \text{diag}(1 - \tanh(\mathbf{W}\mathbf{s}(t_i))^2) \mathbf{W} + (1 - l) * \mathbf{I}. \quad (8)$$

The linear propagator describes the evolution of small perturbations from a reference trajectory. It can be used in DA, in particular 4D-Var, where it is called the tangent linear model (TLM). Further, the Lyapunov exponents of the system are determined by integrating the linear propagator from time $t = 0 \rightarrow \infty$ and computing the eigenvalues of the resulting system. Practical algorithms based on QR decompositions are provided by Geist et al. (1990). For computational efficiency, we implement the TLM and its adjoint as linear operators to avoid matrix multiplications and allow for efficient matrix-vector operations applied within the iterative minimization schemes.

2.2 Data Assimilation

Trevisan et al. (2010), Trevisan & Palatella (2011a), and Palatella et al. (2013) showed that the number of observations needed to constrain any DA system is related to the number of non-negative Lyapunov exponents in the system. Platt et al. (2021) indicated that reproducing the Lyapunov spectrum is critical to generating accurate predictions with reservoir computing models - with deviations from the true spectrum leading to significantly degraded forecast skill. Considering these points, we presume that even if the hidden/reservoir state space is large, if the RNN is trained to be sufficiently accurate such that the true Lyapunov spectrum is well approximated, then the number of observations required to constrain the hidden state space dynamics should be the same as is required to constrain the original system dynamics.

Following this presumption, we apply DA in the hidden/reservoir space of the RNN system, and apply the composition of an observation operator with the readout operator in order to compare hidden/reservoir states with observations of the original system. To test this approach, we apply two well known DA algorithms integrated with the RNN forecast model - the ensemble transform Kalman filter (Bishop et al., 2001; Hunt et al., 2007) and the strong constraint incremental 4-dimensional variational method (4D-Var) (Courtier et al., 1994).

From the perspective of operational forecasting, the RNN provides a simple and low-cost replacement for the production of essential information needed for the online

DA cycle, such as forecast error covariance statistics and the tangent linear and adjoint model dynamics. From the machine learning perspective, the DA algorithms allow the RNN hidden/reservoir dynamics to be driven with a noisy and sparsely observed signal. We will show that in cases where the direct insertion of observations quickly corrupts the hidden/reservoir state and leads to inaccurate forecasts, the DA methods can produce valid reconstructions of the system state as well as viable initial conditions for short-term forecasts.

Kalnay et al. (2006) described the ideal initial ensemble perturbations as those that effectively span the space defined by the analysis error covariance. We use ensemble forecast statistics produced by the RNN model to generate dynamically estimated forecast error covariance statistics, and then apply an ETKF to assimilate noisy observations of the true system state and estimate the analysis error covariance. Bocquet & Carrassi (2017) showed that the minimum ensemble size required to constrain a (non-localized) deterministic ensemble filter such as the ETKF is equal to the number of non-negative Lyapunov exponents of the system dynamics. Thus, we expect that with a well-trained RNN that closely approximates the correct Lyapunov spectrum, the minimum number of ensemble members needed to constrain the ETKF will be the same as the number of members needed to constrain the original system.

We define a new modified observation operator by composing the conventional observation operator $H()$, which maps from the system space to the observation space, with the readout operator $W_{out}()$, which maps from the hidden/reservoir space to the system space (see Figure 1). Our implementation of the ETKF follows the formulation of Hunt et al. (2007). Let $\bar{\mathbf{y}}^b = H(W_{out}(\bar{\mathbf{s}}^b))$, where $\bar{\mathbf{s}}^b$ is the background ensemble mean hidden/reservoir state, and $\mathbf{Y}^b = H(W_{out}(\mathbf{S}^b))$, where the columns of \mathbf{S}^b are ensemble perturbations around the mean, then

$$\tilde{\mathbf{P}}^a = \left[\frac{k-1}{\gamma} \mathbf{I} + (\mathbf{Y}^b)^T \mathbf{R}^{-1} \mathbf{Y}^b \right]^{-1}, \quad (9)$$

$$\mathbf{W}^a = \left[(k-1) \tilde{\mathbf{P}}^a \right]^{\frac{1}{2}}, \quad (10)$$

$$\mathbf{S}^a = \mathbf{S}^b \mathbf{W}^a, \quad (11)$$

$$\bar{\mathbf{w}}^a = \tilde{\mathbf{P}}^a (\mathbf{Y}^b)^T \mathbf{R}^{-1} (\mathbf{y}^o - \bar{\mathbf{y}}^b), \quad (12)$$

$$\bar{\mathbf{s}}^a = \mathbf{S}^b \bar{\mathbf{w}}^a + \bar{\mathbf{s}}^b, \quad (13)$$

where \mathbf{R} is the observation error covariance matrix, k is the ensemble dimension, γ is a multiplicative inflation factor, $\tilde{\mathbf{P}}^a$ is the analysis error covariance matrix represented in the ensemble perturbation subspace, \mathbf{W}^a is applied as a transform operator to map the background ensemble perturbations to analysis ensemble perturbations, and $\bar{\mathbf{w}}^a$ determines the weighting coefficients of the column vectors \mathbf{S}^b , which are used as a linear basis to form the new ensemble mean analysis state vector $\bar{\mathbf{s}}^a$. For reference, the resulting Kalman gain for the integrated RNN-ETKF is of the form,

$$\mathbf{K} = \mathbf{S}^b \left[\frac{k-1}{\gamma} \mathbf{I} + [H(W_{out}(\mathbf{S}^b))]^T \mathbf{R}^{-1} [H(W_{out}(\mathbf{S}^b))] \right]^{-1} [H(W_{out}(\mathbf{S}^b))]^T \mathbf{R}^{-1}. \quad (14)$$

