

Computational Science Laboratory Report
CSL-TR-21-7
November 17, 2021

Austin Chennault, Andrey A. Popov, Amit N. Subrahmanya,
Rachel Cooper, Anuj Karpatne, Adrian Sandu

*“Adjoint-Matching Neural Network Surrogates for Fast
4D-Var Data Assimilation”*

Computational Science Laboratory
“Compute the Future!”

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

Email: achennault@vt.edu, apopov@vt.edu, amitns@vt.edu, cguas@vt.edu,
karpatne@vt.edu, sandu@cs.vt.edu

Web: <http://csl.cs.vt.edu>



Adjoint-Matching Neural Network Surrogates for Fast 4D-Var Data Assimilation

Austin Chennault, *Member, IEEE*, Andrey A. Popov, Amit N. Subrahmanya, Rachel Cooper, Anuj Karpatne, Adrian Sandu

Abstract—The data assimilation procedures used in many operational numerical weather forecasting systems are based around variants of the 4D-Var algorithm. The cost of solving the 4D-Var problem is dominated by the cost of forward and adjoint evaluations of the physical model. This motivates their substitution by fast, approximate surrogate models. Neural networks offer a promising approach for the data-driven creation of surrogate models. The accuracy of the surrogate 4D-Var problem’s solution has been shown to depend explicitly on accurate modeling of the forward and adjoint for other surrogate modeling approaches and in the general nonlinear setting. We formulate and analyze several approaches to incorporating derivative information into the construction of neural network surrogates. The resulting networks are tested on out of training set data and in a sequential data assimilation setting on the Lorenz-63 system. Two methods demonstrate superior performance when compared with a surrogate network trained without adjoint information, showing the benefit of incorporating adjoint information into the training process.

Index Terms—Data assimilation, neural networks, optimization, meteorology, machine learning.

I. INTRODUCTION

MANY areas in science and engineering rely on complex computational models for the simulation of physical systems. A generic model has the form $\mathbf{x} = \mathcal{M}(\theta)$, depends on a set of parameters θ , and produces approximations of the physical quantities \mathbf{x} . The inverse problem consists of using noisy measurements of the physical quantities \mathbf{x} , together with the model operator \mathcal{M} , to obtain improved estimates of the parameter value θ . The model operator of interest and its corresponding adjoint may be expensive to evaluate. An example from operational numerical weather prediction platforms is a nonlinear transformation of a trajectory of partial differential equation solutions computed by finite element discretization at upwards of 10^9 mesh points [10], [35]. Solution of an inverse problem may require thousands of forward model evaluations in the statistical setting or several hundred in the variational setting [15], [30], with solution cost depending primarily on the expense of the forward model \mathcal{M} . Inexpensive surrogate models can then be used in place of the high fidelity model operator \mathcal{M} for fast, approximate inversion. In this paper we focus on data assimilation, i.e., inverse problems

where the underlying models \mathcal{M} are dynamical systems that evolve in time. Surrogate models have a long history within data assimilation. Previously studied approaches have included proper orthogonal decomposition (POD) and variants [34], and various neural network approaches [13], [20], [37].

This work proposes the construction of specialized neural-network based surrogates for accelerating the four dimensional variational (4D-Var) solution of data assimilation problems. The 4D-Var approach calculates a maximum a posteriori estimate of the model parameters θ by solving a constrained optimization problem; the cost function includes the prior information and the mismatch between model predictions and observations, and the constraints are the high fidelity model equations. Our proposed approach is to replace the high fidelity model constraints with the surrogate model equations, thereby considerably reducing the cost of solving the optimization problem. As shown in [24], [34], the quality of the reduced order 4D-Var problem’s solution depends not only on the accuracy of the surrogate model on forward dynamics, but also on the accurate modeling of the adjoint dynamics.

As the model (\mathcal{M}) encapsulates what one knows about the physics of the process at hand, the exact model derivative ($d\mathcal{M}/d\theta$) can also be interpreted as known physical information. We incorporate this information into the neural network surrogate training by appending an adjoint mismatch term to the loss function. Training neural networks using loss functions incorporating the derivative of the network itself has been applied for the solution of partial differential equations [23], but, to our knowledge, this is the first work to use derivative information in the solution of a variational inverse problem.

We formulate several training methods that incorporate derivative information, and demonstrate that the resulting surrogates have superior generalization performance over the traditional approach where training uses only forward model information.

The remainder of the paper is organized as follows. Section II introduces the 4D-Var problem, solution strategies, and the use of surrogates in its solution. Several forms of neural networks and the solution of the training problem are discussed. Section III provides a short theoretical analysis of the solution of 4D-Var with surrogate models, and show the dependence of the 4D-Var solution quality on the accuracy of the surrogate model and its adjoint. Section IV introduces the science-guided machine learning framework, and formulates the science-guided approach to the 4D-Var problem. Numerical results are given in Section V and closing remarks in Section VI.

A. Chennault, A. A. Popov, Amit N. Subrahmanya, and Adrian Sandu are with the Computational Science Laboratory in the Computer Science department at Virginia Tech.

R. Cooper is with MITRE.

A. Karpatne is with the Computer Science department at Virginia Tech.

Manuscript received ...; revised ...

II. BACKGROUND

A. Data Assimilation

Data assimilation [3], [11], [25] is the process of combining imperfect forecasts of a dynamical system from a physics-based model with noisy, typically sparse observations to obtain an improved estimate of the system state. The data assimilation problem is generally solved with statistical or variational approaches [3]. A data assimilation algorithm combines information from a background estimate of the system state and observation information to obtain a more accurate estimate of the system state, known as the analysis. The data assimilation problem is fundamentally Bayesian in nature. The problem assumes a likelihood distribution of the observations given a system state and prior distribution of possible background states. The data assimilation algorithm then aims to produce a sample from the posterior distribution of system states (also called the analysis [3], [29]).

The two main settings for the data assimilation problem are filtering and smoothing. In the filtering setting, observation information from the current time t_0 is used to improve an estimate of the system's current state at time t_0 . In the smoothing setting, current and future observation information from times t_0, \dots, t_n is used to produce an improved estimate of the system state at time t_0 .

We now outline the formal setting for the smoothing problem. Consider a finite dimensional representation \mathbf{x}_i of the state of some physical system at time $t_i, i \geq 0$, a background (prior) estimate \mathbf{x}_0^b of the true system state $\mathbf{x}_0^{\text{true}}$ state at time t_0 , noisy observations (measurements) $\mathbf{y}_i = \mathcal{H}(\mathbf{x}_i)$ of the state at times $t_i, i \geq 0$, the covariance matrices $\mathbf{B}_0, \mathbf{R}_i$, and \mathbf{Q}_i associated with each noise variable, and a high fidelity model operator $\mathcal{M}_{i-1,i}(\mathbf{x})$ which transforms the system state at time t_{i-1} to one at time t_i :

$$\begin{aligned} \mathbf{x}_0 &= \mathbf{x}_0^b = \mathbf{x}_0^{\text{true}} + \eta_0, & \eta_0 &\sim \mathcal{N}(0, \mathbf{B}_0), \\ \mathbf{x}_i &= \mathcal{M}_{i-1,i}(\mathbf{x}_{i-1}) + \eta_i, & \eta_i &\sim \mathcal{N}(0, \mathbf{Q}_i), \\ \mathbf{y}_i &= \mathcal{H}_i(\mathbf{x}_i^{\text{true}}) + \varepsilon_i^{\text{obs}}, & \varepsilon_i^{\text{obs}} &\sim \mathcal{N}(0, \mathbf{R}_i), \\ i &= 1, \dots, n. \end{aligned} \quad (1)$$

This information will be combined in some optimal setting to produce an improved estimate \mathbf{x}_0^a of the true system state $\mathbf{x}_0^{\text{true}}$ at time t_0 .

B. 4D-Var

A solution to the variational data assimilation problem is defined as the numerical solution to some appropriately formulated optimization problem [3]. 4D-Var can be thought of as a variational approach to the smoothing problem. The 4D-Var methodology is an example of a variational inverse problem. In this context, the variable of interest is the state of a dynamical system. We additionally have the notion of successive transformations of the state provided through the model operator and time-distributed observations. The goal of the 4D-Var is to find an initial value for our dynamical system that shadows the true trajectory by weakly matching sparse, noisy observations and prior information about the system's state [16].

