

Physics-informed long short-term memory for forecasting and reconstruction of chaos^{*}

Elise Özalp¹, Georgios Margazoglou¹, and Luca Magri^{1,2}

¹ Department of Aeronautics, Imperial College London, London, SW7 2AZ, UK

² The Alan Turing Institute, London, NW1 2DB, UK

`elise.ozalp@imperial.ac.uk`

Abstract. We present the physics-informed long short-term memory (PI-LSTM) network to reconstruct and predict the evolution of unmeasured variables in a chaotic system. The training is constrained by a regularization term, which penalizes solutions that violate the system's governing equations. The network is showcased on the Lorenz-96 model, a prototypical chaotic dynamical system, for a varying number of variables to reconstruct. First, we show the PI-LSTM architecture and explain how to constrain the differential equations, which is a non-trivial task in LSTMs. Second, the PI-LSTM is numerically evaluated in the long-term autonomous evolution to study its ergodic properties. We show that it correctly predicts the statistics of the unmeasured variables, which cannot be achieved without the physical constraint. Third, we compute the Lyapunov exponents of the network to infer the key stability properties of the chaotic system. For reconstruction purposes, adding the physics-informed loss qualitatively enhances the dynamical behaviour of the network, compared to a data-driven only training. This is quantified by the agreement of the Lyapunov exponents. This work opens up new opportunities for state reconstruction and learning of the dynamics of nonlinear systems.

Keywords: Long short-term memory · Chaos · State reconstruction.

1 Introduction

Chaotic dynamics arise in a variety of disciplines such as meteorology, chemistry, economics, and engineering. Their manifestation emerges because of the exponential sensitivity to initial conditions, which makes long-term time-accurate prediction difficult. In many engineering cases, only partial information about the system's state is available, e.g., because of the computational cost or a limited number of sensors in a laboratory experiment and hence, no data is available on the unmeasured variables. Making predictions of the available observations and reconstructing the unmeasured variables is key to understanding and predicting the behaviour of dynamical systems.

^{*} This research is supported by the ERC Starting Grant No. PhyCo 949388.

Neural networks are powerful expressive tools to extract patterns from data and, once trained, they are fast and efficient at making predictions. Suitable for time series and dynamical evolutions are recurrent neural networks (RNNs) and long short-term memory networks (LSTMs), which have shown promising performance in the inference of dynamical systems with multi-scale, chaotic, or turbulent behaviour [16]. The integration of neural networks with knowledge of governing physical equations has given rise to the field of physics-informed neural networks [7,12]. The use of LSTM approaches that embed domain-specific knowledge was applied to nonlinear structures [15,17] and degradation modelling [14], but their generalization to chaotic systems remains unexplored. Other architectures, such as physics-informed RNNs with reservoir computers have been applied successfully in the short-term reconstruction of unmeasured variables [2,11] of chaotic systems.

In this paper, we propose a physics-informed LSTM (PI-LSTM) to make predictions of observed variables and, simultaneously, infer the unmeasured variables of a chaotic system. The quality of the prediction is evaluated by analysing the autonomous long-term evolution of the LSTM and collecting the statistics of the reconstructed variables. Crucial quantities for characterizing chaos are the Lyapunov exponents (LEs), which provide insight into an attractor's dimension and tangent space. In this paper, we extract the LEs from an LSTM, trained on a prototypical chaotic dynamical system. For state reconstruction, we show that it is necessary to embed prior knowledge in the LSTM, which is in the form of differential equations in this study.

The paper is structured as follows. Section 2 provides a brief introduction to the LEs of chaotic systems and the problem setup of unmeasured variables. The PI-LSTM is proposed in Section 3. In Section 4, we discuss the results for the Lorenz-96 system. Finally, we summarize the work and propose future direction in Section 5.

