



Learning Ergodic Averages in Chaotic Systems

Francisco Huhn¹ and Luca Magri^{1,2}(✉)

¹ Department of Engineering, University of Cambridge, Cambridge, UK
lm547@cam.ac.uk

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

Abstract. We propose a physics-informed machine learning method to predict the time average of a chaotic attractor. The method is based on the hybrid echo state network (hESN). We assume that the system is ergodic, so the time average is equal to the ergodic average. Compared to conventional echo state networks (ESN) (purely data-driven), the hESN uses additional information from an incomplete, or imperfect, physical model. We evaluate the performance of the hESN and compare it to that of an ESN. This approach is demonstrated on a chaotic time-delayed thermoacoustic system, where the inclusion of a physical model significantly improves the accuracy of the prediction, reducing the relative error from 48% to 1%. This improvement is obtained at the low extra cost of solving a small number of ordinary differential equations that contain physical information. This framework shows the potential of using machine learning techniques combined with prior physical knowledge to improve the prediction of time-averaged quantities in chaotic systems.

Keywords: Echo state networks · Hybrid echo state networks · Physics-informed echo state networks · Chaotic dynamical systems

1 Introduction

In the past decade, there has been a proliferation of machine learning techniques applied in various fields, from spam filtering [7] to self-driving cars [3], including the more recent physical applications in fluid dynamics [4, 6]. However, a major hurdle in applying machine learning to complex physical systems, such as those in fluid dynamics, is the high cost of generating data for training [6].

L. Magri—Technical University of Munich, Germany (Visiting)

F.H. acknowledges the support from Fundação para a Ciência e Tecnologia under the Research Studentship no. SFRH/BD/134617/2017. L.M. acknowledges 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.

© Springer Nature Switzerland AG 2020

V. V. Krzhizhanovskaya et al. (Eds.): ICCS 2020, LNCS 12142, pp. 124–132, 2020.

https://doi.org/10.1007/978-3-030-50433-5_10

Nevertheless, this can be mitigated by leveraging prior knowledge (e.g. physical laws). Physical knowledge can compensate for the small amount of training data. These approaches, called physics-informed machine learning, have been applied to various problems in fluid dynamics [4, 6]. For example, [5, 14] improve the predictability horizon of echo state networks by leveraging physical knowledge, which is enforced as a *hard* constraint in [5], without needing more data or neurons. In this study, we use a hybrid echo state network (hESN) [14], originally proposed to time-accurately forecast the evolution of chaotic dynamical systems, to predict the long-term time averaged quantities, i.e., the ergodic averages. This is motivated by recent research in optimization of chaotic multi-physics fluid dynamics problems with applications to thermoacoustic instabilities [8]. The hESN is based on reservoir computing [10], in particular, conventional Echo State Networks (ESNs). ESNs have shown to predict nonlinear and chaotic dynamics more accurately and for a longer time horizon than other deep learning algorithms [10]. However, we stress that the present study is not focused on the accurate prediction of the time evolution of the system, but rather of its ergodic averages, which are obtained by the time averaging of a long time series (we implicitly assume that the system is ergodic, thus, the infinite time average is equal to the ergodic average [2]). Here, the physical system under study is a prototypical time-delayed thermoacoustic system, whose chaotic dynamics have been analyzed and optimized in [8].

2 Echo State Networks

The ESN approach presented in [11] is used here. The ESN is given an input signal $\mathbf{u}(n) \in \mathbb{R}^{N_u}$, from which it produces a prediction signal $\hat{\mathbf{y}}(n) \in \mathbb{R}^{N_y}$ that should match the target signal $\mathbf{y}(n) \in \mathbb{R}^{N_y}$, where n is the discrete time index. The ESN is composed of a reservoir, which can be represented as a directed weighted graph with N_x nodes, called neurons, whose state at time n is given by the vector $\mathbf{x}(n) \in \mathbb{R}^{N_x}$. The reservoir is coupled to the input via an input-to-reservoir matrix, \mathbf{W}_{in} , such that its state evolves according to

$$\mathbf{x}(n) = \tanh(\mathbf{W}_{\text{in}}\mathbf{u}(n) + \mathbf{W}\mathbf{x}(n-1)), \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{N_x} \times \mathbb{R}^{N_x}$ is the weighted adjacency matrix of the reservoir, i.e. W_{ij} is the weight of the edge from node j to node i , and the hyperbolic tangent is the activation function. Finally, the prediction is produced by a linear combination of the states of the neurons