C. 4D-Var Problem Formulation

4D-Var seeks the maximum a posteriori estimate of the state \mathbf{x}_0^a at t_0 subject to the constraints imposed by the high fidelity model equations:

$$\begin{aligned} \mathbf{x}_0^a &= \arg \min \Psi(\mathbf{x}_0) \\ &\text{subject to } \mathbf{x}_i = \mathcal{M}_{i-1,i}(\mathbf{x}_{i-1}), \quad i = 1, \dots, n, \\ \Psi(\mathbf{x}_0) &:= \frac{1}{2} \|\mathbf{x}_0 - \mathbf{x}_0^b\|_{\mathbf{B}_0^{-1}}^2 + \frac{1}{2} \sum_{i=1}^n \|\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i\|_{\mathbf{R}_i^{-1}}^2, \end{aligned} \quad (2)$$

where $\|\mathbf{x}\|_{\mathbf{A}} := \sqrt{\mathbf{x}^T \mathbf{A} \mathbf{x}}$, with \mathbf{A} a symmetric and positive-definite matrix [30].

D. Solving the 4D-Var Optimization Problem

In practice, the constrained minimization problem (2) is solved with gradient-based optimization algorithms. Denote the Jacobian of the model solution operator with respect to model state (called the ‘‘tangent linear model’’), and the Jacobian of the observation operator, by:

$$\mathbf{M}_{i,i+1}(\mathbf{x}_i) := \left. \frac{\partial \mathcal{M}_{i,i+1}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_i} \in \mathbb{R}^{\text{Nstate} \times \text{Nstate}}, \quad (3)$$

$$\mathbf{H}_i(\mathbf{x}_i) := \left. \frac{\partial \mathcal{H}_i(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_i} \in \mathbb{R}^{\text{Nobs} \times \text{Nstate}}, \quad (4)$$

respectively. The gradient of the 4D-Var cost function (2) takes the form:

$$\begin{aligned} \nabla_{\mathbf{x}_0} \Psi(\mathbf{x}_0) &= \mathbf{B}_0^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \\ &\sum_{i=1}^n \left(\prod_{k=1}^i \mathbf{M}_{i-k,i-k+1}^T(\mathbf{x}_{i-k}) \right) \mathbf{H}_i^T \mathbf{R}_i^{-1} (\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i). \end{aligned} \quad (5)$$

This motivates the need for an efficient evaluation of the ‘‘adjoint model’’ $\mathbf{M}_{k,k+1}^T$. For our purposes the adjoint model is the transpose of the tangent linear model, although the method of adjoints is typically derived in a more general setting [3].

The cost of solving the 4D-Var problem (2) is dominated by the computational cost of evaluating Ψ and its gradient (5). Evaluating Ψ requires integrating the model \mathcal{M} forward in time. Evaluating $\nabla_{\mathbf{x}_0} \Psi(\mathbf{x}_0)$ requires evaluating Ψ and running the adjoint model \mathbf{M}^T backward in time, both runs being performed over the entire simulation interval $[t_0, t_n]$. Techniques for efficient computation of the sum in (5) are based on computing only adjoint model-times-vector operations [28], [31], [38]. Each full gradient computation requires a single adjoint run over the simulation interval.

E. 4D-Var with Surrogate Models

Surrogate models for fast, approximate inference have enjoyed great popularity in data assimilation research [5], [21], [22], [32]–[34]. Surrogate models are most often applied in one of two ways: to replace the high fidelity model dynamics constraints in (2), or to supplement the model by estimating the model error term η_i in (1) [13]. In the former approach the derivatives of the surrogate model replace the derivatives of the high fidelity model in (3) and in the 4D-Var gradient calculation in (5). The universal approximation property

of neural networks and their relatively inexpensive forward and derivative evaluations make them a promising candidate architecture for surrogate construction [1], [13], [20].

F. Neural Networks

Neural networks are parametrized function approximation architectures loosely inspired by the biological structure of brains [1]. The canonical example of an artificial neural network is the feedforward network, or multilayer perceptron. Given an element-wise nonlinear activation function ϕ , weight matrices $\{\mathbf{W}_i\}_{i=1,\dots,k}$ and bias vectors $\{\mathbf{b}_i\}_{i=1,\dots,k-1}$, the action \mathbf{u}_{out} of a feedforward network on an input vector \mathbf{u}_{in} is given by the sequence of operations:

$$\begin{aligned} \mathbf{z}_1 &= \phi(\mathbf{W}_1 \mathbf{u}_{\text{in}} + \mathbf{b}_1), \\ \mathbf{z}_i &= \phi(\mathbf{W}_i \mathbf{z}_{i-1} + \mathbf{b}_i), \quad i = 2, \dots, k-1, \\ \mathbf{u}_{\text{out}} &= \mathbf{W}_k \mathbf{z}_{k-1} + \mathbf{b}_k. \end{aligned} \quad (6)$$

Popular choices of activation function include hyperbolic tangent and the rectified linear unit (ReLU) [1]. In general, the dimension of each layer's output may vary. The final bias vector may be omitted to suggest more clearly the idea of linear regression on a learned nonlinear transformation of the data, but we retain it in our implementation. Let $\mathbf{z}_i, \mathbf{b}_i \in \mathbb{R}^{N_i}$ for $i = 1, \dots, k-1$, $\mathbf{b}_k \in \mathbb{R}^{N_k}$ and $\mathbf{W}_i \in \mathbb{R}^{N_i \times N_{i-1}}$, $i = 2, \dots, k$. The input dimension $\mathbf{u}_{\text{in}} \in \mathbb{R}^{N_{\text{in}}}$ and the output dimension $\mathbf{u}_{\text{out}} \in \mathbb{R}^{N_{\text{out}}}$ are specified by the problem at hand.

G. Training Neural Networks

In the supervised setting, networks are trained on loss functions which are additive on a provided set of input/output pairs $\{(\mathbf{u}_{\text{in}}^\ell, \mathbf{u}_{\text{out}}^\ell)\}_{\ell=1,\dots,N_{\text{data}}}$, called the training data. The output and optionally the input may be replaced by a sequence in the typical RNN setting. The superscript here is the index variable and does not denote exponentiation. The structure of neural networks and additive loss function allows use of the backpropagation algorithm for efficient gradient computation [1]. In addition, it provides motivation for the typical method of training neural networks, stochastic gradient descent, where an estimate of the loss function across the entire dataset and its gradient with respect to the network weights is computed using only a sample of training data. Stochastic gradient descent is often combined with an acceleration method such as Adam, which makes use of estimates of the stochastic gradient's first and second order moments to accelerate convergence [17]. Constraints can be enforced in the neural network training process through incorporation of a penalty term or the more elaborate process of modifying a traditional constrained optimization algorithm for the stochastic gradient descent setting [7].

H. Neural Networks in Numerical Weather Prediction

Approximate models based on neural networks in the context of numerical weather prediction have been explored. Previously studied approaches include learning of the model operator by feedforward networks [9] and learning repeated

application of the model operator by recurrent neural network for use in 4D-Var [20], approaches based on learning the results of the data assimilation process [2], and online and offline approaches for model error correction [12]. While to the authors' knowledge, deep-learning based models have failed to match operational weather-prediction models in prediction skill [9], [12], [36] they nonetheless have offered comparatively inexpensive approximations of operational operators which can be applied to specific areas of interest and have thus been identified as an area of application for machine learning by organizations such as the European Centre for Medium-Range Weather Forecasts [14].

III. THEORETICAL MOTIVATION

It has been shown in [24], [32], [34] that, when the optimization (2) is performed with a surrogate model \mathcal{N} in place of the inner loop model operator \mathcal{M} , the accuracy of the resulting solution depends on accuracy of the forward model as well as its adjoint. A rigorous derivation of error estimates in the resulting 4D-Var solution upon the forward and adjoint surrogate solutions has developed in [24]. We will follow simplified analysis in the same setting to demonstrate the requirement for accurate adjoint model dynamics. We note that the theory based on first-order necessary conditions predicts an increase in solution accuracy regardless of the optimization algorithm used to solve (2).