2 Chaotic dynamical systems

We consider a nonlinear autonomous dynamical system

$$\frac{d}{dt}\mathbf{y}(t) = f(\mathbf{y}(t)), \quad (1)$$

where $\mathbf{y}(t) \in \mathbb{R}^N$ is the state vector of the physical system and $f : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is a smooth nonlinear function. The dynamical system (1) is chaotic if infinitesimally nearby trajectories diverge at an exponential rate. This behaviour is quantified by the LEs, which measure the average rate of stretching of the trajectories in the phase space. The LEs, $\lambda_1 \geq \dots \geq \lambda_N$, provide fundamental insight into the chaoticity and geometry of an attractor. Chaotic dynamical systems have at least one positive Lyapunov exponent. In chaotic systems, the Lyapunov time $\tau_\lambda = \frac{1}{\lambda_1}$ defines a characteristic timescale for two nearby trajectories to separate, which gives an estimate of the system's predictability horizon. LEs can be computed numerically by linearizing the equations of the dynamical system

around a reference point, defining the tangent space, and by extracting the LEs based on the Gram-Schmidt orthogonalization procedure from the corresponding Jacobian [3].

2.1 State reconstruction

Let $\mathbf{y}(t) = [\mathbf{x}(t); \boldsymbol{\xi}(t)]$ be the state of a chaotic dynamical system, where $\mathbf{x}(t) \in \mathbb{R}^{N_x}$ are the observed variables and $\boldsymbol{\xi}(t) \in \mathbb{R}^{N_\xi}$ are the unmeasured variables with $N = N_x + N_\xi$. Specifically, let us assume that $\mathbf{x}(t_i)$ is measured at times $t_i = i\Delta t$ with $i = 0, \dots, N_t$ and constant time step Δt . Based on these observations, we wish to predict the full state $\mathbf{y}(t_i) = [\mathbf{x}(t_i), \boldsymbol{\xi}(t_i)]$, whilst respecting the governing equation (1), using the PI-LSTM.

3 Physics-informed long short-term memory

LSTMs have been successfully applied to time forecasting of dynamical systems when full observations are available [13, 16]. They are characterized by a cell state $\mathbf{c}_{i+1} \in \mathbb{R}^{N_h}$ and a hidden state $\mathbf{h}_{i+1} \in \mathbb{R}^{N_h}$ that are updated at each step. In the case of partial observations, the states are updated using the observed variables $\mathbf{x}(t_i)$ as follows

$$\begin{aligned} \mathbf{i}_{i+1} &= \sigma(\mathbf{W}^i[\mathbf{x}(t_i); \mathbf{h}_i] + \mathbf{b}^i), & \tilde{\mathbf{c}}_{i+1} &= \tanh(\mathbf{W}^g[\mathbf{x}(t_i); \mathbf{h}_i] + \mathbf{b}^g), \\ \mathbf{f}_{i+1} &= \sigma(\mathbf{W}^f[\mathbf{x}(t_i); \mathbf{h}_i] + \mathbf{b}^f), & \mathbf{c}_{i+1} &= \sigma(\mathbf{f}_{i+1} * \mathbf{c}_i + \mathbf{i}_{i+1} * \tilde{\mathbf{c}}_{i+1}), \\ \mathbf{o}_{i+1} &= \sigma(\mathbf{W}^o[\mathbf{x}(t_i); \mathbf{h}_i] + \mathbf{b}^o), & \mathbf{h}_{i+1} &= \tanh(\mathbf{c}_{i+1}) * \mathbf{o}_{i+1}, \end{aligned}$$

where $\mathbf{i}_{i+1}, \mathbf{f}_{i+1}, \mathbf{o}_{i+1} \in \mathbb{R}^{N_h}$ are the input, forget and output gates. The matrices $\mathbf{W}^i, \mathbf{W}^f, \mathbf{W}^o, \mathbf{W}^g \in \mathbb{R}^{N_h \times (N_x + N_h)}$ are the corresponding weight matrices, and $\mathbf{b}^i, \mathbf{b}^f, \mathbf{b}^o, \mathbf{b}^g \in \mathbb{R}^{N_h}$ are the biases. The full prediction $\tilde{\mathbf{y}}(t_{i+1}) = [\tilde{\mathbf{x}}(t_{i+1}), \tilde{\boldsymbol{\xi}}(t_{i+1})]$ is obtained by concatenating the hidden state \mathbf{h}_{i+1} with a dense layer