The control vector for 4D-Var can similarly be formed in the hidden/reservoir space. We use the strong constraint incremental 4D-Var, implemented using an outer and inner loop. In the outer loop, a nonlinear forecast $\mathbf{s}_t^f = \mathcal{M}_{[t,0]}(\mathbf{s}_0)$ is generated over a short optimization period, called the analysis window. In the inner loop, the linearized dynamics are used to find an improved guess for the initial state \mathbf{s}_0 using an iterative linear solver, and then the outer loop is repeated. If we let $\delta\mathbf{s}_0 = (\mathbf{s}_0 - \mathbf{s}_0^f)$, $\delta\mathbf{s}_0^b = (\mathbf{s}_0^b - \mathbf{s}_0^f)$, and $\mathbf{d}_t = (\mathbf{y}_t^o - H_t \circ W_{out}(\mathbf{s}_t^f))$, then the objective function is,

$$J(\delta\mathbf{s}_0) = J_b(\delta\mathbf{s}_0) + J_o(\delta\mathbf{s}_0), \quad (15)$$

where,

$$J_b(\delta\mathbf{s}_0) = \frac{1}{2}(\delta\mathbf{s}_0 - \delta\mathbf{s}_0^b)^T \mathbf{B}^{-1}(\delta\mathbf{s}_0 - \delta\mathbf{s}_0^b), \quad (16)$$

$$J_o(\delta\mathbf{s}_0) = \frac{1}{2} \sum_{t=0}^{N_t} (\mathbf{d}_t - H_t \circ W_{out}(\mathbf{M}_{[t,0]}\delta\mathbf{s}_0))^T \mathbf{R}_t^{-1} (\mathbf{d}_t - H_t \circ W_{out}(\mathbf{M}_{[t,0]}\delta\mathbf{s}_0)), \quad (17)$$

$$\mathbf{M}_{[0,0]} = \mathbf{I}, \quad (18)$$

$$\mathbf{M}_{[t+1,t]} = l * \text{diag}(1 - \tanh(\mathbf{W}\mathbf{s}(t))^2) \mathbf{W} + (1 - l) * \mathbf{I}, \quad (19)$$

as in equation (8), and the initial condition in the original system space can be recovered by,

$$\mathbf{x}_0 = W_{out}(\mathbf{s}_0). \quad (20)$$

In our implementation, $H_t()$ and $W_{out}()$ are linear, so we replace them with their matrix notation. When these operators are not linear, a linear approximation via Taylor series expansion is typically applied. The minimum is found when the gradient with respect to the control vector $\delta\mathbf{s}_0$ equals 0,

$$\nabla_{\delta\mathbf{s}_0} J = \mathbf{B}^{-1} (\delta\mathbf{s}_0 - \delta\mathbf{s}_0^b) \sum_{t=0}^{N_t} \mathbf{M}_{[t,0]}^T \mathbf{H}_t^T \mathbf{R}_t^{-1} (\mathbf{d}_t - \mathbf{H}_t \mathbf{M}_{[t,0]} \mathbf{W}_{out} \delta\mathbf{s}_0) = 0. \quad (21)$$

We solve this by separating terms into the form ' $\mathbf{Ax}=\mathbf{b}$ ', where $\delta\mathbf{s}_0$ is the only unknown quantity,

$$\left(\mathbf{I} + \mathbf{B} \sum_{t=0}^{N_t} \mathbf{H}_t \mathbf{M}_{[t,0]} \right) \delta\mathbf{s}_0 = \mathbf{B} \sum_{t=0}^{N_t} \mathbf{M}_{[t,0]}^T \mathbf{H}_t^T \mathbf{R}_t^{-1} \mathbf{d}_t + \delta\mathbf{s}_0^b \quad (22)$$

and then applying the biconjugate gradient stabilized method (Van der Vorst, 1992). Alternative forms of equation (22) are available, for example making the ' \mathbf{A} ' matrix symmetric so that the conjugate gradient method can be applied. Returning to the outer loop, a new nonlinear forecast $\mathbf{s}^f(t) = \mathcal{M}_{[t,0]}(\mathbf{s}_0^f + \delta\mathbf{s}_0)$ is generated and the entire process is repeated with the goal of converging to the optimal nonlinear trajectory.

2.3 Localization

For the RNN model itself, scalability is enabled by partitioning the model system domain into smaller local patches, with a separate RNN trained for each patch (Pathak et al., 2018). Each local patch is assigned a small radius of ‘halo’ points that allow information from neighboring patches as inputs to the RNN model, while computing a forecast only for the points within the patch. This follows a similar paradigm to that used for the domain decomposition of general circulation models. Localization of a geophysical forecast model is motivated by the presence of locally low dimensional chaotic dynamics (Oczkowski et al., 2005). In previous works, Arcomano et al. (2020) demonstrated the use of RC for prediction of global scale atmospheric dynamics by applying the localization scheme described above, while Lin & Penny (2021-submitted) further showed that this spatial localization approach could be improved for geophysical systems by applying transformations into Fourier space.