$$\hat{\mathbf{y}}(n) = \mathbf{W}_{\text{out}}\mathbf{x}(n), \quad (2)$$

where $\mathbf{W}_{\text{out}} \in \mathbb{R}^{N_y} \times \mathbb{R}^{N_x}$. In this work, we are interested in dynamical system prediction. Thus, the target at time step n is the input at time step $n+1$, i.e. $\mathbf{y}(n) = \mathbf{u}(n+1)$ [14]. We wish to learn ergodic averages, given by

$$\langle \mathcal{J} \rangle = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mathcal{J}(\mathbf{u}(t)) dt, \quad (3)$$

where \mathcal{J} is a cost functional, of a dynamical system governed by

$$\dot{\mathbf{u}} = \mathbf{F}(\mathbf{u}), \quad (4)$$

where $\mathbf{u} \in \mathbb{R}^{N_u}$ is the state vector and \mathbf{F} is a nonlinear operator. The training data is obtained via numerical integration of Eq. (4), resulting in the time series $\{\mathbf{u}(1), \dots, \mathbf{u}(N_t)\}$, where the different samples are taken at equally spaced time intervals Δt , and N_t is the length of the training data set. In the conventional ESN approach, \mathbf{W}_{in} and \mathbf{W} are generated once and fixed. Then, \mathbf{W} is re-scaled to have the desired spectral radius, ρ , to ensure that the network satisfies the Echo State Property [9]. Only \mathbf{W}_{out} is trained to minimize the mean-squared-error

$$E_d = \frac{1}{N_y} \sum_{i=1}^{N_y} \frac{1}{N_t} \sum_{n=1}^{N_t} (\hat{y}_i(n) - y_i(n))^2. \quad (5)$$

To avoid overfitting, we use ridge regularization, so the optimization problem is

$$\min_{\mathbf{W}_{\text{out}}} E_d + \gamma \|\mathbf{W}_{\text{out}}\|^2, \quad (6)$$

where γ is the regularization factor. Because the prediction $\hat{\mathbf{y}}(n)$ is a linear combination of the reservoir state $\mathbf{x}(n)$, the optimal \mathbf{W}_{out} can be explicitly obtained with

$$\mathbf{W}_{\text{out}} = \mathbf{Y} \mathbf{X}^T (\mathbf{X} \mathbf{X}^T + \gamma \mathbf{I})^{-1}, \quad (7)$$

where \mathbf{I} is the identity matrix and \mathbf{Y} and \mathbf{X} are the column-concatenation of the various time instants of the output data, \mathbf{y} , and corresponding reservoir states, \mathbf{x} , respectively. After the optimal \mathbf{W}_{out} is found, the ESN can be used to predict the time evolution of the system. This is done by looping back its output to its input, i.e. $\mathbf{u}(n) = \hat{\mathbf{y}}(n-1) = \mathbf{W}_{\text{out}} \mathbf{x}(n-1)$, which, on substitution into Eq. (1), results in

$$\mathbf{x}(n) = \tanh(\widetilde{\mathbf{W}} \mathbf{x}(n-1)), \quad (8)$$

with $\widetilde{\mathbf{W}} = \mathbf{W} + \mathbf{W}_{\text{in}} \mathbf{W}_{\text{out}}$. Interestingly, Eq. (8) shows that if the reservoir follows an evolution of states $\mathbf{x}(1), \dots, \mathbf{x}(N_p)$, where N_p is the number of prediction steps, then $-\mathbf{x}(1), \dots, -\mathbf{x}(N_p)$ is also possible, because flipping the sign of \mathbf{x} in Eq. (8) results in the same equation. This implies that either the attractor of the ESN (if any) is symmetric, i.e. if some \mathbf{x} is in the ESN's attractor, then so is $-\mathbf{x}$; or the ESN has a co-existing symmetric attractor. While this seemed not to have been an issue in short-term prediction, such as in [5], it does pose a problem in the long-term prediction of statistical quantities. This is because the ESN, in its present form, *can not* generate non-symmetric attractors. This symmetry needs to be broken to work with a general non-symmetric dynamical system. This can be done by including biases [10]. However, the addition of a bias can make the reservoir prone to saturation (results not shown), i.e. $x_i \rightarrow \pm 1$, and thus care needs to be taken in the choice of hyperparameters. In this paper, we break the symmetry by exploiting prior knowledge on the physics of the problem under investigation with a hybrid ESN.