$$\begin{bmatrix} \tilde{\mathbf{x}}(t_{i+1}) \\ \tilde{\boldsymbol{\xi}}(t_{i+1}) \end{bmatrix} = \mathbf{W}^{dense} \mathbf{h}_{i+1} + \mathbf{b}^{dense},$$

where $\mathbf{W}^{dense} \in \mathbb{R}^{(N_x + N_\xi) \times N_h}$ and $\mathbf{b}^{dense} \in \mathbb{R}^{N_x + N_\xi}$.

LSTMs are universal approximators for an arbitrary continuous objective function [4, 5]; however, practically, the network's performance is dependent on the parameters, such as weights and biases, which are computed during the training phase. To train the weights and biases, a data-driven loss is defined on the observed data via the mean-squared error $\mathcal{L}_{dd} = \frac{1}{N_t} \sum_{i=1}^{N_t} (\mathbf{x}(t_i) - \tilde{\mathbf{x}}(t_i))^2$. To constrain the network for the unmeasured dynamics, we add a penalization term

$$\mathcal{L}_{pi} = \frac{1}{N_t} \sum_{i=1}^{N_t} \left(\frac{d}{dt} \tilde{\mathbf{y}}(t_i) - f(\tilde{\mathbf{y}}(t_i)) \right)^2.$$

This loss regularizes the network's training to provide predictions that fulfil the governing equation from (1) (up to a numerical tolerance).

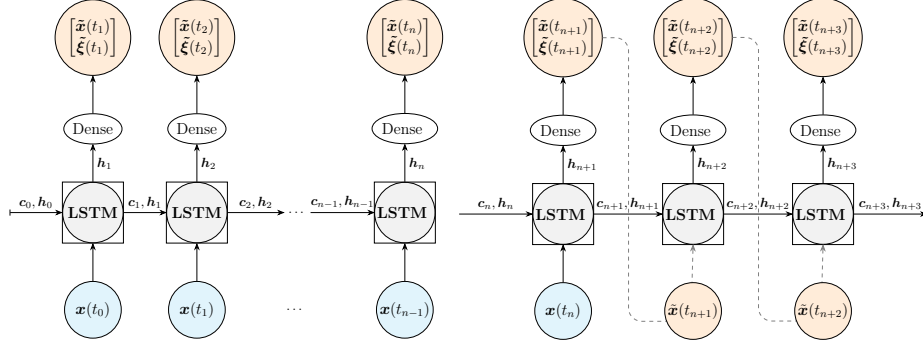


Fig. 1: Open-loop configuration for training and validation.

Fig. 2: Closed-loop configuration for autonomous evolution and testing.

For simplicity, the time derivative $\frac{d}{dt}\tilde{\mathbf{y}}$ is computed using a forward difference scheme

$$\frac{d}{dt}\tilde{\mathbf{y}}(t_i) \approx \frac{\tilde{\mathbf{y}}(t_{i+1}) - \tilde{\mathbf{y}}(t_i)}{\Delta t}.$$

(Alternatively, the time derivative can be computed with a higher-order scheme.) Combining the data-driven loss and weighing the physics-informed loss leads to the total loss

$$\mathcal{L} = \mathcal{L}_{dd} + \alpha_{pi}\mathcal{L}_{pi}, \quad \alpha_{pi} \in \mathbb{R}^+, \quad (2)$$

where α_{pi} is a penalty hyperparameter. If $\alpha_{pi} = 0$, the network is not constrained by the governing equations, which is referred to as the ‘data-driven LSTM’. The selection of an appropriate weighting for the penalty term is a nontrivial task, and a grid search is employed in this paper. The weights and biases are optimized by minimizing the loss \mathcal{L} with the Adam optimizer [6]. Early stopping is employed to avoid overfitting. During the training and validation, the network is in an open-loop configuration, as shown in Fig. 1. After training, the network is evaluated on test data with fixed weights and biases, operating in a closed-loop configuration, as shown in Fig. 2. In this mode, the network predicts the observed variables and the unmeasured variables, while the observed variables are used as input for the next time step, allowing for an autonomous evolution of the LSTM. This effectively defines a dynamical system and allows for stability analysis to be performed on the LSTM.