Localization has also been an important tool for scaling DA methods to enable application to high dimensional systems (Greybush et al., 2011). The same localization procedure used to scale the RNN model can be applied in the context of DA, which provides a path to scaling the RNN-based DA methods to more realistic high-dimensional applications. We apply localization of the DA using an approach similar to the Local Ensemble Transform Kalman Filter (LETKF) (Hunt et al., 2007). In its original formulation, the LETKF computes a separate ETKF analysis at each model grid point, while only assimilating observations within a prescribed localization radius around that grid point. In our RNN-LETKF, we instead choose a radius around local patches. Observations are selected from an area larger than the patch itself based on the RNN localization in order to promote consistency with the analyses computed for neighboring patches. We make the design decision to maintain correspondence with the local RNN model architecture by using a radius that aligns with the input field of the local RNN, which includes the local patch and its halo points. As with the LETKF, the RNN-LETKF local analyses can be computed in parallel, after observation innovations are computed globally and distributed to each local patch. The results of the local analyses are then communicated to the neighboring patches in order to initialize the next forecast.

2.4 Source Data for Training, Validation, and Testing

Here we describe the underlying model equations that we use to generate data for training, validating, and testing the RNN models. Lorenz (1996) developed a simple model (L96) that includes advection, dissipation, and external forcing to describe basic wave-like dynamics in the atmosphere around a latitude ring. The L96 model is a frequently used test system for DA studies (Abarbanel et al., 2010; Penny, 2014, 2017; Goodliff et al., 2017; Chen & Kalnay, 2019; Brajard et al., 2020), and multiple varieties of RNNs have been applied successfully for emulation of L96 dynamics (Vlachas et al., 2020). The L96 system is defined by a set of ordinary differential equations on a discrete finite cyclic domain,

$$f_{L96}(\mathbf{x}_i) = \frac{d\mathbf{x}_i}{dt} = \mathbf{x}_{i-1}(\mathbf{x}_{i+1} - \mathbf{x}_{i-2}) - \mathbf{x}_i + F_{L96}. \quad (23)$$

We use forcing $F_{L96} = 8.0$, which is sufficient to achieve chaotic dynamics, meaning that at least one Lyapunov exponent is greater than 0. This implies initial errors will grow exponentially on average. The model is integrated with a timestep of $\delta t = 0.01$ model time units (MTUs). Lorenz (1996) scaled the coefficients of the model so that the error growth over 1 MTU is roughly equivalent to 5 days in the atmosphere, relative to the state-of-the-art atmospheric models of the time. In operational prediction centers, analyses are often produced with 6-hour, 12-hour, or 24-hour update cycles, thus we will focus on DA cycles ranging up to 0.2 MTU (≈ 24 hours).

2.5 Experiment Design

We compare results using a variety of DA configurations. All model integrations are computed with a timestep of $\delta t=0.01$. Unless otherwise noted, we use the following parameter settings: each DA experiment is integrated for 100 MTU (or 10,000 time steps); we use observation noise of $\sigma_{noise}=0.5$ and a corresponding estimated observation error of $\sigma_{obs}=0.5$ to form \mathbf{R} ; observations are sampled every $\tau_{obs}=0.02$ MTU (≈ 2.4 hours); and we use an analysis cycle window of $\tau_{da}=0.2$ MTU (≈ 24 hours).

We use a 10-member ensemble for the RNN-ETKF, a 30-member ensemble for the RNN-LETKF applied to the higher dimensional L96-40D model configuration, and for the purposes of this discussion a 1-member ‘ensemble’ for the RNN-4DVar. Both the RNN-ETKF and the RNN-4DVar were applied to the L96-6D model. All DA experiments are initialized by first preparing a set of perturbed spinup datasets, one for each ensemble member, applying Gaussian random noise with standard deviation $\sigma_{init}=0.5$ to the true state over a 1000 time step window (2000 for the L96-40D system). Each RNN ensemble member is synchronized with its corresponding perturbed dataset in order to produce an initial ensemble of hidden/reservoir states that reflect the uncertainty present in the noisy input data.

As noted by Lorenc (2003), additional covariance inflation is needed in the presence of model error if that error is not addressed explicitly. Covariance inflation is also typically needed due to the use of a finite ensemble size. We found an inflation parameter of 1-5% (i.e. $\gamma = 1.01$ to 1.05) to be effective for the ETKF when applied with the ‘perfect’ numerical model. To account for model error in the RNN, we increase the inflation parameter to 20% ($\gamma=1.2$) for the RNN-ETKF and 30% ($\gamma=1.3$) for the RNN-LETKF to account for errors in the RNN model. The 4D-Var uses an empirically chosen static error covariance that is diagonal with standard error equal to $\sigma_b = \sigma_{obs}$, an analysis time at the start of the window, and uses 2 outer loops.

To mimic a realistic scenario of geophysical prediction, we focus mainly on cases where the variables are sparsely observed. If not otherwise stated, the L96-6D model is observed only at the first, second, and fourth nodes. For the L96-40D model, we limit the observing network to only 15 nodes.

3 Results

3.1 Assessment of error growth rates

An essential consideration of DA is the behavior of long and short term error growth, which can be characterized by the Lyapunov exponent (LE) spectrum and finite-time (also known as ‘local’) Lyapunov exponents (FTLEs) (Abarbanel et al., 1992; Abarbanel, 1996). If one considers DA as the synchronization of a model with the natural process from which measurements are drawn, then the conditional LE spectrum of this coupled model-nature system must be driven negative to ensure the model synchronizes with the observed system (Penny, 2017). Previous studies (Pathak et al., 2018; Griffith et al., 2019; Platt et al., 2021) have already shown that reservoir computing can be used to reproduce the Lyapunov spectrum of the source system. This spectrum characterizes the long time average exponential growth rates of small errors in the system trajectory. However, at very short timescales, we find that the error growth rates of our trained RNNs are not well representative of the error growth rates produced by the source system dynamics. We find instead that the FTLEs of the RNN dynamics converge towards the Lyapunov exponents of the source system dynamics over a transient period of a few Lyapunov timescales (Figure 2). While the growth rates of errors in the RNN models take some time to converge to the true growth rates, the forecast error correlations appear to be estimated relatively accurately at short lead times (an example is shown in Figure 3). This indicates that while it may be desirable to improve the convergence rates of the FTLEs produced

by the RNN, the effect can be compensated for by using a scalar multiplicative inflation applied to the forecast error covariance matrix.