The constrained optimization problem (2) can be solved analytically by the method of Lagrange multipliers. The Lagrangian for (2) is:

$$\begin{aligned} \mathcal{L}(\mathbf{x}_0; \boldsymbol{\lambda}) &= \frac{1}{2} \|\mathbf{x}_0^{\text{b}} - \mathbf{x}_0\|_{\mathbf{B}_0^{-1}}^2 + \frac{1}{2} \sum_{i=1}^n \|\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i\|_{\mathbf{R}_i^{-1}}^2 \\ &\quad + \sum_{i=0}^{n-1} \boldsymbol{\lambda}_{i+1}^T (\mathbf{x}_{i+1} - \mathcal{M}_{i,i+1}(\mathbf{x}_i)), \end{aligned} \quad (7)$$

where $\boldsymbol{\lambda}$ is the vector containing all Lagrange multipliers $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \dots, \boldsymbol{\lambda}_n$. A local optimum to the constrained optimization problem (2) is stationary point of (7) and fulfills the first order optimality necessary conditions:

High fidelity forward model:

$$\begin{aligned} \mathbf{x}_{i+1} &= \mathcal{M}_{i,i+1}(\mathbf{x}_i), \quad i = 0, \dots, n-1; \\ \mathbf{x}_0 &= \mathbf{x}_0^{\text{a}}; \end{aligned} \quad (8\text{a})$$

High fidelity adjoint model:

$$\begin{aligned} \boldsymbol{\lambda}_n &= \mathbf{H}_n^T \mathbf{R}_n^{-1} (\mathbf{y}_n - \mathcal{H}_n(\mathbf{x}_n)); \\ \boldsymbol{\lambda}_i &= \mathbf{M}_{i+1,i}^T \boldsymbol{\lambda}_{i+1} + \mathbf{H}_i^T \mathbf{R}_i^{-1} (\mathbf{y}_i - \mathcal{H}_i(\mathbf{x}_i)), \\ i &= n-1, \dots, 0; \end{aligned} \quad (8\text{b})$$

High fidelity gradient:

$$\nabla_{\mathbf{x}_0} \Psi(\mathbf{x}_0^{\text{a}}) = -\mathbf{B}_0^{-1} (\mathbf{x}_0^{\text{b}} - \mathbf{x}_0^{\text{a}}) - \boldsymbol{\lambda}_0 = 0. \quad (8\text{c})$$

Suppose we have some differentiable approximation \mathcal{N} to our high fidelity model operator \mathcal{M} . We have $\mathcal{N}_{i,i+1}(\mathbf{x}) = \mathcal{M}_{i,i+1}(\mathbf{x}) + \mathbf{e}_{i,i+1}(\mathbf{x})$, where the surrogate model approximation error is $\mathbf{e}_{i,i+1}(\mathbf{x}) = \mathcal{N}_{i,i+1}(\mathbf{x}) - \mathcal{M}_{i,i+1}(\mathbf{x})$. Replacing

\mathcal{M} by $\mathcal{N} = \mathcal{M} + \mathbf{e}$ leads to the perturbed 4D-Var problem:

$$\begin{aligned} \mathbf{x}_0^{a*} &= \arg \min \Psi(\mathbf{x}_0) \\ \text{subject to } \mathbf{x}_i &= \mathcal{M}_{i,i+1}(\mathbf{x}_i) + \mathbf{e}_{i,i+1}(\mathbf{x}_i), \\ & i = 0, \dots, n-1, \end{aligned} \quad (9)$$

and its corresponding optimality conditions:

Surrogate forward model:

$$\begin{aligned} \mathbf{x}_{i+1}^* &= \mathcal{M}_{i,i+1}(\mathbf{x}_i^*) + \mathbf{e}_{i,i+1}(\mathbf{x}_i^*), \\ i = 0, \dots, n-1, \quad \mathbf{x}_0^* &= \mathbf{x}_0^{a*}; \end{aligned} \quad (10a)$$

Surrogate adjoint model:

$$\begin{aligned} \boldsymbol{\lambda}_n^* &= \mathbf{H}_n^T \mathbf{R}_n^{-1}(\mathbf{y}_n - \mathcal{H}_n(\mathbf{x}_n^*)), \\ \boldsymbol{\lambda}_i^* &= (\mathbf{M}_{i,i+1} + \mathbf{e}'_{i,i+1}(\mathbf{x}_i^*))^T \boldsymbol{\lambda}_{i+1}^* \\ &+ \mathbf{H}_i^T \mathbf{R}_i^{-1}(\mathbf{y}_i - \mathcal{H}_i(\mathbf{x}_i^*)), \\ i = n-1, \dots, 0; \end{aligned} \quad (10b)$$

Surrogate gradient:

$$\nabla_{\mathbf{x}_0} \Psi(\mathbf{x}_0) = -\mathbf{B}_0^{-1}(\mathbf{x}_0^b - \mathbf{x}_0^{a*}) - \boldsymbol{\lambda}_0^* = 0. \quad (10c)$$

Let \mathbf{x}_0^{a*} be a solution to the reduced 4D-Var problem (9) satisfying its corresponding optimality conditions and \mathbf{x}^a a solution to the full 4D-Var problem (2). We can examine the quality of \mathbf{x}_0^{a*} as a solution to the original 4D-Var problem (2) by examining the difference of two solutions satisfying their respective first-order optimality conditions:

Forward model error:

$$\begin{aligned} \mathbf{x}_1^* - \mathbf{x}_1 &= \mathcal{M}_{0,1}(\mathbf{x}_0^{a*}) - \mathcal{M}_{0,1}(\mathbf{x}_0^a) \\ &+ \mathbf{e}_{0,1}(\mathbf{x}_0^{a*}), \\ \mathbf{x}_i^* - \mathbf{x}_i &= \mathcal{M}_{i-1,i}(\mathbf{x}_{i-1}^*) - \mathcal{M}_{i-1,i}(\mathbf{x}_{i-1}) \\ &+ \mathbf{e}_{i-1,i}(\mathbf{x}_{i-1}^*), \\ i = 2, \dots, n. \end{aligned} \quad (11a)$$

Adjoint model error:

$$\begin{aligned} \boldsymbol{\lambda}_n^* - \boldsymbol{\lambda}_n &= \mathbf{H}_n^T \mathbf{R}_n^{-1}(\mathcal{H}_n(\mathbf{x}_n) - \mathcal{H}_n(\mathbf{x}_n^*)), \\ \boldsymbol{\lambda}_i^* - \boldsymbol{\lambda}_i &= \mathbf{M}_{i+1,i}^T(\boldsymbol{\lambda}_{i+1}^* - \boldsymbol{\lambda}_{i+1}) \\ &+ \mathbf{H}_i^T \mathbf{R}_i^{-1}(\mathcal{H}_i(\mathbf{x}_i) - \mathcal{H}_i(\mathbf{x}_i^*)) \\ &+ \mathbf{e}'_{i,i+1}(\mathbf{x}_i^*)^T \boldsymbol{\lambda}_{i+1}^*, \\ i = n-1, \dots, 0. \end{aligned} \quad (11b)$$

Solution error:

$$\mathbf{x}_0^{a*} - \mathbf{x}_0^a = \mathbf{B}_0(\boldsymbol{\lambda}_0^* - \boldsymbol{\lambda}_0). \quad (11c)$$

We see that the additional error at each step of the forward evaluation depends on the difference of the model operator applied to \mathbf{x}_i , the accumulated error, and the mismatch function \mathbf{e} itself. The error in the adjoint variables depends on the forward error as transformed by the observation operator and the adjoint of the mismatch function applied to specific vectors, namely the perturbed adjoint variables $\boldsymbol{\lambda}_i^*$. Since the error of the solution the mismatch in final adjoint variables, we see that the quality of the solution depends directly on both the mismatch function \mathbf{e} and its derivative.

IV. SCIENCE-GUIDED ML FRAMEWORK

A. Traditional Neural Network Approach

Neural networks offer one approach to building computationally inexpensive surrogate models for the model solution operator \mathcal{M} in eq. (2), which can be used in place of the high fidelity model in the inversion procedure [20], [37]. The neural network surrogate \mathcal{N} takes as input the system state \mathbf{x}_i and outputs an approximation of the state advanced by a traditional time stepping method, i.e. $\mathcal{N}(\mathbf{x}_i; \theta) \approx \mathcal{M}_{i,i+1}(\mathbf{x}_i)$, where θ denote the parameters of the neural network. In our methodology, the surrogate \mathcal{N} is fixed does not vary between time steps, therefore we do not use time subscripts. Moreover, the same set of parameters θ is used for all time steps i .

The standard approach is to train the surrogate model \mathcal{N} on input-output pairs $\{(\mathbf{x}_{t_i}, \mathcal{M}_{t_i,t_{i+1}}(\mathbf{x}_{t_i}))\}_{i=1}^{N_{\text{data}}}$ resulting from the full scientific model. We have replaced the subscript i priorly used with a double subscript t_i to indicate that the data used in the training process need not consist of all snapshots collected from a single model trajectory. Training is accomplished by minimizing squared two-norm mismatch summed across the training data set:

$$\mathcal{L}_{\text{Standard}}(\theta) := \sum_{i=1}^{N_{\text{data}}} \|\mathcal{N}(\mathbf{x}_{t_i}; \theta) - \mathcal{M}_{t_i,t_{i+1}}(\mathbf{x}_{t_i})\|_2^2. \quad (12)$$

B. Adjoint-Match Training

Incorporating known physical quantities into the training of a neural network by an additional term in the loss function is one of the basic methods of increasing model performance on complex scientific data [37]. In the variational data assimilation context, we assume access to the high fidelity model operator's adjoint. Our goal is to devise a training method that incorporates this adjoint information and produces more accurate solutions to the 4D-Var problem when the trained network is used as a surrogate. To this end, we denote the Jacobian of the neural network surrogate by

$$\mathbf{N}(\mathbf{x}_i; \theta) := \left. \frac{\partial \mathcal{N}(\mathbf{x}; \theta)}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_i} \in \mathbb{R}^{N_{\text{state}} \times N_{\text{state}}}. \quad (13)$$