Previous work showed that when echo state networks, gated recurrent units or convolutional neural networks are employed to learn and predict full states from chaotic systems, the LEs of the network align with those of the dynamical system [1, 9, 10, 16] (to a numerical tolerance), allowing to gain valuable insight into the behaviour of the network. Thus, we analyse the LEs of the proposed PI-LSTM.

4 State reconstruction and Lyapunov exponents of the Lorenz-96 model

The Lorenz-96 model is a system of coupled ordinary differential equations that describe the large-scale behaviour of the mid-latitude atmosphere, and the transfer of a scalar atmospheric quantity [8]. The explicit Lorenz-96 formulation of dynamical system (1) is

$$\frac{d}{dt}y_i(t) = (y_{i+1}(t) - y_{i-2}(t))y_{i-1}(t) - y_i(t) + F, \quad i = 1, \dots, N, \quad (3)$$

where F is an external forcing. It is assumed that $y_{-1}(t) = y_{N-1}(t)$, $y_0(t) = y_N(t)$ and $y_{N+1}(t) = y_1(t)$. The state $\mathbf{y}(t) \in \mathbb{R}^N$ describes an atmospheric quantity in N sectors of one latitude circle. In this study, we set $F = 8$ and $N = 10$, for which the system exhibits chaos with three positive LEs. Both the numerical solution and the reference LEs are computed using a first-order Euler scheme with a time step of $\Delta t = 0.01$. The largest Lyapunov exponent is $\lambda_1 \approx 1.59$. The training set consists of $N_t = 20000$ points, which is equivalent to $125\tau_\lambda$. Hyperparameter tuning was utilized to optimise key network parameters. In particular, the dimension of the hidden and cell state N_h was selected from $\{20, 50, 100\}$ and the weighing of the physics-informed loss α_{pi} varied from 10^{-9} to 1 in powers of 10.

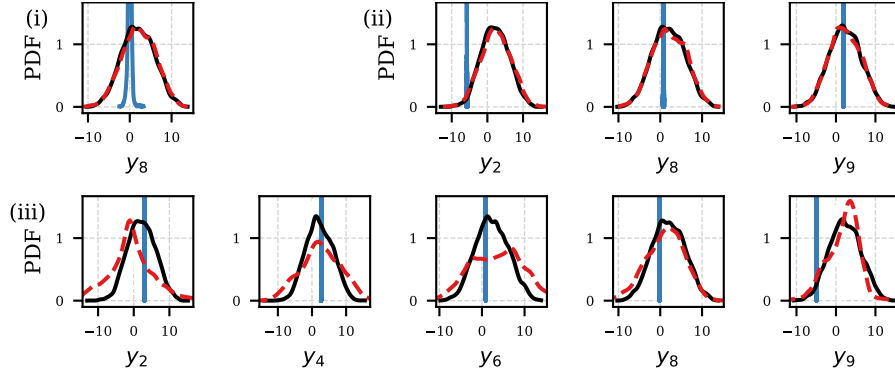


Fig. 3: Statistics reconstruction of unmeasured variables. Comparison of the target (black line), PI-LSTM (red dashed line) and data-driven LSTM (blue line) probability density functions (PDF) of (i) $N_\xi = 1$, (ii) $N_\xi = 3$, (iii) $N_\xi = 5$ unmeasured variables over a $1000\tau_\lambda$ trajectory in closed-loop configuration.