Figure 2. Convergence of the leading FTLE (λ_1) for a trained RNN (model 1 in Table A2) averaged over 100 initial conditions of the L96-6D system. As the RNN is integrated for longer periods of time, the error growth rates generated by the RNN model become more accurate. However, over the same period there is an exponential growth of errors in initial conditions.

3.2 Control RNN case

To demonstrate the need for DA, we commence our cycled forecast experiments with a control case that sets a baseline for the performance of the RNN without using DA. Here, observations are inserted directly into the RNN as defined by as in equations (1) and (2). If the system is fully observed, then this amounts to replacing $\mathbf{x}(t_i)$ in equation (1) with observed data.

Lu et al. (2017) examined a ‘sparse in space’ case that limited the forcing of an RC model to only a subset of inputs. Their results showed that synchronization of the full state can be achieved even when observing only a subset of the variables of the system. However, we find that this type of direct insertion method for synchronization fails as the observations become more sparse in time (see Figure 4 and the middle and bottom rows of Figure 5).

Before further evaluating the RNN, we first consider for reference the case of direct insertion of observations into the original L96-6D numerical model, using the update equation,

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{H}^T(\mathbf{y}^o - \mathbf{H}\mathbf{x}^b). \quad (24)$$

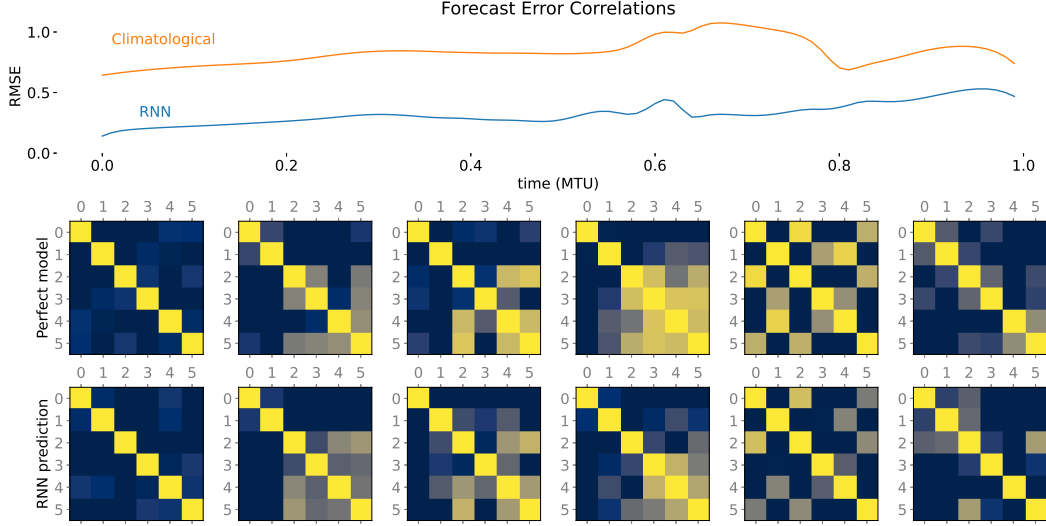


Figure 3. (Top) RMSE of the forecast error correlation matrix over time for the L96-6D system, comparing an example RNN ensemble forecast (blue) to the climatological error correlation matrix (orange), both evaluated versus a 100-member ensemble forecast using the perfect reference model. The initial conditions are sampled from the test dataset, and initial ensemble perturbations are generated using a Gaussian distribution with standard deviation 0.1. (Bottom) Forecast error correlation matrix for the reference perfect numerical model and the RNN model over the same forecast period, valid at times 0.0, 0.2, 0.4, 0.6, 0.8, and 1.0. The color scale ranges from 0 to 1.

As should be expected, providing perfect observations of all variables at every model time step ($\delta t=0.01$) produces exact synchronization between the driver signal and the numerical model trajectory. Increasing the timestep between observations as high as 0.2 MTU does not significantly degrade the state estimates, with errors peaking at $2.5e-7$. Increasing the observational noise generally increases the error in the state estimates by a similar magnitude. When the number of observed variables is reduced (e.g. to 3 or 2 out of 6), the system experiences transient synchronization with occasional bursting. While still using perfect observations, combining reduced observations (e.g. 50%) and using longer timesteps (e.g. $\delta t=0.1$) actually improves stability compared to using a time step of $\delta t=0.01$, and leads to synchronization. However, when observation noise is added to this combination of sparseness in space and time, the bursting phenomena return, particularly in the unobserved variables (see top row of Figure 5).