A natural means of incorporating adjoint information into the training process is via the cost function

$$\begin{aligned} \mathcal{L}_{\text{Adj}}(\theta) &:= \mathcal{L}_{\text{Standard}}(\theta) \\ &+ \lambda \sum_{i=1}^{N_{\text{data}}} \|\mathbf{N}^T(\mathbf{x}_{t_i}; \theta) - \mathbf{M}_{t_i,t_{i+1}}^T(\mathbf{x}_{t_i})\|_F^2 \\ &\equiv \mathcal{L}_{\text{Standard}}(\theta) \\ &+ \lambda \sum_{i=1}^{N_{\text{data}}} \|\mathbf{N}(\mathbf{x}_{t_i}; \theta) - \mathbf{M}_{t_i,t_{i+1}}(\mathbf{x}_{t_i})\|_F^2, \end{aligned} \quad (14)$$

where the regularization parameter λ determines how heavily the adjoint mismatch term is weighted. Training neural networks using loss functions incorporating the derivative of the network itself has been applied for approximate solution of partial differential equations [23], but, to our knowledge, this paper represents the first attempt to use of this information in the solution of a variational inverse problem. In addition to

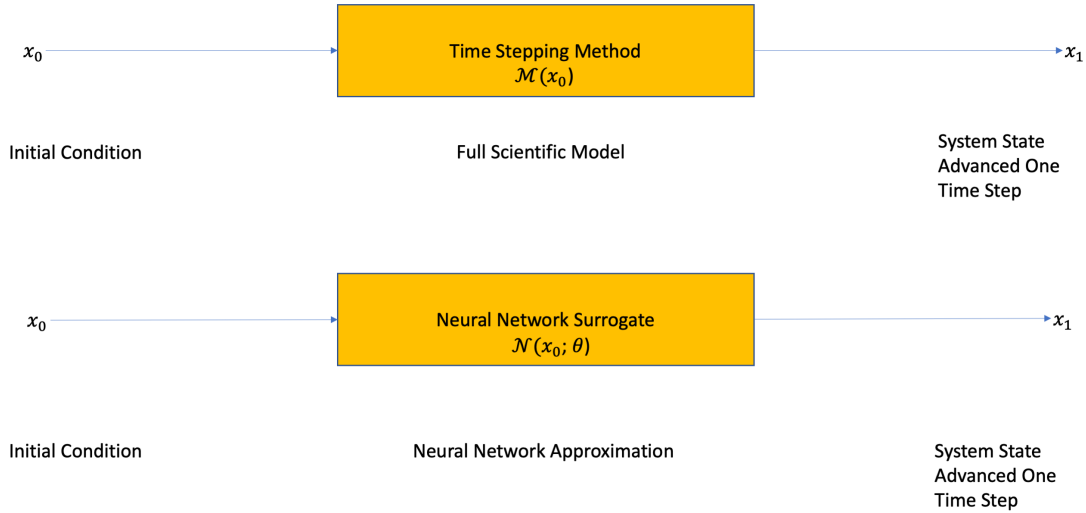


Fig. 1. Diagram of the neural network surrogate approach. A time stepping method advances the system state according to a scientifically derived formula. The neural network learns to provide a cheap imitation of the timestepping method through the training process.

the motivation for incorporating adjoint information into \mathcal{N} derived above, we may hope that in ensuring that \mathcal{N} linearly responds to perturbations in the input similarly to \mathcal{M} , we will obtain improved performance when the network is tested on out of training set data.

We note that the term $\|\mathbf{N}^T(\mathbf{x}_{t_i}; \theta) - \mathbf{M}_{t_i, t_{i+1}}^T(\mathbf{x}_{t_i})\|_F^2$ in (14) is exactly equal to $\|\mathbf{N}(\mathbf{x}_{t_i}; \theta) - \mathbf{M}_{t_i, t_{i+1}}(\mathbf{x}_{t_i})\|_F^2$, and that computation of the model adjoint $\mathbf{M}_{t_i, t_{i+1}}^T(\mathbf{x}_{t_i})$ requires considerably more computational complexity than computation of the model tangent linear model (TLM) $\mathbf{M}_{t_i, t_{i+1}}(\mathbf{x}_{t_i})$ due in part to the additional memory requirements [27]. In the sequel we will explore training with adjoint-vector products rather than the full adjoint matrix itself. Accurate adjoint-vector products at the solution (along with accuracy of the forward model) is the sufficient condition for accurate solution to the surrogate-optimized problem derived in [24]. Therefore we retain the formulation (14). Although we do not consider the time of the training process in this paper, we note that evaluation of (14) requires evaluation of the neural network's adjoint, and should thus increase the computational expense of the training problem. We note that the cost evaluating the network and its derivative after the training process is not affected by the choice of training cost function.

C. Training with Adjoint-Vector Products

In operational settings adjoint-vector products $\mathbf{M}^T \mathbf{v}$ are more readily available than high fidelity model adjoint operator. Using the figure from the Introduction section I, one dense Jacobian matrix of the full model operator represented in double precision floating point format would require $10^9 \cdot 10^9 \cdot 64$ bits, or 8 exabytes of memory to store. Instead of training our network in terms using adjoint operator mismatch, the network can be trained to match adjoint-vector products of the physical

model on a given set of vectors $\{\mathbf{v}_i\}_{i=1}^{N_{\text{data}}}$, which results in the following loss function:

$$\mathcal{L}_{\text{AdjVec}}(\theta) := \mathcal{L}_{\text{Standard}}(\theta) + \lambda \sum_{i=1}^{N_{\text{data}}} \|\mathbf{N}(\mathbf{x}_{t_i}; \theta)^T \mathbf{v}_i - \mathbf{M}_{t_i, t_{i+1}}^T \mathbf{v}_i\|_2^2. \quad (15)$$

In operational settings, the most readily available adjoint-vector products will be those calculated in computation of the term

$$\sum_{i=1}^n \left(\prod_{k=1}^i \mathbf{M}_{i-k, i-k+1}^T \right) \mathbf{H}_i^T \mathbf{R}_i^{-1} (\mathcal{H}(\mathbf{x}_i) - \mathbf{y}_i)$$

from the 4D-Var gradient (5).

We note that like the adjoint training procedure above, evaluation of (15) requires evaluation of neural network adjoint-vector products, and thus should increase computational expense of the training procedure over the standard network. The cost of neural network and neural network adjoint evaluation after the training process is not affected.

D. Independent Forward/Adjoint Surrogate Training

Another machine learning-based approach to model reduction is the construction of separate models of the forward and adjoint dynamics. Use of a regularized cost function as in our original formulation (14) results in an inherent trade-off between minimizing the data mismatch term and the weighted penalty term. In principle separate models may allow a more accurate modeling of the forward and adjoint dynamics. We construct two neural networks: one, $\mathcal{N}_{\text{IndepFwd}}(\mathbf{x}; \theta_{\text{Fwd}})$, modeling only the forward dynamics and trained with loss

function:

$$\mathcal{L}_{\text{IndepFwd}}(\theta) := \sum_{i=1}^{N_{\text{data}}} \|\mathcal{N}_{\text{IndepFwd}}(\mathbf{x}_{t_i}; \theta) - \mathcal{M}_{t_i, t_i+1}(\mathbf{x}_{t_i})\|_2^2 \quad (16)$$

and another model, $\mathcal{N}_{\text{IndepAdj}}(\mathbf{x}; \theta_{\text{Adj}})$ which mimics the adjoint of the scientific model with respect to the system state. Although the adjoint $\mathbf{M}_{t_i}^T(\mathbf{x}_{t_i})$ is itself a linear operator, the function $\mathbf{x}_{t_i} \mapsto \mathbf{M}_{t_i}^T(\mathbf{x}_{t_i})$ in general is not, justifying use of a nonlinear approximation. The network is trained with loss function:

$$\mathcal{L}_{\text{IndepAdj}}(\theta) := \sum_{i=1}^{N_{\text{data}}} \|\mathcal{N}_{\text{IndepAdj}}(\mathbf{x}_{t_i}; \theta) - \mathbf{M}_{t_i, t_i+1}^T(\mathbf{x}_{t_i})\|_F^2. \quad (17)$$

The input to the trained network $\mathcal{N}_{\text{IndepAdj}}$ is a state vector \mathbf{x}_{t_i} and its output a $N_{\text{state}} \times N_{\text{state}}$ matrix which approximates $\mathbf{M}_{t_i, t_i+1}^T(\mathbf{x}_{t_i})$. Although not explored in our paper, the adjoint network could also be constructed by training on the loss function formulated in terms of adjoint-vector products:

$$\mathcal{L}_{\text{IndepAdjVec}}(\theta) := \sum_{i=1}^{N_{\text{data}}} \|\mathcal{N}_{\text{IndepAdj}}(\mathbf{x}_{t_i}; \theta) \mathbf{v}_i - \mathbf{M}_{t_i, t_i+1}^T(\mathbf{x}_{t_i}) \mathbf{v}_i\|_2^2. \quad (18)$$

V. NUMERICAL EXPERIMENTS

A. Lorenz-63 System and 4D-Var Problem

In [18], Edward N. Lorenz introduced the first formally identified continuous chaotic dynamical system and common data assimilation test problem

$$\begin{aligned} x' &= \sigma(y - x), \\ y' &= x(\rho - z) - y, \\ z' &= xy - \beta z, \end{aligned} \quad (19)$$

with $\sigma = 10, \rho = 28, \beta = 8/3$.

