

Learning Hidden States in a Chaotic System: A Physics-Informed Echo State Network Approach^{*}

N.A.K. Doan^{1,2}, W. Polifke¹, and L. Magri^{3,4}

¹ Department of Mechanical Engineering, Technical University of Munich, Germany

² Institute for Advanced Study, Technical University of Munich, Germany

³ Department of Engineering, University of Cambridge, United Kingdom

⁴ (visiting) Institute for Advanced Study, Technical University of Munich, Germany

Abstract. We extend the Physics-Informed Echo State Network (PI-ESN) framework to reconstruct the evolution of an unmeasured state (hidden state) in a chaotic system. The PI-ESN is trained by using (i) data, which contains no information on the unmeasured state, and (ii) the physical equations of a prototypical chaotic dynamical system. Non-noisy and noisy datasets are considered. First, it is shown that the PI-ESN can accurately reconstruct the unmeasured state. Second, the reconstruction is shown to be robust with respect to noisy data, which means that the PI-ESN acts as a denoiser. This paper opens up new possibilities for leveraging the synergy between physical knowledge and machine learning to enhance the reconstruction and prediction of unmeasured states in chaotic dynamical systems.

Keywords: Echo State Networks · Physics-Informed Echo State Networks · Chaotic Dynamical Systems · State Reconstruction.

1 Introduction

In physical science experiments, it is often difficult to measure all the physical states, whether it be because the instruments have a finite resolution, or because the measurement techniques have some limitations. Consequently, we are typically able to infer only a few states of the system from the measured observable quantities. The states that cannot be measured are *hidden*, that is, they may affect the system's evolution, but they cannot be straightforwardly measured. The accurate reconstruction of hidden states is crucial in many fields such as cardiac blood flow modelling [13], climate science [6], and fluid dynamics [2], to name only a few. For example, in fluid dynamics, measurements of the velocity field with particle image velocimetry may be limited to the in-plane two-dimensional

^{*} The authors acknowledge the support of the Technical University of Munich - Institute for Advanced Study, funded by the German Excellence Initiative and the European Union Seventh Framework Programme under grant agreement no. 291763. L.M. also acknowledges the Royal Academy of Engineering Research Fellowship Scheme.

velocity, although the three-dimensional velocity is the quantity of interest. The reconstruction of unmeasured quantities from experimental measurements has been the subject of recent studies, that used a variety of data assimilation and/or machine learning techniques. For example, spectral nudging, which combines data assimilation with physical equations, was used to infer temperature and rotation rate in 3D isotropic rotating turbulence [3]. Alternatively, [5] reconstructed the fine-scale features of an unsteady flow from large scale information by using a series of Convolutional Neural Networks. Using a similar approach, the reconstruction of the velocity from hydroxyl-radical planar laser induced fluorescence images in a turbulent flame was also performed [1]. Another approach based on echo state networks has also been used for the reconstruction of time series of unmeasured states of chaotic systems [9]. While effective in reconstructing the unmeasured states, these approaches required training data with both the measured *and* unmeasured states. In this paper, we propose using physical knowledge to reconstruct hidden states in a chaotic system *without the need of any data of the unmeasured states during the training*. This is performed with the Physics-Informed Echo State Network (PI-ESN), which has been shown to be suited to the accurate forecasting of chaotic systems [4]. The PI-ESN, and more generally *Physics-Informed* Machine Learning, relies on using the physical knowledge of the system under study, in the form of its conservation equations, and adding the computations of these conservation equations in the loss function during the training of the machine learning framework [4, 12]. These approaches, which combine physical knowledge and machine learning, have been shown to be efficient in improving the accuracy of neural networks [4, 12]. The PI-ESN approach will be, here, applied to the Lorenz system, which is a prototypical chaotic system [8].

The paper is organized as follows. The problem statement and the methodology based on PI-ESN are detailed in Sect. 2. Then, results are presented and discussed in Sect. 3 and final comments are summarized in Sect. 4.

2 Methodology: Physics-Informed Echo State Network for learning of hidden states

We consider a dynamical system whose governing equations are:

$$\mathcal{F}(\mathbf{y}) \equiv \dot{\mathbf{y}} + \mathcal{N}(\mathbf{y}) = 0 \quad (1)$$

where $\mathcal{F}$ is a non-linear operator, $\dot{\cdot}$ is the time derivative and $\mathcal{N}$ is a nonlinear differential operator. Eq. (1) represents a formal ordinary differential equation, which governs the dynamics of a nonlinear system. It is assumed that only a subset of the system states can be observed, which is denoted $\mathbf{z} \in \mathbb{R}^{N_z}$, while the hidden states are denoted $\mathbf{h} \in \mathbb{R}^{N_h}$. The full state vector is $\mathbf{y} \in \mathbb{R}^{N_y}$, which is the concatenation of $\mathbf{z}$ and $\mathbf{h}$, i.e., $\mathbf{y} = [\mathbf{z}; \mathbf{h}]$. The vectors' dimensions are related by $N_y = N_z + N_h$. The objective is to train a PI-ESN to reconstruct the hidden states, $\mathbf{h}$. We assume that we have training data of the measured states