The situation is quite different with the RNN model. Increasing the time step between the (noise-free) observations significantly degrades the RNN estimates, first adding high frequency oscillations (e.g. with $\delta t=0.02$ to $\delta t=0.1$), and then leading to trajectories with little discernible connection to the L96 dynamics (e.g. at $\delta t=0.2$). Recall that for these experiments the RNN is trained on data that have a temporal resolution matching the underlying model timestep $\delta t=0.01$. Reducing the frequency of the input driving signal allows the hidden/reservoir state to drift. Observing frequently ($\delta t=0.01$) but removing the observation of one variable degrades the estimates of that variable without noticeably affecting the rest, while removing the observation of more than one variable can cause occasional degradation of the remaining observed variables. Due to the presence of the hidden/reservoir state, which maintains a memory of the past trajectory, the RNN itself is relatively insensitive to the introduction of noise to the observations. When observing the full system state, noisy observations supplied as driving data to the

RNN simply increase high frequency noise in the analyzed state estimate, without resulting in divergence between the RNN and the true signal. Combining any of these constraints on the RNN appears to have additive effects. A comparison of the errors in the RNN using direct insertion with a range of observation noise and observing frequency (of which Figure 4 is one instance) is shown in Figure 5.

We note that even for this simple L96 model, the total set of direct insertion experiments using RNN model 1 was about 10% faster, and RNN model 2 about 30-40% faster, than the total set of direct insertion experiments using the conventional numerical integration of the L96 differential equation (23). We do not claim that these results can be easily extrapolated to other applications. However, we do emphasize that the projection of the system dynamics to the higher dimensional hidden/reservoir state does not necessarily imply that the computations become more costly.

To summarize - simply providing the sparse and noisy observation data directly to the RNN is not adequate for initializing forecasts, which provides motivation for the use of a more sophisticated DA strategy.




Figure 4. Direct insertion using the RNN with perfect ($\sigma_{noise}=0$) observations of the system at points (1,2,4), sampled with $\tau_{obs}=0.05$ (≈ 6 hours), which is every 5 model time steps. The RNN alone cannot successfully recover the true trajectory when observations are sparse and noisy. This figure provides a corresponding entry in figure 5.

3.3 RNN-DA with sparsely observed dynamics

DA methods provide most of their value when observations are sparse and noisy. Here, we restrict the observing network to only the first, second, and fourth nodes of the 6-dimensional cyclic L96 system (L96-6D). This leaves two patches that are unobserved for the duration of each experiment - the third node and the combined fifth and sixth node. We noticed no qualitative differences in other configurations of the observing system layout at the same 50% coverage level.

We first implement the RNN-ETKF and compare to the conventional ETKF using the ‘perfect’ numerical model. We find that with the exception of the case in which analyses are updated frequently ($\tau_{DA}=\tau_{obs}=\delta t = 0.01$), the RNN-ETKF using both RNN models 1 and 2 performs quite well in comparison (Figures 6 and 7). We note that the

Figure 5. Normalized RMSE using direct insertion using (Top) the 'perfect' numerical model and (Bottom) the RNN model 1. Both are applied to the L96-6D system integrated over 100 MTU, varying the observation noise σ_{noise} and the observation timestep Δt_{obs} . (Left) RMSE of observed points (indexes 1,2,4). (Right) RMSE of points that are not observed (indexes 3,5,6). A normalized RMSE of 1.0 equals the L96 system's climatological standard deviation. Note that the conventional model is more sensitive to increased noise, while the RNN model is more sensitive to the observing frequency.

total set of ETKF experiments using RNN model 1 had about equal run time, while RNN model 2 was about 25% faster, compared to the run time of the total set of ETKF experiments using the conventional numerical integration of the L96 differential equation (23).

The RNN-4DVar method performs well when the observational noise is small, but is sensitive to increasingly sparse and noisy observing sets (Figures 8 and 9). As the underlying RNN model is improved (from model 2 to model 1 in Table A2), this appears to improve the performance of the 4D-Var correspondingly. The sensitivity of the RNN-4DVar to observational noise may be exacerbated by errors in the RNN model equations from which the TLM and adjoint operators are derived, and also the approximated background error covariance matrix. A further drawback is that the experiments using the RNN-4DVar required 1-2 orders of magnitude more computational time than the conventional 4D-Var applied using the numerical integration of the L96 differential equation (23) and its TLM and adjoint equations.

The difference between FTLEs estimated by the RNN and the numerical model at short lead times indicates that the linearized dynamics (i.e. the TLM and adjoint) are not well represented at these timescales. The RNN models used here generally under-represent the magnitude of error growth at short timescales. This affects the ETKF as well but is alleviated by the application of multiplicative inflation, and gives some explanation for why the ETKF is more stable than 4D-Var. The strong-constraint 4D-Var used here assumes a 'perfect' model. We expect that transitioning from the strong-constraint 4D-Var formulation to the weak-constraint 4D-Var approach should further improve the RNN-4DVar performance, as the latter explicitly accounts for errors in the model.

Figure 10. Normalized error of the RNN-LETKF based state estimation for the L96-40D system using RNN Model 3. (Top) Normalized Root Square Error (NRSE) shown for each node of the L96 system (y-axis). (Bottom) NRMSE computed separately for the observed and unobserved nodes of the Lorenz system, with 15 nodes of the system observed. Note the y-axis is logarithmic in the lower plot. The error in both plots are normalized by the temporal standard deviation of the true trajectory. The RNN-LETKF uses a 30 member ensemble, with $\sigma_{\text{obs}} = \sigma_{\text{noise}} = 0.5$, and macro-scale parameters indicated by RNN model 3 in Table A2. The observed nodes of the system are [Q 3; 5; 8; 10; 14; 16; 19; 20; 25; 27; 30; 34; 36; 39].

ner loop, or alternatively the RNN could be used as the nonlinear model in the outer loop to reduce computational costs while still using a numerical tangent linear model and adjoint in the inner loop. We also note that while the RNN models were trained once on historical data and held fixed during the RNN-DA cycling, it is straightforward to perform online retraining of the RNN model during the DA cycle, e.g. as suggested by Brjard et al. (2020).