We use (19) as our test problem as implemented in the ODE Test Problems software package [6], [26]. Time integration is accomplished with fourth order explicit Runge Kutta, over integration intervals $\Delta T = 0.12$ (model time units) with a fixed step size $\Delta t = \Delta T/50$. Adjointns are obtained from the corresponding discrete adjoint method [27].

In the 4D-Var problem, we observe all variables and use observation noise covariance matrix $\mathbf{R} = \mathbf{I}_3$, the 3×3 identity matrix. We use a scaled version of the estimated climatological covariance matrix

$$\mathbf{B}_0 = \frac{1}{5} \begin{bmatrix} 62.1471 & 62.1614 & -1.0693 \\ 62.1614 & 80.4184 & -0.2497 \\ -1.0693 & -0.2497 & 73.8169 \end{bmatrix}. \quad (20)$$

The 4D-Var window is set to $n = 2$. The 4D-Var cost function (2) looks forward two intervals of length ΔT .

All models are used to solve the data assimilation problem sequentially, where the analysis at time t_i is used as the background state for time $t_{i+1} = t_i + \Delta T$. Runs of length $N_{\text{time}} = 550$ are computed using by solving the 4D-Var problem with analysis window $n = 2$ using each model to obtain an analysis estimate. Analysis propagation for each

method is done by the corresponding model itself, meaning that a solution \mathbf{x}^a to the 4D-Var optimization procedure at step t_i solved with surrogate model \mathcal{N} in place of \mathcal{M} sets $\mathbf{x}^b = \mathcal{N}(\mathbf{x}^a)$ at step t_{i+1} . These runs of 550 steps of length ΔT (with the first 50 steps discarded in error calculations) are used for all 4D-Var solution accuracy tests. This problem setting is indendent of the data used in the neural network training process, described in Section V-C.

The 4D-Var optimization problem (2) is solved using BFGS and strong Wolfe condition line search [8], [19]. Gradients for the 4D-Var problem are computed using the integration method's discrete adjoint [27]. The data assimilation setting is the same for each trained model. Each trained model is used in place of the time integration procedure as illustrated in Figure 1. `Exact` refers to the solution of the 4D-Var problem (2) solved with the time integration and discrete adjoints described above. Results for `Adj`, `AdjVec`, `Standard`, and `Indep` are computed using surrogate 4D-Var and gradient equations where where the full model operator \mathcal{M} in (2) has been replaced by \mathcal{N} and the full model adjoint \mathbf{M}^T by \mathbf{N}^T when its gradient is computed. Explicitly, the surrogate 4D-Var problem and gradient are given by

$$\begin{aligned} \mathbf{x}_0^a &= \arg \min \Psi(\mathbf{x}_0) \\ \text{subject to } \mathbf{x}_i &= \mathcal{N}(\mathbf{x}_{i-1}), \quad i = 1, \dots, n, \\ \Psi(\mathbf{x}_0) &:= \frac{1}{2} \|\mathbf{x}_0 - \mathbf{x}_0^b\|_{\mathbf{B}_0^{-1}}^2 + \frac{1}{2} \sum_{i=1}^n \|\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i\|_{\mathbf{R}_i^{-1}}^2, \end{aligned} \quad (21)$$

and

$$\begin{aligned} \nabla_{\mathbf{x}_0} \Psi(\mathbf{x}_0) &= \mathbf{B}_0^{-1} (\mathbf{x}_0 - \mathbf{x}_0^b) + \\ &\sum_{i=1}^n \left(\prod_{k=1}^i \mathbf{N}_{i-k, i-k+1}^T(\mathbf{x}_{i-k}) \right) \mathbf{H}_i^T \mathbf{R}_i^{-1} (\mathcal{H}_i(\mathbf{x}_i) - \mathbf{y}_i). \end{aligned} \quad (22)$$

\mathbf{N}^T in (22) is the adjoint of the neural network itself for `Adj`, `AdjVec`, and `Standard`, and is provided by the separate adjoint network for `Indep`. The trained surrogate model does not vary between timesteps and only depends on the input \mathbf{u}_{in} . This is our primary validation of the trained model's accuracy, since it is constructed specifically to provide cheap solutions to the 4D-Var problem.

B. Network Architectures and Training

All neural networks except the `Indep`'s $\mathcal{N}_{\text{IndepAdj}}$ model (17) are two layer hidden feedforward networks with fixed hidden dimension and hyperbolic tangent activation functions:

$$\mathcal{N}(\mathbf{u}_{\text{in}}; \theta) = \mathbf{W}_2 \tanh(\mathbf{W}_1 \mathbf{u}_{\text{in}} + \mathbf{b}_1) + \mathbf{b}_2, \quad (23)$$

where θ is the network's parameter vector which specifies \mathbf{W}_1 , \mathbf{W}_2 , \mathbf{b}_1 , and \mathbf{b}_2 . The `Standard` network is trained using the loss function (12) and the `Adj` network with loss function (14). `AdjVec` is trained with loss function (15) with the adjoint applied to vectors derived from 4D-Var gradient calculations (5). `Indep`'s forward model $\mathcal{N}_{\text{IndepFwd}}$ is trained using loss function (16). In training `Adj`, the value of the regularization parameter λ in (14) is set to 1.1785. The value

of λ in (15) used to train `AdjVec` is set to 1.0526. The value of the regularization parameter for `AdjVec` was determined by a parameter sweep over values for the value resulting in the optimal forward model in terms of performance on out of training set test data. The same procedure applied to `Adj` resulted in very large values of λ which produced bad solutions to the sequential 4D-Var problem. The chosen value for `Adj` was chosen by experimentation by the authors.

In our experiments, the dimensions specified by our test problem are $N_{\text{in}} = N_{\text{out}} = 3$. In this work we use two layers ($k = 2$), a hidden dimension $N_{\text{hidden}} = 25$, and

$$\mathbf{W}_1 \in \mathbb{R}^{N_{\text{hidden}} \times N_{\text{state}}}, \quad \mathbf{W}_2 \in \mathbb{R}^{N_{\text{state}} \times N_{\text{hidden}}}, \quad (24a)$$

$$\mathbf{b}_1 \in \mathbb{R}^{N_{\text{hidden}}}, \quad \mathbf{b}_2 \in \mathbb{R}^{N_{\text{state}}}, \quad (24b)$$

$$\mathbf{z} \in \mathbb{R}^{N_{\text{state}}} \quad \phi(\mathbf{x}) = \tanh(\mathbf{x}). \quad (24c)$$

Because the neural network is used as part of an optimization procedure which estimates the function's second derivative (BFGS) during offline inference, the authors speculate that the choice of \tanh or activation such as sine [4] may be preferable to ReLU, which has zero second derivative everywhere that the second derivative is defined.

For the forward independent model $\mathcal{N}_{\text{IndepFwd}}$ we use the architecture (23). The Jacobian of our feedforward network architecture (23) with regard to its input \mathbf{u}_{in} is

$$\mathbf{N}(\mathbf{u}_{\text{in}}; \theta) = \mathbf{W}_2 \text{diag}(\text{sech}^2(\mathbf{W}_1 \mathbf{u}_{\text{in}} + \mathbf{b}_1)) \mathbf{W}_1. \quad (25)$$

To analyze the benefits of training the network and its derivative independently, we introduce a nonstandard network architecture for our adjoint model. Specifically, we use the structure specified by (25) for `Indep`'s adjoint model $\mathcal{N}_{\text{IndepAdj}}$, but with independently selected weights. The model is trained with loss function (17). The architecture of the forward and adjoint models exactly matches those resulting from the single model, but retains independent weights for forward and adjoint calculations. In addition, the special structure allows for direct comparability of the weights θ_{Adj} and θ_{Fwd} .

Training for all models is done with Adam, 100 batches of size 5 per epoch, and 200 epochs. Learning rate is scheduled over training epochs from a maximum of $1e-2$ to a minimum of $1e-5$

C. Training and Test Data

To train all networks, three groups of data are collected. Forward model data $\{(\mathbf{x}_{t_i}, \mathcal{M}_{t_i, t_{i+1}}(\mathbf{x}_{t_i}))\}_{i=1}^{N_{\text{data}}}$ is used in training `Standard`, `Adj`, `AdjVec`, and `IndepFwd`. The adjoint data $\{\mathbf{M}_{t_i, t_{i+1}}^T\}_{i=1}^{N_{\text{data}}}$ corresponding to the same model runs for each i is required for training `Adj` and `IndepAdj`. We forward data and adjoint vector product data of the form $\{\mathbf{M}_{t_i, t_{i+1}}^T \mathbf{v}_i\}_{i=1}^{N_{\text{data}}}$ for training `AdjVec`.