We deploy the PI-LSTM to reconstruct the unmeasured variables in three test cases: reconstruction of (i) $N_\xi = 1$, (ii) $N_\xi = 3$, (iii) $N_\xi = 5$ unmeasured variables, which are removed from the full state $\mathbf{y}(t)$ prior to training. We display networks with parameters (i) $N_h = 100, \alpha_{pi} = 0.01$, (ii) $N_h = 100, \alpha_{pi} = 0.01$, and (iii) $N_h = 50, \alpha_{pi} = 0.001$. We choose test case (i) and (ii) to highlight the

capabilities of the PI-LSTM and select (iii) to demonstrate how the network behaves with further limited information. (We remark that, when the full state is available, i.e. $N_\xi = 0$, both data-driven LSTM and PI-LSTM perform equally well in learning the long-term statistics and the LEs.) The reconstruction is based on an autonomous $1000\tau_\lambda$ long trajectory in closed-loop mode.

In Fig. 3 we show the statistics of the reconstructed variables corresponding to $\xi(t)$, which are computed based on the autonomous evolution of the network. To illustrate, in case(i) with $N_\xi = 1$, the reconstructed variable $\xi(t)$ is $y_8(t)$, and accordingly for cases (ii) and (iii). The data-driven LSTM fails to reproduce the solution, in particular, the corresponding delta-like distribution (in blue) indicates a fixed-point. In test cases (i) and (ii), the PI-LSTM accurately reproduces the long-term behaviour of the dynamical system. At each time step, it successfully extrapolates from the partial input to the full state. Case (iii) shows that, by increasing the number of unmeasured observations, the complexity of the reconstruction is increased, and the accuracy of the reconstruction is decreased. However, the PI-LSTM provides a markedly more accurate statistical reconstruction of the target compared to the data-driven LSTM. The PI-LSTM predicts the observed variables well (not shown here), which indicates that incorporating knowledge of the underlying physics enables accurate long-term prediction of the full state.




Fig. 4: Comparison of the target (black squares), PI-LSTM (red dots) and data-driven LSTM (blue dots) LEs for (i) $N_\xi = 1$, (ii) $N_\xi = 3$, (iii) $N_\xi = 5$ unmeasured variables. All the vertical axes are in logarithmic scale.

In Fig. 4, we compare the reference LEs (in black squares) with LEs extracted from the data-driven LSTM (in blue circles) and PI-LSTM (in red circles) in the three test cases. By reconstructing the unmeasured variables the networks effectively reconstruct the tangent space, the properties of which are encapsulated in the LEs. In all cases, the LEs of the data-driven LSTM deviate significantly from the reference LEs, differing more from the target when fewer observations are available. For test cases (i) and (ii), the PI-LSTM reproduces the target LEs with high accuracy, with an error of 0.28% and 8% in λ_1 , respectively. When

reducing the number of observations further, as in (iii), the accuracy of the PI-LSTM LEs is limited by the lack of information, with an error of 11.6% in λ_1 . Figure 4 also shows that in cases (ii) and (iii) of the data-driven LSTM, the leading LE is negative ($\lambda_1 < 0$), resulting in a completely incorrect solution (fixed point solution). This means that in (ii) and (iii) the data-driven LSTM displays no chaotic dynamics, whereas the PI-LSTM reproduces the chaotic behaviour.

5 Conclusions and future directions

We propose the physics-informed long short-term memory (PI-LSTM) network to embed the knowledge of the governing equations into an LSTM, by using the network’s prediction to compute a first-order derivative. In contrast to physics-informed neural networks, which have no internal recurrent connections, the PI-LSTMs capture temporal dynamics whilst penalizing predictions that violate the system’s governing equations. We deploy the PI-LSTM to reconstruct the unmeasured variables of the Lorenz-96 system, which is a chaotic system with three positive Lyapunov exponents. The long-term prediction of the PI-LSTM in closed-loop accurately reconstructs the statistics of multiple unmeasured variables. By computing the Lyapunov exponents of the PI-LSTM, we show the key role of the physics-informed loss in learning the dynamics. This exemplifies how leveraging knowledge of the physical system can be advantageous to reconstruct and predict data in line with the fundamental chaotic properties. Future work should focus on reconstructing unmeasured variables from experimental data and exploring alternative physics-informed loss formulations.