Further methods are needed to optimize the design and training of the ML model used in this study to replace the numerical forecast model. More development is needed in ML modeling of chaotic dynamics to ensure a rapid convergence of the finite time Lyapunov exponents toward the true values. While a simplified RNN resembling the reservoir computing approach was applied here, we expect that our approach could be applied with more sophisticated types of RNN that also use hidden state representations, such as Long Short Term Memory (LSTM) or Gated Recurrent Unit (GRU) architectures.

Figure A1. Normalized macro loss function, L_{macro} (see equation 6) for various M visualized in the $(\theta; \phi)$ plane. Increasing M regularizes the loss function, revealing an approximate global minimum in parameter space. The loss function shown here is based on an 800 dimensional hidden/reservoir state RNN model prediction of the L96-6D system.

Section 2.4). When $M=1$, the macro loss function exhibits many local minima. This indicates that the resulting trained RNN may produce accurate forecasts for a particular set of initial conditions, but does not generalize well to other points on the system attractor. Evaluating forecasts from a greater number of initial conditions results in a smoother loss function landscape. However, increasing M also carries a corresponding increase in cost for the evaluation of the macro loss function. We select $M=100$ to balance the computational cost and the diminishing returns observed beyond this point.

We establish a ‘Valid Prediction Time’ (VPT) as the length of time that the RMSE of any forecast starting at time t_0 remains below a threshold ϵ . The VPT is defined precisely as,




Figure A2. A free forecast of the RNN model 1 initialized from a random point in the test dataset, compared to the true trajectory. The valid prediction time (VPT) using $\epsilon=0.2$ is marked as a vertical line.




Figure A2. A free forecast of the RNN model 1 initialized from a random point in the test dataset, compared to the true trajectory. The valid prediction time (VPT) using $\epsilon=0.2$ is marked as a vertical line.




Figure A3. Histogram of valid prediction time (VPT) for RNN model 1 forecasts drawn from 100,000 initial conditions in the test dataset.




Figure A4. As in Figure A3, but using RNN model 2. Note the model is skillful, but the mean VPT is reduced compared to model 1.




Figure A5. As in Figures A3 and A4, but using RNN Model 3. Due to the computational cost of Model 3, only 1,000 sample forecasts are used. The model is still skillful, but the mean VPT is reduced when compared to Models 1 and 2.

- 631 Arcomano, T., Szunyogh, I., Pathak, J., Wikner, A., Hunt, B. R., & Ott, E. (2020).
 632 A machine learning-based global atmospheric forecast model. *Geophysical Re-*
 633 *search Letters*, 47(9), e2020GL087776.
- 634 Arcucci, R., Zhu, J., Hu, S., & Guo, Y.-K. (2021). Deep data assimilation: integrat-
 635 ing deep learning with data assimilation. *Applied Sciences*, 11(3), 1114.
- 636 Atiya, A., & Parlos, A. (2000). New results on recurrent network training: unifying
 637 the algorithms and accelerating convergence. *IEEE Transactions on Neural Net-*
 638 *works*, 11(3), 697-709. doi: 10.1109/72.846741
- 639 Bishop, C. H., Etherton, B. J., & Majumdar, S. J. (2001). Adaptive sampling with
 640 the ensemble transform kalman filter. part i: Theoretical aspects. *Monthly weather*
 641 *review*, 129(3), 420–436.
- 642 Bocquet, M., & Carrassi, A. (2017). Four-dimensional ensemble variational data
 643 assimilation and the unstable subspace. *Tellus A: Dynamic Meteorology and*
 644 *Oceanography*, 69(1), 1304504. Retrieved from [https://doi.org/10.1080/](https://doi.org/10.1080/16000870.2017.1304504)
 645 [16000870.2017.1304504](https://doi.org/10.1080/16000870.2017.1304504) doi: 10.1080/16000870.2017.1304504
- 646 Bonavita, M., & Laloyaux, P. (2020). Machine learning for model error infer-
 647 ence and correction. *Journal of Advances in Modeling Earth Systems*, 12(12),
 648 e2020MS002232.
- 649 Bouhlel, M. A., Hwang, J. T., Bartoli, N., Lafage, R., Morlier, J., & Martins,
 650 J. R. R. A. (2019). A python surrogate modeling framework with deriva-
 651 tives. *Advances in Engineering Software*, 102662. doi: [https://doi.org/10.1016/](https://doi.org/10.1016/j.advengsoft.2019.03.005)
 652 [j.advengsoft.2019.03.005](https://doi.org/10.1016/j.advengsoft.2019.03.005)
- 653 Boukabara, S., Krasnopolsky, V., Penny, S. G., Stewart, J. Q., McGovern, A., Hall,
 654 D., ... Hoffman, R. N. (2021). Outlook for exploiting artificial intelligence in the
 655 earth and environmental sciences. *Bulletin of the American Meteorological Society*,
 656 102(5), E1016-E1032.