We use explicit fourth order Runge-Kutta to collect forward model data, and its corresponding discrete adjoint method to generate adjoint data. Integration window and time integration step lengths are kept the same as in Section V-A, namely integration interval $\Delta T = 0.12$ and fixed step length $\Delta t = \Delta T/50$ for the integration method. Data is generated by one extended integration of $N_{\text{data}} = 500$ intervals of length ΔT with system state perturbed by normally




Fig. 2. First order optimality, computed as 2-norm of the surrogate cost function gradients (22) during minimization of the surrogate 4D-Var cost function (21) from a fixed starting point with different surrogate models. The two-norm of the approximate gradient is divided by its initial value. Values for `Exact` are computed during minimization of the full cost function (2) and gradient (5). Optimization is stopped when the surrogate 4D-Var cost function meets first order two-norm optimality tolerance $1e-9$ or 200 iterations. `Indep` fails to converge after 200 iterations. Remaining iterations are omitted from the plot for clarity.

distributed random noise having mean zero and covariance \mathbf{B}_0 (20) every five time intervals of length ΔT . Forward data $\{(\mathbf{x}_{t_i}, \mathcal{M}_{t_i, t_{i+1}}(\mathbf{x}_{t_i}))\}_{i=1}^{N_{\text{data}}}$ and adjoint data $\{\mathbf{M}_{t_i, t_{i+1}}^T\}_{i=1}^{N_{\text{data}}}$ is shared across all model training procedures. Adjoint-vector data is generated as follows. Vectors $\{\mathbf{v}_i\}_{i=1}^{N_{\text{data}}}$ are chosen to be of the form

$$\mathbf{v}_i = \mathbf{H}_i^T \mathbf{R}_i^{-1} (\mathcal{H}(\mathbf{x}_i) - \mathbf{y}_i). \quad (26)$$

These are saved from an independent sequential 4D-Var run using `Exact` on the same Lorenz-63 system described in Section V-A with independent noise realizations. At each timestep t_i of the sequential run, multiple gradient evaluations are processed to solve the 4D-Var optimization problem. At the first iteration of the optimization algorithm with the background state used as initial optimization condition, the vector (26) is saved. Adjoint vector product data $\{\mathbf{M}_{t_i, t_{i+1}}^T \mathbf{v}_i\}_{i=1}^{N_{\text{data}}}$ is generated by multiplying each saved \mathbf{v}_i by the previously saved adjoint $\mathbf{M}_{t_i, t_{i+1}}^T$.

Data used for the generalization test is generated the same manner. We collect 10000 data generated independently from the training data generation process are collected in one extended run of 10000 time steps of length ΔT , with system state perturbed by mean zero normally distributed random noise with covariance \mathbf{B}_0 every five steps. This amounts to the same generation procedure as the training data with independent noise realizations. Forward and adjoint information is collected for forward and adjoint model generalization tests.

D. Results and Discussion

1) *First Order Optimality*: The theoretical analysis in Section III indicates that the gradient of the 4D-Var cost function (2) depends both on the accuracy of the forward model and the accuracy of the adjoint model applied to certain vectors. This is seen clearly by examining the first-order optimality condition in (8c). We compute the norm of the high fidelity gradient

TABLE I

NORM OF THE HIGH FIDELITY GRADIENT (8c) EVALUATED AT THE OPTIMUM 4D-VAR INITIAL CONDITIONS \mathbf{x}_0^{a*} OBTAINED BY EACH ALGORITHMS. MEAN AND STANDARD DEVIATION OF 1,000 OPTIMIZATIONS FROM RANDOMLY CHOSEN STARTING POINTS.

Method	Exact	Adj	AdjVec	Standard	Indep
$\ \nabla_{\mathbf{x}_0} \Psi(\mathbf{x}_0^{a*})\ _2$	$3.42e-10 \pm 7.18e-10$	1.20 ± 0.66	5.14 ± 4.01	5.65 ± 4.40	5.50 ± 3.53

TABLE II

SPATIOTEMPORAL RMSE AND OF 4D-VAR ANALYSES OBTAINED WITH DIFFERENT SURROGATE MODELS. NEURAL NETWORKS WITH A HIDDEN DIMENSION $N_{\text{hidden}} = 25$ AND $N_{\text{data}} = 500$ DATA POINTS ARE USED IN TRAINING FOR ALL MODELS. WE REPORT THE MEAN RMSE AND STANDARD DEVIATIONS OF FIVE RUNS OF 550 INTERVALS OF LENGTH $\Delta T = 0.12$, WITH THE FIRST 50 SOLUTIONS DISREGARDED FOR ERROR CALCULATIONS.

Method	Exact	Adj	AdjVec	Standard	Indep
Analysis RMSE	0.73 ± 0.01	1.29 ± 0.04	1.50 ± 0.02	2.10 ± 0.16	6.91 ± 0.43




Fig. 3. First order optimality, computed as 2-norm of the cost function gradients (22) during minimization of approximate 4D-Var cost function (21) from a fixed starting point with different surrogate models. The two-norm of the approximate gradient is divided by its initial value. Values for *Exact* are computed during minimization of the full cost function (2) and gradient (5). Optimization is stopped when the surrogate 4D-Var cost function meets first order two-norm optimality tolerance $1e-9$ or 200 iterations. For legibility the first value of each method (normalized high resolution gradient two-norm of 1 at time 0) is omitted from plotting.

TABLE III

GENERALIZATION PERFORMANCE ON 10,000 OUT OF DATASET TEST POINTS. NEURAL NETWORK WITH $N_{\text{hidden}} = 25$, $N_{\text{data}} = 500$. GENERALIZATION PERFORMANCE OF THE MODELS INCREASES WITH THE INCORPORATION OF ADJOINT INFORMATION INTO THE TRAINING DATA. MEAN TOTAL RMSE AND STANDARD DEVIATIONS OF FIVE TEST DATA REALIZATIONS. *Indep*'S FORWARD MODEL IS SIMPLY A SECOND TRAINING REALIZATION OF *Standard*'S AND THUS IS NOT CONSIDERED.

Method	Standard	Adj	AdjVec	Indep
RMSE	1.85 ± 0.05	1.01 ± 0.01	1.15 ± 0.08	—

(8c) evaluated at the optimum 4D-Var initial conditions \mathbf{x}_0^{a*} obtained by each algorithm. The two-norm of the gradient at each algorithm's solution gives some measure of its suitability as an approximate local minimum of the original problem. Table I confirms that *Adj* and *AdjVec* solutions achieve better first order optimality with respect to the full 4D-Var cost function than the standard approach by this measure.

2) *Accuracy of 4D-Var Solutions Using Different Surrogates*: As a measure of analysis accuracy we compute the spatiotemporal root mean square error (RMSE) between the 4D-Var analyses using different surrogates and the reference

TABLE IV

MODEL ADJOINT GENERALIZATION PERFORMANCE ON 10,000 OUT OF DATASET ADJOINT TEST POINTS. NEURAL NETWORK WITH $N_{\text{hidden}} = 25$, $N_{\text{data}} = 500$. GENERALIZATION PERFORMANCE OF THE MODELS INCREASES WITH THE INCORPORATION OF ADJOINT INFORMATION INTO THE TRAINING DATA. MEAN TOTAL RMSE AND STANDARD DEVIATIONS OF FIVE TEST DATA REALIZATIONS. *Indep*'S ADJOINT MODEL SHOWS GREATLY IMPROVED ACCURACY OVER *Standard*, *Adj*, AND *AdjVec*, BUT PERFORMS WORSE IN THE CONTEXT OF SOLVING THE 4D-VAR OPTIMIZATION PROBLEM.