Fig. 1. Schematic of the hybrid echo state network. In training mode, the input of the network is the training data (switch is horizontal). In prediction mode, the input of the network is its output from the previous time step (switch is vertical).

3 Physics-Informed and Hybrid Echo State Network

The ESN's performance can be increased by incorporating physical knowledge during training [5] or during training and prediction [14]. This physical knowledge is usually present in the form of a reduced-order model (ROM) that can generate (imperfect) predictions. The authors of [5] introduced a physics-informed ESN (PI-ESN), which constrains the physics as a *hard* constraint with a physics loss term. The prediction is consistent with the physics, but the training requires nonlinear optimization. The authors of [14] introduced a hybrid echo state network (hESN), which incorporates incomplete physical knowledge by feeding the prediction of the physical model into the reservoir and into the output. This requires ridge regression. Here, we use an hESN (Fig. 1) because we are not interested in constraining the physics as a hard constraint for an accurate short-term prediction [5]. In the hESN, similarly to the conventional ESN, the input is fed to the reservoir via the input layer \mathbf{W}_{in} , but also to a physical model, which is usually a set of ordinary differential equations that approximately describe the system that is to be predicted. In this work, that model is a reduced-order model (ROM) of the full system. The output of the ROM is then fed to the reservoir via the input layer and into the output of the network via the output layer.

4 Learning the Ergodic Average of an Energy

We use a prototypical time-delayed thermoacoustic system composed of a longitudinal acoustic cavity and a heat source modelled with a nonlinear time-delayed model [8, 12, 16], which has been used to optimize ergodic averages in [8] with a dynamical systems approach. The non-dimensional governing equations are

$$\partial_t u + \partial_x p = 0, \quad \partial_t p + \partial_x u + \zeta p - \dot{q} \delta(x - x_f) = 0, \quad (9)$$

where u , p , ζ and \dot{q} are the non-dimensionalized acoustic velocity, pressure, damping and heat-release rate, respectively. δ is the Dirac delta. These equations




Fig. 2. Acoustic velocity at the flame location.

are discretized by using N_g Galerkin modes

$$u(x, t) = \sum_{j=1}^{N_g} \eta_j(t) \cos(j\pi x), \quad p(x, t) = - \sum_{j=1}^{N_g} \mu_j(t) \sin(j\pi x), \quad (10)$$

which results in a system of $2N_g$ oscillators, which are nonlinearly coupled through the heat released by the heat source

$$\dot{\eta}_j - j\pi\mu_j = 0, \quad \dot{\mu}_j + j\pi\eta_j + \zeta_j\mu_j + 2\dot{q} \sin(j\pi x_f) = 0, \quad (11)$$

where $x_f = 0.2$ is the heat source location and $\zeta_j = 0.1j + 0.06j^{1/2}$ is the modal damping [8]. The heat release rate, \dot{q} , is given by the modified King's law [8], $\dot{q}(t) = \beta[(1 + u(x_f, t - \tau))^{1/2} - 1]$, where β and τ are the heat release intensity parameter and the time delay, respectively. With the nomenclature of Sect. 2, $\mathbf{y}(n) = (\eta_1; \dots; \eta_{N_g}; \mu_1; \dots; \mu_{N_g})$. Using 10 Galerkin modes ($N_g = 10$), $\beta = 7.0$ and $\tau = 0.2$ results in a chaotic motion (Fig. 2), with the leading Lyapunov exponent being $\lambda_1 \approx 0.12$ [8]. (The leading Lyapunov exponent measures the rate of (exponential) separation of two close initial conditions, i.e. an initial separation $\|\delta\mathbf{u}_0\|$ grows asymptotically like $\|\delta\mathbf{u}_0\|e^{\lambda_1 t}$.) However, for the same choice of parameter values, the solution with $N_g = 1$ is a limit cycle (i.e. a periodic solution).