References

1. Brajard, J., Carrassi, A., Bocquet, M., Bertino, L.: 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**, 101171 (2020). <https://doi.org/10.1016/j.jocs.2020.101171>
2. Doan, N., Polifke, W., Magri, L.: Physics-informed echo state networks. *Journal of Computational Science* **47**, 101237 (2020). <https://doi.org/10.1016/j.jocs.2020.101237>
3. Eckmann, J.P., Ruelle, D.: Ergodic theory of chaos and strange attractors. *Reviews of Modern Physics* **57**, 617–656 (1985). <https://doi.org/10.1103/RevModPhys.57.617>
4. Funahashi, K.I.: On the approximate realization of continuous mappings by neural networks. *Neural Networks* **2**(3), 183–192 (1989). [https://doi.org/10.1016/0893-6080\(89\)90003-8](https://doi.org/10.1016/0893-6080(89)90003-8)
5. Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. *Neural networks* **2**(5), 359–366 (1989). [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8)
6. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015). <https://doi.org/10.48550/ARXIV.1412.6980>

7. Lagaris, I., Likas, A., Fotiadis, D.: Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks* **9**(5), 987–1000 (1998). <https://doi.org/10.1109/72.712178>
8. Lorenz, E.N.: Predictability: a problem partly solved. In: *Seminar on Predictability*, 4–8 September 1995. vol. 1, pp. 1–18. ECMWF, Shinfield Park, Reading (1996). <https://doi.org/10.1017/CB09780511617652.004>
9. Margazoglou, G., Magri, L.: Stability analysis of chaotic systems from data. *Nonlinear Dynamics* **111**, 8799–8819 (2023). <https://doi.org/10.1007/s11071-023-08285-1>
10. Pathak, J., Lu, Z., Hunt, B.R., Girvan, M., Ott, E.: Using machine learning to replicate chaotic attractors and calculate lyapunov exponents from data. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **27**(12), 121102 (2017). <https://doi.org/10.1063/1.5010300>
11. Racca, A., Magri, L.: Automatic-differentiated physics-informed echo state network (api-esn). In: *International Conference on Computational Science*. pp. 323–329. Springer (2021). <https://doi.org/10.48550/ARXIV.2101.00002>
12. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* **378**, 686–707 (2019). <https://doi.org/10.1016/j.jcp.2018.10.045>
13. Sangiorgio, M., Dercole, F., Guariso, G.: Forecasting of noisy chaotic systems with deep neural networks. *Chaos, Solitons & Fractals* **153**, 111570 (2021). <https://doi.org/10.1016/j.chaos.2021.111570>
14. Shi, J., Rivera, A., Wu, D.: Battery health management using physics-informed machine learning: Online degradation modeling and remaining useful life prediction. *Mechanical Systems and Signal Processing* **179**, 109347 (2022)
15. Tsai, L.W., Alipour, A.: Physics-informed long short-term memory networks for response prediction of a wind-excited flexible structure. *Engineering Structures* **275**, 114968 (2023). <https://doi.org/10.1016/j.engstruct.2022.114968>
16. Vlachas, P., Pathak, J., Hunt, B., Sapsis, T., Girvan, M., Ott, E., Koumoutsakos, P.: Backpropagation algorithms and reservoir computing in recurrent neural networks for the forecasting of complex spatiotemporal dynamics. *Neural Networks* **126**, 191–217 (2020). <https://doi.org/10.1016/j.neunet.2020.02.016>
17. Zhang, R., Liu, Y., Sun, H.: Physics-informed multi-lstm networks for metamodelling of nonlinear structures. *Computer Methods in Applied Mechanics and Engineering* **369**, 113226 (2020). <https://doi.org/10.1016/j.cma.2020.113226>