Method	Standard	Adj	AdjVec	Indep
RMSE	0.60 ± 0.01	0.22 ± 0.03	0.27 ± 0.01	0.07 ± 0.00

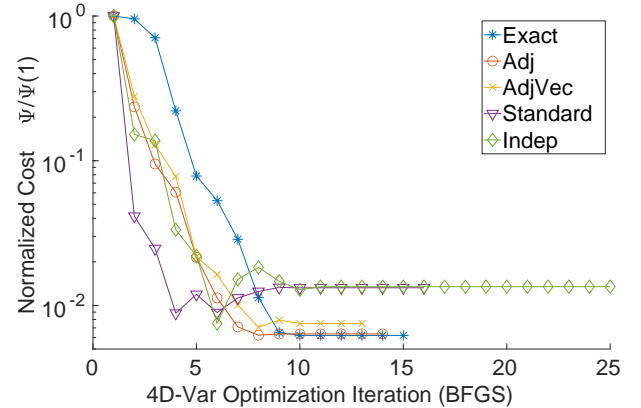


Fig. 4. Cost values of the full 4D-Var cost function (2) from a fixed starting point during optimization of the surrogate 4D-Var problem (21) with different surrogate models. The cost is calculated by calculating (2)'s value when applied to the intermediate values computed during the optimization of (21). The cost is divided by its initial value. Values for *Exact* are computed during minimization of the cost function (2) and full gradient (5). Optimization ceases when the surrogate 4D-Var cost function meets first order two-norm optimality tolerance $1e-9$ or 200 iterations. *Indep* fails to converge after 200 iterations. Iterations after 25 are omitted for clarity.

trajectory, as follows. Given a sequence of analysis vectors $\mathbf{x}_1^a, \dots, \mathbf{x}_{N_{\text{time}}}^a$, and state vectors from the reference trajectory $\mathbf{x}_1^{\text{true}}, \dots, \mathbf{x}_{N_{\text{time}}}^{\text{true}}$, the RMSE is:

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{N_{\text{time}}} \|\mathbf{x}_i^a - \mathbf{x}_i^{\text{true}}\|_2^2}{N_{\text{time}} \cdot N_{\text{state}}}}. \quad (27)$$

Table II illustrates the accuracy of the solution to the 4D-Var problem for each method of surrogate model construction.

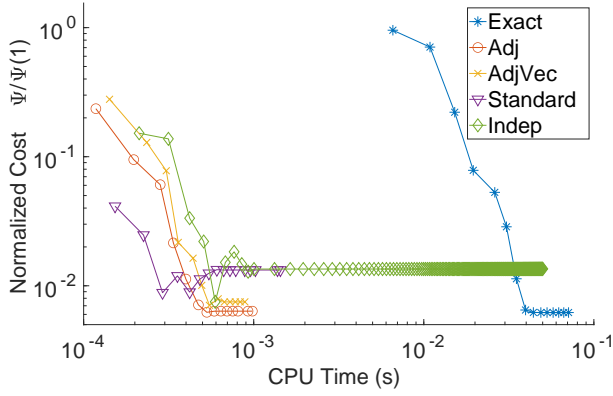


Fig. 5. Cost values of the full 4D-Var cost function (2) from a fixed starting point during optimization of the surrogate 4D-Var problem (21) with different surrogate models. The cost is calculated by calculating (2)’s value when applied to the intermediate values computed during the optimization of (21). The cost is divided by its initial value. Values for `Exact` are computed during minimization of the cost function (2) and full gradient (5). Optimization ceases when the surrogate 4D-Var cost function meets first order two-norm optimality tolerance $1e-9$ or 200 iterations. `Indep` fails to converge after 200 iterations. Iterations after 25 are omitted for clarity.




Fig. 6. Forward loss of each network during the training process. Forward loss is calculated using equation (12), where the summation is over a random batch of size 5 of training data per epoch.

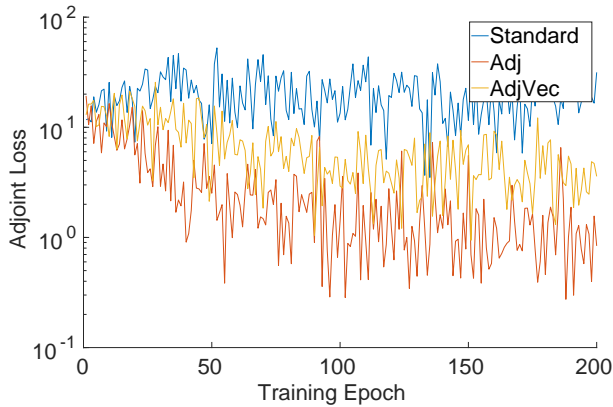


Fig. 7. Adjoint loss of each network during the training process. Adjoint loss is calculated using equation (17), where the summation is over a random batch of size 5 of training data per epoch.

We see from the theoretical analysis (11c) that the difference between the surrogate-derived and the exact analysis solutions depends on the accuracy of both the surrogate model and its adjoint model. This is confirmed in Table II. Both `Adj` and `AdjVec` demonstrate better approximate solutions than the standard approach. This indicates that the incorporation of derivative information into the training of the model surrogate improves performance when the surrogate is used for approximate solution of the variational inverse problem. `Indep` provides the least accurate solution to sequential 4D-Var problem. In the authors’ experimentation, `Indep` only occasionally arrives at catastrophically bad solutions. These sporadic failures seem to have been sufficient to cause failure in the sequential setting.

3) *Generalization Performance*: Table III shows generalization performance compared between the `Standard`, `Adj`, and `AdjVec` on out-of-training-set data. Out-of-training-set test data is generated as described in Section V-C. RMSE is calculated using the formula (27). The table reports average RMSE from five independent test data set generations. The incorporation of adjoint information into the training process appears to improve the accuracy of the forward model. Incorporation of full adjoint matrices as in `Adj`’s loss function (14) provides more benefit than incorporation only of adjoint-vector products in `AdjVec`’s loss function (15), but improved accuracy is obtained in both cases.

4) *Adjoint Model Generalization Performance*: Table IV shows generalization performance compared between the `Standard`, `Adj`, `AdjVec`, and `Indep` on out-of-training-set data. Out-of-training-set test data is generated as described in Section V-C. RMSE is calculated using the formula

$$\text{RMSE} = \sqrt{\frac{\sum_{i=1}^{N_{\text{time}}} \|\mathbf{N}^T(\mathbf{x}_{t_i}; \theta) - \mathbf{M}_{t_i, t_{i+1}}^T(\mathbf{x}_{t_i})\|_F^2}{N_{\text{time}} \cdot N_{\text{state}}^2}}.$$

The table reports average RMSE from five independent test data set generations. The incorporation of adjoint information into the training process appears to improve the accuracy of the neural network adjoint model. As is intuitive, incorporation of full adjoint matrices as in `Adj`’s loss function (14) provides more benefit than incorporation only of adjoint-vector products in `AdjVec`’s loss function (15), but benefit is obtained in both cases. `Indep` shows significantly improved accuracy in matching the full model’s adjoint dynamics.

5) *4D-Var Optimization Convergence and Timing*: We now assess the impact of different surrogates on the convergence of the 4D-Var optimization process. For this we plot the decrease in the cost function, and in the gradient norm, with the number of optimization iterations and optimization CPU time in figures 2, 3, 4, and 5. In figures 2, 3, 4, and 5, cost values are computed using the full cost function (2). Gradients used to compute first-order optimality are computed using the gradients (22) of the surrogate cost function (21).

Figures 2 and 3 show the reduction in two-norm of the approximate gradient (22) derived from each surrogate model versus BFGS iteration and time, respectively, during the optimization of (5). All of the neural networks show speedup in time to solution over `Exact`, including `Indep`.

Figures 4 and 5 demonstrate the behavior of the optimizer using various surrogates on a single minimization of the 4D-Var cost function from a common starting point. While `AdjVec`'s adjoints should not match `Exact`'s as closely as `Adj`'s, it eventually arrives at a better solution than `Standard`, providing some confirmation for the theoretical result derived in Section III. Although `Indep` does not arrive at a catastrophically bad solution, the optimization does fail to converge after 200 iterations. All neural network methods show speedup in time to solution over `Exact`.

6) *Loss During the Training Process*: Figures 6 and 7 show decrease in the data mismatch and adjoint mismatch during the training process for `Standard`, `Adj`, and `AdjVec`. Forward mismatch is calculated by modifying (12) to sum over one random batch $\tilde{5}$ of training data per epoch. Adjoint mismatch is calculated by modifying (17) to sum over one random batch $\tilde{5}$ of training data per epoch. All methods show roughly comparable change in the forward cost, while `Adj` and `AdjVec` show decrease in the adjoint mismatch.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

This work constructs science-guided neural network surrogates of dynamical systems for use in variational data assimilation. Replacing the high fidelity model with inexpensive surrogates in the inner optimization loop can speed up the solution process considerably. We propose several novel science-guided training methodologies for the neural surrogates that incorporate model adjoint information in the loss function. Our results suggest that, for a small number of training data, making use of adjoint information in the surrogate construction results in significantly improved solutions to the 4D-Var problem. The quality of the forward model, as measured by its generalization to out of training set data, also increases. Adjoint-matched neural networks thus present a promising method of training surrogates in situations where adjoint information is available.

As expected, the network performs better when full adjoint matrix information is used, than when only adjoint-vector products are incorporated. However, the adjoint-vector product training approach provides a more computationally tractable loss function, and still leads to considerable improvements in the solution when compared to the standard approach. Therefore the adjoint-vector training method holds promise for applicability to larger systems.