The echo state network is trained on data generated with $N_g = 10$, while the physical knowledge (ROM in Fig. 1) is generated with $N_g = 1$ only. We wish to predict the time average of the instantaneous acoustic energy,


$$E_{ac}(t) = \int_0^1 \frac{1}{2}(u^2 + p^2) dx, \quad (12)$$

which is a relevant metric in the optimization of thermoacoustic systems [8]. The reservoir is composed of 100 units, a modest size, half of which receive their input from \mathbf{u} , while the other half receives it from the output of the ROM, $\hat{\mathbf{y}}_{\text{ROM}}$. The entries of \mathbf{W}_{in} are randomly generated from the uniform distribution $\text{unif}(-\sigma_{\text{in}}, \sigma_{\text{in}})$, where $\sigma_{\text{in}} = 0.2$. The matrix \mathbf{W} is highly sparse, with only 3% of non-zero entries from the uniform distribution $\text{unif}(-1, 1)$. Finally, \mathbf{W} is scaled such that its spectral radius, ρ , is 0.1 and 0.3 for the ESN and the hESN,


respectively. The time step is $\Delta t = 0.01$. The network is trained for $N_t = 5000$ units, which corresponds to 6 Lyapunov times, i.e. $6\lambda_1^{-1}$. The data is generated by integrating Eq. (11) in time with $N_g = 10$, resulting in $N_u = N_y = 20$. In the hESN, the ROM is obtained by integrating the same equations, but with $N_g = 1$ (one Galerkin mode only) unless otherwise stated. Ridge regression is performed with $\gamma = 10^{-7}$. The values of the hyperparameters are taken from the literature [5, 14] and a grid search, which, while not the most efficient, is well suited when there are few hyperparameters, such as this work’s ESN architecture.

On the one hand, Fig. 3a shows the instantaneous error of the first modes of the acoustic velocity and pressure ($\eta_1; \mu_1$) for the ESN, hESN and ROM. None of these can accurately predict the instantaneous state of the system. On the other hand, Fig. 3b shows the error of the prediction of the average acoustic energy. Once again, the ROM alone does a poor job at predicting the statistics of the system, with an error of 50%. This should not come as a surprise since, as discussed previously, the ROM does not even produce a chaotic solution. The ESN, trained on data only, performs marginally better, with an error of 48%. In contrast, the hESN predicts the time-averaged acoustic energy satisfactorily, with an error of about 7%. This is remarkable, since both the ESN and the ROM do a poor job at predicting the average acoustic energy. However, when the ESN is combined with prior knowledge from the ROM, the prediction becomes significantly better. Moreover, while the hESN’s error still decreases at the end of the prediction period, $t = 250$, which is 5 times the training data time, the ESN and the ROM stabilize much earlier, at a time similar to that of the training data. This result shows that complementing the ESN with a cheap physical model (only 10% the number of degrees of freedom of the full system) can greatly improve the accuracy of the predictions, with no need for more data or neurons. Figure 3c shows the relative error as a function of the number of Galerkin modes in the ROM, which is a proxy for the quality of the model. For each N_g , we take the median of 16 reservoir realizations. As expected, as the quality of the model increases, so does the quality of the prediction. This effect is most noticeable from $N_g = 1$ to 4, with the curve presenting diminishing returns. The downside of increasing N_g is obviously the increase in computational cost. At $N_g = 10$, the original system is recovered. However, the error does not tend exactly to 0 because \mathbf{W}_{out} can not combine the ROM’s output only (i.e. 0 entries for reservoir nodes) due to: i) the regularization factor in ridge regression that penalizes large entries; ii) numerical error. This graph further strengthens the point previously made that cheap physical models can greatly improve the prediction of physical systems with data techniques.

We stress that the optimal values of hyperparameters for a certain set of physical parameters, e.g. (β_1, τ_1) , might not be optimal for a different set of physical parameters (β_2, τ_2) . This should not be surprising, since different physical parameters will result in different attractors. For example, Fig. 4 shows that changing the physical parameters from $(\beta = 7.0, \tau = 0.2)$ to $(\beta = 6.0, \tau = 0.3)$ results in a change of type of attractor from chaotic to limit cycle. For the hESN to predict the limit cycle, the value of σ_{in} must change from 0.2 to 0.03. Thus,




(a) Absolute error of $(\eta_1; \mu_1)$ prediction. (b) Relative error of $\langle E_{ac} \rangle$ prediction.



(c) Relative error vs number of Galerkin modes of the ROM.

Fig. 3. Errors on the prediction from ESN (blue), hESN (red) and ROM (green). (Color figure online)

if the hESN (or any deep learning technique in general) is to be used to predict the dynamics of various physical configurations (e.g. the generation of a bifurcation diagram), then it should be coupled with a robust method for the automatic selection of optimal hyperparameters [1], with a promising candidate being Bayesian optimization [15, 17].