$\mathbf{z}(n)$ only, where $n = 0, 1, 2, \dots, N_t - 1$ are the discrete time instants that span from 0 to $T = (N_t - 1)\Delta t$, where Δt is the sampling time. Thus, the specific goal for the PI-ESN is to reconstruct the hidden time series, $\mathbf{h}(n)$, for the same time instants. To solve this problem, the PI-ESN of [4], which is based on the *data-only* ESN of [10], needs to be extended, as explained next.

The PI-ESN is composed of three main parts (Fig. 1): (i) an artificial high dimensional dynamical system, i.e., the reservoir, whose neurons' (or units') states at time n are represented by a vector, $\mathbf{x}(n) \in \mathbb{R}^{N_x}$, representing the reservoir neuron activations; (ii) an input matrix, $\mathbf{W}_{in} \in \mathbb{R}^{N_x \times (1+N_u)}$, and (iii) an output matrix, $\mathbf{W}_{out} \in \mathbb{R}^{N_y \times (N_x+N_u+1)}$. The reservoir is coupled to the input signal, $\mathbf{u} \in \mathbb{R}^{N_u}$, via $\mathbf{W}_{in}$. A bias term is added to the input to excite the reservoir with a constant signal. The output of the PI-ESN, $\hat{\mathbf{y}}$, is a linear combination of the reservoir states, inputs and an additional bias:

$$\hat{\mathbf{y}}(n) = [\hat{\mathbf{z}}(n); \hat{\mathbf{h}}(n)] = \mathbf{W}_{out}[\mathbf{x}(n); \mathbf{u}(n); 1] \quad (2)$$

where $[\cdot]$ indicates a vertical concatenation and $\hat{\cdot}$ denotes the predictions from the PI-ESN. The PI-ESN outputs both the measured states, $\hat{\mathbf{z}}$, and the hidden states, $\hat{\mathbf{h}}$ (Eq. (2)). The reservoir states evolve as:

$$\mathbf{x}(n) = \tanh(\mathbf{W}_{in}[\mathbf{u}(n); 1] + \mathbf{W}\mathbf{x}(n-1)) \quad (3)$$

where $\mathbf{W} \in \mathbb{R}^{N_x \times N_x}$ is the recurrent weight matrix and the (element-wise) $\tanh$ function is the activation function for the reservoir neurons. Because we wish to predict a dynamical system, the input data for the PI-ESN corresponds to the measured system state at the previous time instant, $\mathbf{u}(n) = \mathbf{z}(n-1)$, which is only a subset of the state vector. In the ESN approach [10], the input and recurrent matrices, $\mathbf{W}_{in}$ and $\mathbf{W}$, are randomly initialized once and are not trained. Only $\mathbf{W}_{out}$ is trained. The sparse matrices $\mathbf{W}_{in}$ and $\mathbf{W}$ are constructed to satisfy the Echo State Property [10]. Following [11], $\mathbf{W}_{in}$ is generated such that each row of the matrix has only one randomly chosen nonzero element, which is independently taken from a uniform distribution in the interval $[-\sigma_{in}, \sigma_{in}]$. Matrix $\mathbf{W}$ is constructed with an average connectivity $\langle d \rangle$, and the non-zero elements are taken from a uniform distribution over the interval $[-1, 1]$. All the coefficients of $\mathbf{W}$ are then multiplied by a constant coefficient for the largest absolute eigenvalue of $\mathbf{W}$, i.e. the spectral radius, to be equal to a value Λ , which is typically smaller than (or equal to) 1. To train the PI-ESN, hence $\mathbf{W}_{out}$, a




Fig. 1. Schematic of the ESN. ① indicates the bias.

combination of the data available and the physical knowledge of the system is used: the components of $\mathbf{W}_{out}$ are computed such that they minimize the sum of (i) the error between the PI-ESN prediction and the measured system states, E_d , and (ii) the physical residual, $\mathcal{F}(\hat{\mathbf{y}}(n))$, on the prediction of the ESN, E_p :

$$E_{tot} = \underbrace{\frac{1}{N_t} \sum_{n=0}^{N_t-1} \frac{1}{N_z} \sum_{i=1}^{N_z} \|\hat{z}_i(n) - z_i(n)\|^2}_{E_d} + \underbrace{\frac{1}{N_t} \sum_{n=0}^{N_t-1} \frac{1}{N_y} \sum_{i=1}^{N_y} \|\mathcal{F}(\hat{y}_i(n))\|^2}_{E_p} \quad (4)$$