The methodology presented in this paper in the context of feed-forward networks can be extended to other network architectures, such as recurrent neural networks. A working implementation was created by the authors, but the difficulties with training RNNs, combined with the difficulty of efficiently evaluating the loss function (14) over batches of data rather than single points, resulted in poor experimental performance. For this reason the RNN results are not reported in the paper.

Numerical results show that two of the newly proposed approaches, namely the inclusion of adjoint operator mismatch and adjoint-vector product mismatch in the cost function, result in good approximate solution accuracy when applied to 4D-Var problems. The less accurate results obtained with independently trained models suggest that forward and adjoint

model surrogates need to be coupled. The unique architecture (25) and the comparability of the respective weights suggests that a natural way to weakly couple the two models during training is by applying the constraint

$$\|\theta_{\text{Fwd}} - \theta_{\text{Adj}}\|_2 < \epsilon$$

to the training process. This constraint can be enforced either through a penalty term or by employing a constrained optimization algorithm.

Future work will consider a nested approach to 4D-Var optimization, involving an outer loop and an inner loop [34]. After each inner optimization loop completes the full model is run again, and the neural surrogate is retrained to reflect the current high fidelity solution; the outer loop then repeats the inner optimization with the new surrogate model, and so on.

VII. ACKNOWLEDGMENTS

This work was supported by NSF through grants CDS&E-MSS-1953113 and CCF-1613905, by DOE through grant ASCR DE-SC0021313, and by the Computational Science Laboratory at Virginia Tech.

REFERENCES

- [1] C. C. AGGARWAL ET AL., *Neural networks and deep learning*, Springer, 2018.
- [2] R. ARCUCCI, J. ZHU, S. HU, AND Y.-K. GUO, *Deep data assimilation: Integrating deep learning with data assimilation*, Applied Sciences, 11 (2021).
- [3] M. ASCH, M. BOCQUET, AND M. NODET, *Data assimilation: methods, algorithms, and applications*, SIAM, 2016.
- [4] N. BORREL-JENSEN, A. P. ENGSIG-KARUP, AND C.-H. JEONG, *Physics-informed neural networks (pinns) for sound field predictions with parameterized sources and impedance boundaries*, 2021.
- [5] J. BRAJARD, A. CARRASSI, M. BOCQUET, AND L. BERTINO, *Combining data assimilation and machine learning to emulate a dynamical model from sparse and noisy observations: A case study with the lorenz 96 model*, Journal of Computational Science, 44 (2020), p. 101171.
- [6] COMPUTATIONAL SCIENCE LABORATORY, *ODE test problems*, 2021.
- [7] A. DENER, M. A. MILLER, R. M. CHURCHILL, T. MUNSON, AND C.-S. CHANG, *Training neural networks under physical constraints using a stochastic augmented lagrangian approach*, arXiv preprint, (2020).
- [8] J. E. DENNIS, JR. AND J. J. MORÉ, *Quasi-newton methods, motivation and theory*, SIAM Review, 19 (1977), pp. 46–89.
- [9] P. D. DUEBEN AND P. BAUER, *Challenges and design choices for global weather and climate models based on machine learning*, Geoscientific Model Development, 11 (2018), pp. 3999–4009.
- [10] *IFS Documentation CY47R1 - Part III: Dynamics and Numerical Procedures*, no. 3 in IFS Documentation, ECMWF, 2020, ch. 3.
- [11] G. EVENSEN, *Data assimilation: the ensemble Kalman filter*, Springer Science & Business Media, 2009.
- [12] A. FARCHI, M. BOCQUET, P. LALOYAUX, M. BONAVIDA, AND Q. MALARTIC, *A comparison of combined data assimilation and machine learning methods for offline and online model error correction*, Journal of Computational Science, (2021), p. 101468.
- [13] A. FARCHI, P. LALOYAUX, M. BONAVIDA, AND M. BOCQUET, *Exploring machine learning for data assimilation*, 2020.
- [14] E. C. FOR MEDIUM-RANGE WEATHER FORECASTS, *AI and machine learning at ECMWF*.
- [15] M. FRANGOS, Y. MARZOUK, K. WILLCOX, AND B. VAN BLOEMEN WAANDERS, *Surrogate and Reduced-Order Modeling: A Comparison of Approaches for Large-Scale Statistical Inverse Problems*, John Wiley & Sons, Ltd, 2010, ch. 7, pp. 123–149.
- [16] J. GUCKENHEIMER AND P. HOLMES, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, vol. 42, Springer Science & Business Media, 2013.
- [17] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, Proceedings of the 3rd International Conference on Learning Representations (ICLR), (2014).

- [18] E. N. LORENZ, *Deterministic nonperiodic flow*, Journal of the atmospheric sciences, 20 (1963), pp. 130–141.
- [19] J. NOCEDAL AND S. WRIGHT, *Numerical optimization*, Springer Science & Business Media, 2006.
- [20] N. PANDA, M. G. FERNÁNDEZ-GODINO, H. C. GODINEZ, AND C. DAWSON, *A data-driven non-linear assimilation framework with neural networks*, Computational Geosciences, (2020).
- [21] A. A. POPOV, C. MOU, A. SANDU, AND T. ILIESCU, *A multifidelity ensemble Kalman filter with reduced order control variates*, SIAM Journal on Scientific Computing, 43 (2021), pp. A1134–A1162.
- [22] A. A. POPOV AND A. SANDU, *Multifidelity ensemble Kalman filtering using surrogate models defined by physics-informed autoencoders*, arXiv preprint arXiv:2102.13025, (2021).
- [23] M. RAISSI, P. PERDIKARIS, AND G. E. KARNIADAKIS, *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, Journal of Computational Physics, 378 (2019), pp. 686–707.
- [24] V. RAO AND A. SANDU, *A posteriori error estimates for inverse problems*, SIAM/ASA Journal on Uncertainty Quantification, 3 (2015), pp. 737–761.
- [25] S. REICH AND C. COTTER, *Probabilistic forecasting and Bayesian data assimilation*, Cambridge University Press, 2015.
- [26] S. ROBERTS, A. A. POPOV, A. SARSHAR, AND A. SANDU, *Ode test problems: a matlab suite of initial value problems*, arXiv preprint arXiv:1901.04098, (2021).
- [27] A. SANDU, *On the properties of runge-kutta discrete adjoints*, in Computational Science – ICCS 2006, V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, eds., Berlin, Heidelberg, 2006, Springer Berlin Heidelberg, pp. 550–557.
- [28] A. SANDU, *Solution of inverse ODE problems using discrete adjoints*, in Large Scale Inverse Problems and Quantification of Uncertainty, L. Biegler, G. Biros, O. Ghattas, M. Heinkenschloss, D. Keyes, B. Mallick, L. Tenorio, B. van Bloemen Waanders, and K. Willcox, eds., John Wiley & Sons, 2010, ch. 12, pp. 345–364.
- [29] A. SANDU, *Introduction to data assimilation: Solving inverse problems with dynamical systems*, 2019.
- [30] ———, *Variational smoothing: 4d-var data assimilation*, 2019.
- [31] A. SANDU, D. DAESCU, G. CARMICHAEL, AND T. CHAI, *Adjoint sensitivity analysis of regional air quality models*, Journal of Computational Physics, 204 (2005), pp. 222–252.
- [32] K. SINGH AND A. SANDU, *Variational chemical data assimilation with approximate adjoints*, Computers & Geosciences, 40 (2012), pp. 10–18.
- [33] R. STEFANESCU AND A. SANDU, *Efficient approximation of sparse Jacobians for time-implicit reduced order models*, International Journal of Numerical Methods in Fluids, 83 (2017), pp. 175–204.
- [34] R. STEFANESCU, A. SANDU, AND I. NAVON, *POD/DEIM strategies for reduced data assimilation systems*, Journal of Computational Physics, 295 (2015), pp. 569–595.
- [35] P. J. VAN LEEUWEN, *Nonlinear Data Assimilation for high-dimensional systems*, Springer International Publishing, Cham, 2015, pp. 1–73.
- [36] J. A. WEYN, D. R. DURRAN, AND R. CARUANA, *Can machines learn to predict weather? using deep learning to predict gridded 500-hpa geopotential height from historical weather data*, Journal of Advances in Modeling Earth Systems, 11 (2019), pp. 2680–2693.
- [37] J. WILLARD, X. JIA, S. XU, M. STEINBACH, AND V. KUMAR, *Integrating physics-based modeling with machine learning: A survey*, arXiv preprint, (2020), pp. 883–894.
- [38] L. ZHANG, E. CONSTANTINESCU, A. SANDU, Y. TANG, T. CHAI, G. CARMICHAEL, D. BYUN, AND E. OLAGUER, *An adjoint sensitivity analysis and 4D-Var data assimilation study of Texas air quality*, Atmospheric Environment, 42 (2008), pp. 5787–5804.