- 657 Brajard, J., Carrassi, A., Bocquet, M., & Bertino, L. (2020). Combining data assim-
658 ilation and machine learning to emulate a dynamical model from sparse and noisy
659 observations: a case study with the lorenz 96 model. *Journal of Computational*
660 *Science*, 44, 101171.
- 661 Chattopadhyay, A., Hassanzadeh, P., & Subramanian, D. (2020). Data-driven
662 predictions of a multiscale lorenz 96 chaotic system using machine-learning meth-
663 ods: reservoir computing, artificial neural network, and long short-term memory
664 network. *Nonlinear Processes in Geophysics*, 27(3), 373–389.
- 665 Chen, T.-C., & Kalnay, E. (2019, January). Proactive quality control: Observ-
666 ing system simulation experiments with the lorenz '96 model. *Monthly Weather*
667 *Review*, 147, 53–67. Retrieved from [http://journals.ametsoc.org/view/](http://journals.ametsoc.org/view/journals/mwre/147/1/mwr-d-18-0138.1.xml)
668 journals/mwre/147/1/mwr-d-18-0138.1.xml (Publisher: American Meteorolog-
669 ical Society Section: Monthly Weather Review) doi: 10.1175/MWR-D-18-0138.1
- 670 Courtier, P., & Talagrand, O. (1987). Variational assimilation of meteorological
671 observations with the adjoint vorticity equation. ii: Numerical results. *Quarterly*
672 *Journal of the Royal Meteorological Society*, 113(478), 1329–1347. Retrieved from
673 <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49711347813>
674 doi: <https://doi.org/10.1002/qj.49711347813>
- 675 Courtier, P., Thépaut, J. N., & Hollingsworth, A. (1994, July). A strategy for
676 operational implementation of 4D-Var, using an incremental approach. *Quar-*
677 *terly Journal of the Royal Meteorological Society*, 120(519), 1367–1387. doi:
678 10.1002/qj.49712051912
- 679 Elman, J. L. (1990). Finding structure in time. *Cognitive science*, 14(2), 179–211.
- 680 Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-
681 geostrophic model using monte carlo methods to forecast error statistics. *Jour-*
682 *nal of Geophysical Research: Oceans*, 99(C5), 10143–10162. Retrieved from
683 <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/94JC00572>
684 doi: <https://doi.org/10.1029/94JC00572>
- 685 Geist, K., Parlitz, U., & Lauterborn, W. (1990). Comparison of different meth-
686 ods for computing lyapunov exponents. *Progress of theoretical physics*, 83(5), 875–
687 893.
- 688 Ginsbourger, D., Le Riche, R., & Carraro, L. (2010). Kriging is well-suited to par-
689 allelize optimization. In *Computational intelligence in expensive optimization prob-*
690 *lems* (pp. 131–162). Springer.
- 691 Goodliff, M., Amezcua, J., & Leeuwen, P. J. V. (2017). A weak-constraint 4densem-
692 blevar. part ii: experiments with larger models. *Tellus A: Dynamic Meteorology*
693 *and Oceanography*, 69(1), 1271565. Retrieved from [https://doi.org/10.1080/](https://doi.org/10.1080/16000870.2016.1271565)
694 [16000870.2016.1271565](https://doi.org/10.1080/16000870.2016.1271565) doi: 10.1080/16000870.2016.1271565
- 695 Greybush, S. J., Kalnay, E., Miyoshi, T., Ide, K., & Hunt, B. R. (2011). Balance and
696 ensemble kalman filter localization techniques. *Monthly Weather Review*, 139(2),
697 511–522.
- 698 Griffith, A., Pomerance, A., & Gauthier, D. J. (2019). Forecasting chaotic systems
699 with very low connectivity reservoir computers. *Chaos: An Interdisciplinary Jour-*
700 *nal of Nonlinear Science*, 29(12), 123108.
- 701 Hatfield, S. E., Link to external site, t. l. w. o. i. a. n. w., Chantry, M., Link to ex-
702 ternal site, t. l. w. o. i. a. n. w., Dueben, P. D., Lopez, P., ... Link to external
703 site, t. l. w. o. i. a. n. w. (2021, February). Building tangent-linear and adjoint
704 models for data assimilation with neural networks. *Earth and Space Science Open*
705 *Archive ESSOAr*. Retrieved 2021-03-18, from [http://search.proquest.com/](http://search.proquest.com/docview/2493694654/abstract/51615480403B4848PQ/1)
706 [docview/2493694654/abstract/51615480403B4848PQ/1](http://search.proquest.com/docview/2493694654/abstract/51615480403B4848PQ/1) (Place: Washington,
707 United States Publisher: American Geophysical Union Section: Atmospheric Sci-
708 ences) doi: <http://dx.doi.org.ezproxy.lib.utexas.edu/10.1002/essoar.10506310.1>
- 709 Hunt, B. R., Kostelich, E. J., & Szunyogh, I. (2007). Efficient data assimilation for
710 spatiotemporal chaos: A local ensemble transform kalman filter. *Physica D: Non-*