where $\|\cdot\|$ is the Euclidean norm. The training of the PI-ESN for the reconstruction of hidden states is initialized as follow. Matrix $\mathbf{W}_{out}$ is split into two partitions $\mathbf{W}_{z,out}$ and $\mathbf{W}_{h,out}$, i.e. $\mathbf{W}_{out} = [\mathbf{W}_{z,out}; \mathbf{W}_{h,out}]$, which are responsible for the prediction of the observed states, $\hat{\mathbf{z}} = \mathbf{W}_{z,out}[\mathbf{x}(n); \mathbf{u}(n); 1]$, and the hidden states, $\hat{\mathbf{h}} = \mathbf{W}_{h,out}[\mathbf{x}(n); \mathbf{u}(n); 1]$, respectively. $\mathbf{W}_{z,out}$ is initialized by Ridge regression of the data available for the measured states

$$\mathbf{W}_{z,out} = \mathbf{Z}\mathbf{X}^T (\mathbf{X}\mathbf{X}^T + \gamma\mathbf{I})^{-1} \quad (5)$$

where $\mathbf{Z}$ and $\mathbf{X}$ are respectively the horizontal concatenation of the measured states, $\mathbf{z}(n)$, and associated ESN states, inputs signals and biases, $[\mathbf{x}(n); \mathbf{u}(n); 1]$ at the different time instants during training; γ is the Tikhonov regularization factor [10]; and $\mathbf{I}$ is the identity matrix. Matrix $\mathbf{W}_{h,out}$ is randomly initialized to provide an initial guess for the optimization of $\mathbf{W}_{out}$. The optimization process modifies the components of $\mathbf{W}_{out}$ to obtain the hidden states, while ensuring that the predictions on the hidden states satisfy the physical equations. The optimization is performed with a stochastic gradient method (the Adam-optimizer [7]) with a learning rate of 0.0001.

3 Results and Discussions

The approach described in Sect. 2 is tested for the reconstruction of the chaotic Lorenz system, which is described by [8]:

$$\dot{\phi}_1 = \sigma(\phi_2 - \phi_1), \quad \dot{\phi}_2 = \phi_1(\rho - \phi_3) - \phi_2, \quad \dot{\phi}_3 = \phi_1\phi_2 - \beta\phi_3 \quad (6)$$

where $\rho = 28$, $\sigma = 10$ and $\beta = 8/3$. The size of the training dataset is $N_t = 20000$ with a timestep between two time instants of $\Delta t = 0.01$. An explicit Euler scheme is used to obtain this dataset. We assume that only measurements of ϕ_1 and ϕ_2 are available for the training of the PI-ESN and the state ϕ_3 is to be reconstructed. The parameters of the reservoir of the PI-ESN are taken to be: $\sigma_{in} = 1.0$, $\Lambda = 1.0$ and $\langle d \rangle = 20$. For the initialization of $\mathbf{W}_{z,out}$ via Ridge regression, a value of $\gamma = 10^{-6}$ is used for the Tikhonov regularization. These values of the hyperparameters are taken from previous studies [9], who performed a grid search.

3.1 Reconstruction of hidden states

In Fig. 2 where the time is normalized by the largest Lyapunov exponent, $\lambda_{\max} = 0.934$, the reconstructed ϕ_3 time series is shown for the last 10% of the training data for PI-ESNs with reservoirs of 50 and 600 units. (The dominant Lyapunov exponent is the exponential divergence rate of two system trajectories, which are initially infinitesimally close to each other.) The small PI-ESN (50 units) can satisfactorily reconstruct the hidden state, ϕ_3 . The accuracy slightly deteriorates when ϕ_3 has very large minima or maxima (e.g., $\lambda_{\max}t = 202$). However, the large PI-ESN (600 units) shows an improved accuracy. The ability of the PI-ESN to reconstruct ϕ_3 , which is not present in the training data, is a key-result. The reconstruction is enabled exclusively by the knowledge of the physical equation, which is constrained into the training of the PI-ESN. This constraint allows the PI-ESN to deduce the evolution of ϕ_3 from ϕ_1 and ϕ_2 . Conversely, with neither the physical equation nor *training data for ϕ_3* , a data-only ESN cannot learn and reconstruct ϕ_3 because it has no information on it.




Fig. 2. Reconstruction of ϕ_3 .