(a) $(\beta = 7.0, \tau = 0.2)$: $\sigma_{in} = 0.2$ (b) $(\beta = 6.0, \tau = 0.3)$: $\sigma_{in} = 0.03$

Fig. 4. Phase plot of system (black) and hESN (red) for two sets of (β, τ) . (a) chaotic solution; (b) periodic solution. (Color figure online)

5 Conclusion and Future Directions

We propose the use of echo state networks informed with incomplete prior physical knowledge for the prediction of time averaged cost functionals in chaotic dynamical systems. We apply this to chaotic acoustic oscillations, which is relevant to aeronautical propulsion. The inclusion of physical knowledge comes at a low cost and significantly improves the performance of conventional echo state networks from a 48% error to 1%, without requiring additional data or neurons. This improvement is obtained at the low extra cost of solving a small number of ordinary differential equations that contain physical information. The ability of the proposed ESN can be exploited in the optimization of chaotic systems by accelerating computationally expensive shadowing methods [13]. For future work, (i) the performance of the hybrid echo state network should be compared against those of other physics-informed machine learning techniques; (ii) robust methods for hyperparameters' search should be coupled for a "hands-off" autonomous tool; and (iii) this technique is currently being applied to larger scale problems. In summary, the proposed framework is able to learn the ergodic average of a fluid dynamics system, which opens up new possibilities for the optimization of highly unsteady problems.

References

1. Bengio, Y.: Practical recommendations for gradient-based training of deep architectures. In: Montavon, G., Orr, G.B., Müller, K.-R. (eds.) *Neural Networks: Tricks of the Trade*. LNCS, vol. 7700, pp. 437–478. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35289-8_26
2. Birkhoff, G.D.: Proof of the ergodic theorem. *Proc. Natl. Acad. Sci.* **17**(12), 656–660 (1931)
3. Bojarski, M., et al.: End to end learning for self-driving cars. arXiv e-prints, [arXiv:1604.07316](https://arxiv.org/abs/1604.07316) (2016)
4. Brunton, S.L., Noack, B.R., Koumoutsakos, P.: Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**(1), 477–508 (2020)
5. 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., Slot, P.M.A. (eds.) *ICCS 2019*. LNCS, vol. 11539, pp. 192–198. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22747-0_15
6. Duraisamy, K., Iaccarino, G., Xiao, H.: Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51**(1), 357–377 (2019)
7. Guzella, T.S., Caminhas, W.M.: A review of machine learning approaches to spam filtering. *Expert Syst. Appl.* **36**(7), 10206–10222 (2009)
8. Huhn, F., Magri, L.: Stability, sensitivity and optimisation of chaotic acoustic oscillations. *J. Fluid Mech.* **882**, A24 (2020)
9. Jaeger, H., Haas, H.: Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**(5667), 78–80 (2004)
10. Lukoševičius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**(3), 127–149 (2009)

11. Lukoševičius, M.: A practical guide to applying echo state networks. In: Montavon, G., Orr, G.B., Müller, K.-R. (eds.) *Neural Networks: Tricks of the Trade*. LNCS, vol. 7700, pp. 659–686. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35289-8_36
12. Nair, V., Sujith, R.: A reduced-order model for the onset of combustion instability: physical mechanisms for intermittency and precursors. *Proc. Combust. Inst.* **35**(3), 3193–3200 (2015)
13. Ni, A., Wang, Q.: Sensitivity analysis on chaotic dynamical systems by non-intrusive least squares shadowing (NILSS). *J. Comput. Phys.* **347**, 56–77 (2017)
14. Pathak, J.: Hybrid forecasting of chaotic processes: using machine learning in conjunction with a knowledge-based model. *Chaos: Interdisc. J. Nonlinear Sci.* **28**(4), 41101 (2018)
15. Reinier Maat, J., Gianniotis, N., Protopapas, P.: Efficient optimization of echo state networks for time series datasets. arXiv e-prints, [arXiv:1903.05071](https://arxiv.org/abs/1903.05071) (2019)
16. Traverso, T., Magri, L.: Data assimilation in a nonlinear time-delayed dynamical system with lagrangian optimization. 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.A. (eds.) *ICCS 2019*. LNCS, vol. 11539, pp. 156–168. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-22747-0_12
17. Yperman, J., Becker, T.: Bayesian optimization of hyper-parameters in reservoir computing. arXiv e-prints, [arXiv:1611.05193](https://arxiv.org/abs/1611.05193) (2016)