- linear Phenomena, 230(1-2), 112–126.
- Jaeger, H. (2001). The “echo state” approach to analysing and training recurrent neural networks-with an erratum note. *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, 148(34), 13.
- Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4), 455–492.
- Kalnay, E., Hunt, B., Ott, E., & Szunyogh, I. (2006). Ensemble forecasting and data assimilation: two problems with the same solution. *Predictability of weather and climate*, 157, 180.
- Konkoli, Z. (2017). Reservoir computing. In R. A. Meyers (Ed.), *Encyclopedia of complexity and systems science* (pp. 1–12). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from https://doi.org/10.1007/978-3-642-27737-5_683-1 doi: 10.1007/978-3-642-27737-5_683-1
- Lin, H.-Y., & Penny, S. (2021-submitted). Fourier reservoir computing for data-driven prediction of multi-scale coupled quasi-geostrophic dynamics. *Journal of Advances in Modeling Earth Systems*, X(X), XX–XX.
- Lu, Z., Pathak, J., Hunt, B., Girvan, M., Brockett, R., & Ott, E. (2017). Reservoir observers: Model-free inference of unmeasured variables in chaotic systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(4), 041102.
- Lukoševičius, M. (2012). A practical guide to applying echo state networks. In G. Montavon, G. B. Orr, & K.-R. Müller (Eds.), *Neural networks: Tricks of the trade: Second edition* (pp. 659–686). Berlin, Heidelberg: Springer Berlin Heidelberg. Retrieved from https://doi.org/10.1007/978-3-642-35289-8_36 doi: 10.1007/978-3-642-35289-8_36
- Lukoševičius, M. (2012). A Practical Guide to Applying Echo State Networks. In G. Montavon, G. B. Orr, & K.-R. Müller (Eds.), *Neural Networks: Tricks of the Trade: Second Edition* (pp. 659–686). Berlin, Heidelberg: Springer. Retrieved 2021-01-04, from https://doi.org/10.1007/978-3-642-35289-8_36 doi: 10.1007/978-3-642-35289-8_36
- Lukoševičius, M., & Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3), 127–149. Retrieved from <https://www.sciencedirect.com/science/article/pii/S1574013709000173> doi: <https://doi.org/10.1016/j.cosrev.2009.03.005>
- Maass, W., Natschläger, T., & Markram, H. (2002, 11). Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*, 14(11), 2531–2560. Retrieved from <https://doi.org/10.1162/089976602760407955> doi: 10.1162/089976602760407955
- Moćkus, J. (1975). On bayesian methods for seeking the extremum. In G. I. Marchuk (Ed.), *Optimization techniques ifip technical conference novosibirsk, july 1–7, 1974* (pp. 400–404). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Oczkowski, M., Szunyogh, I., & Patil, D. (2005). Mechanisms for the development of locally low-dimensional atmospheric dynamics. *Journal of the atmospheric sciences*, 62(4), 1135–1156.
- Palatella, L., Carrassi, A., & Trevisan, A. (2013). Lyapunov vectors and assimilation in the unstable subspace: theory and applications. *Journal of Physics A: Mathematical and Theoretical*, 46(25), 254020.
- Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018). Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical review letters*, 120(2), 024102.
- Penny, S. G. (2014). The hybrid local ensemble transform kalman filter. *Monthly Weather Review*, 142(6), 2139–2149.
- Penny, S. G. (2017). Mathematical foundations of hybrid data assimilation from a synchronization perspective. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(12), 126801.

- 765 Platt, J. A., Wong, A., Clark, R., Penny, S. G., & Abarbanel, H. D. (2021). Fore-
766 casting using reservoir computing: The role of generalized synchronization. *arXiv*
767 *preprint arXiv:2102.08930*.
- 768 Ruckstuhl, Y., Janjić, T., & Rasp, S. (2021). Training a convolutional neural net-
769 work to conserve mass in data assimilation. *Nonlinear Processes in Geophysics*,
770 28(1), 111–119. Retrieved from [https://npg.copernicus.org/articles/28/](https://npg.copernicus.org/articles/28/111/2021/)
771 111/2021/ doi: 10.5194/npg-28-111-2021
- 772 Schrauwen, B., Verstraeten, D., & Campenhout, J. V. (2007). An overview of reser-
773 voir computing: theory, applications and implementations. In *Proceedings of the*
774 *15th european symposium on artificial neural networks* (pp. 471–482).
- 775 Steil, J. (2004). Backpropagation-decorrelation: online recurrent learning with o(n)
776 complexity. In *2004 ieee international joint conference on neural networks (ieee*
777 *cat. no.04ch37541)* (Vol. 2, p. 843–848 vol.2). doi: 10.1109/IJCNN.2004.1380039
- 778 Talagrand, O., & Courtier, P. (1987). Variational assimilation of meteorological ob-
779 servations with the adjoint vorticity equation. i: Theory. *Quarterly Journal of the*
780 *Royal Meteorological Society*, 113(478), 1311–1328. Retrieved from [https://rmets](https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49711347812)
781 [.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49711347812](https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49711347812) doi: [https://](https://doi.org/10.1002/qj.49711347812)
782 doi.org/10.1002/qj.49711347812
- 783 Trevisan, A., D’Isidoro, M., & Talagrand, O. (2010). Four-dimensional variational
784 assimilation in the unstable subspace and the optimal subspace dimension. *Quar-*
785 *terly Journal of the Royal Meteorological Society: A journal of the atmospheric*
786 *sciences, applied meteorology and physical oceanography*, 136(647), 487–496.
- 787 Trevisan, A., & Palatella, L. (2011a). On the kalman filter error covariance collapse
788 into the unstable subspace. *Nonlinear Processes in Geophysics*, 18(2), 243–250.
789 Retrieved from <https://npg.copernicus.org/articles/18/243/2011/> doi:
790 10.5194/npg-18-243-2011
- 791 Trevisan, A., & Palatella, L. (2011b). On the kalman filter error covariance collapse
792 into the unstable subspace. *Nonlinear Processes in Geophysics*, 18(2), 243–250.
- 793 Van der Vorst, H. A. (1992). Bi-cgstab: A fast and smoothly converging variant of
794 bi-cg for the solution of nonsymmetric linear systems. *SIAM Journal on scientific*
795 *and Statistical Computing*, 13(2), 631–644.
- 796 Vlachas, P. R., Pathak, J., Hunt, B. R., Sapsis, T. P., Girvan, M., Ott, E., &
797 Koumoutsakos, P. (2020). Backpropagation algorithms and reservoir comput-
798 ing in recurrent neural networks for the forecasting of complex spatiotemporal
799 dynamics. *Neural Networks*, 126, 191–217.