3.2 Effect of noise

As the ultimate objective is to work with real-world experimental data, the effect of noise on the results is investigated. The training data for ϕ_1 and ϕ_2 are modified by adding Gaussian noise to the original signal to imitate additive measurements noise. Two Signal-to-Noise Ratios (SNRs) of 20 dB and 40 dB are considered. The results of the reconstructed ϕ_3 time series from the PI-ESN trained with the noisy training data are presented in Fig. 3. Despite the presence of noise in the training data, the PI-ESN well reconstructs the non-noisy ϕ_3 signal. This means that the physical constraints in Eq. (4) act as a physics-based smoother (or denoiser) of the noisy data. This can be appreciated also in the *prediction* of measured states. Figure 3b shows the prediction of state

ϕ_1 : the non-noisy original data (full black line) and the prediction from the PI-ESN (dashed red line) overlap. This means that the PI-ESN provides a denoised prediction after training. Finally, Fig. 4 shows the root mean squared error of




Fig. 3. (a) Reconstruction of ϕ_3 with PI-ESN of 600 units trained from noisy data (with zoomed inset). (b) Prediction of ϕ_1 .

the reconstructed hidden state $\widehat{\phi}_3$, $RMSE = \sqrt{\frac{1}{N_t} \sum_{n=0}^{N_t-1} (\phi_3(n) - \widehat{\phi}_3(n))^2}$, for PI-ESNs of different reservoir sizes and noise levels, where $\phi_3(n)$ is the reference non-noisy data, which we wish to recover. For the non-noisy case, there is a large decrease in the RMSE when the PI-ESNs has 300 units (or more). With noise, the performance between the non-noisy and low-noise (SNR = 40 dB) cases are similar, whereas for a larger noise level (SNR = 20 dB), a larger reservoir is required to keep the RMSE small, as it may be expected. This suggests that the PI-ESN approach may be robust with respect to noise.




Fig. 4. RMSE of the reconstructed ϕ_3 time series in the training data.

4 Conclusions and future directions

We extend the Physics-Informed Echo State Network to reconstruct the hidden states in a chaotic dynamical system. The approach combines the knowledge of

the system's physical equations and a small dataset. It is shown, on a prototypical chaotic system, that this method can (i) accurately reconstruct the hidden states; (ii) accurately reconstruct the states with training data contaminated by noise; and (iii) provide a physics-based smoothing of the noisy measured data. Compared to other reconstruction approaches, the proposed framework does not require any data of the hidden states during training. This has the potential to enable the reconstruction of unmeasured quantities in experiments of higher dimensional chaotic systems, such as fluids. This is being explored in ongoing studies. Future work also aims at assessing the effect of imperfect physical knowledge on the reconstruction of the hidden states.

This paper opens up new possibilities for the reconstruction and prediction of unsteady dynamics from partial and noisy measurements.

References

1. Barwey, S., Hassanaly, M., Raman, V., Steinberg, A.: Using machine learning to construct velocity fields from OH-PLIF images. *Combust. Sci. Technol.* (available online) (2019)
2. Brenner, M.P.: Perspective on machine learning for advancing fluid mechanics. *Phys. Rev. Fluids* **4**(10), 100501 (2019)
3. Clark Di Leoni, P., Mazzino, A., Biferale, L.: Inferring flow parameters and turbulent configuration with physics-informed data assimilation and spectral nudging. *Phys. Rev. Fluids* **3**, 104604 (2018)
4. Doan, N.A.K., Polifke, W., Magri, L.: Physics-Informed Echo State Networks for Chaotic Systems Forecasting. In: Rodrigues, J.M.F., Cardoso, P.J.S., Monteiro, J., Lam, R., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J.J., Sloot, P.M. (eds.) *Comput. Sci. – ICCS 2019*. pp. 192–198. Springer International Publishing (2019)
5. Fukami, K., Fukagata, K., Taira, K.: Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120 (2019)
6. Kalnay, E.: *Atmospheric Modeling, Data Assimilation, and Predictability*. Cambridge University Press (2003)
7. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic optimization. 3rd Int. Conf. Learn. Represent. ICLR 2015 - Conf. Track Proc. pp. 1–15 (2015)
8. Lorenz, E.N.: Deterministic Nonperiodic Flow. *J. Atmos. Sci.* **20**(2), 130–141 (1963)
9. Lu, Z., Pathak, J., Hunt, B., Girvan, M., Brockett, R., Ott, E.: Reservoir observers: Model-free inference of unmeasured variables in chaotic systems. *Chaos* **27**(4), 041102 (2017)
10. Lukoševičius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**(3), 127–149 (2009)
11. Pathak, J., Wikner, A., Fussell, R., Chandra, S., Hunt, B.R., Girvan, M., Ott, E.: Hybrid forecasting of chaotic processes: Using machine learning in conjunction with a knowledge-based model. *Chaos* **28**(4) (2018)
12. Raissi, M., Perdikaris, P., Karniadakis, G.: Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. *J. Comput. Phys.* **378**, 686–707 (2019)
13. Sankaran, S., Moghadam, M.E., Kahn, A.M., Tseng, A.M., Guccione, J.M.: Patient-Specific Multiscale Modeling of Blood Flow for Coronary Artery Bypass Graft Surgery. *Ann. Biomed. Eng.* **40**(10), 2228–2242 (2012)