

Adjoint Sensitivities of Chaotic Flows without Adjoint Solvers: A Data-Driven Approach

Defne E. Ozan¹[0009-0006-4208-9981] and Luca Magri^{1,2,3}[0000-0002-0657-2611]

¹ Imperial College London, Department of Aeronautics, Exhibition Road, London SW7 2BX, UK

d.ozan@imperial.ac.uk, l.magri@imperial.ac.uk

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

³ Politecnico di Torino, DIMEAS, Corso Duca degli Abruzzi, 24 10129 Torino, Italy

Abstract. In one calculation, adjoint sensitivity analysis provides the gradient of a quantity of interest with respect to all system’s parameters. Conventionally, adjoint solvers need to be implemented by differentiating computational models, which can be a cumbersome task and is code-specific. To propose an adjoint solver that is not code-specific, we develop a data-driven strategy. We demonstrate its application on the computation of gradients of long-time averages of chaotic flows. First, we deploy a parameter-aware echo state network (ESN) to accurately forecast and simulate the dynamics of a dynamical system for a range of system’s parameters. Second, we derive the adjoint of the parameter-aware ESN. Finally, we combine the parameter-aware ESN with its adjoint version to compute the sensitivities to the system parameters. We showcase the method on a prototypical chaotic system. Because adjoint sensitivities in chaotic regimes diverge for long integration times, we analyse the application of ensemble adjoint method to the ESN. We find that the adjoint sensitivities obtained from the ESN match closely with the original system. This work opens possibilities for sensitivity analysis without code-specific adjoint solvers.

Keywords: Reservoir computing · Adjoint methods · Sensitivity · Chaotic flows

1 Introduction

Many computational applications including optimization, data assimilation, and uncertainty quantification require the information about how a quantity of interest is influenced by the system’s parameters and initial conditions e.g., [9]. Precisely, sensitivity analysis provides the gradient of an objective functional with respect to all system’s parameters. A direct method for estimating the sensitivity is to perturb the parameter and approximate the gradient with a finite difference between the base and the perturbed solutions. This procedure suffers from numerical errors and needs to be repeated for each parameter. Therefore, for systems with many parameters, the computational cost increases linearly with the number of parameters. Adjoint methods turn this around. With an

adjoint formulation, the sensitivity to all parameters is computed with a single operation, which requires the adjoint system.

Whilst the adjoint is a powerful tool, there are some practical limitations to its use. The computation of the adjoint sensitivity relies on (i) the system equations to be known, and (ii) the system to be linearized, i.e., the computation of the Jacobian [4, 13]. First, the adjoint sensitivity is only as accurate as the model assumptions. Second, the system might be high-dimensional and nonlinear, and in some cases the computational code might not be differentiable, all of which make the derivation of the Jacobian cumbersome. To address the challenges associated with developing adjoint solvers, we propose a data-driven approach, and apply this to a prototypical chaotic flow, which is a qualitative example of complex behaviour such as turbulence.

In chaotic systems, in which nearby trajectories diverge at an exponential rate, the adjoint system becomes unstable and the sensitivities diverge after some integration time. This poses a challenge as the quantities of interest are usually long-term time-averages. A possible approach that we also exploit in this paper is to take an ensemble of short-term trajectories [6].

The objective of this paper is two-fold; (i) we introduce the parameter-aware echo state network to learn the parametrized dynamics of a chaotic system with, and (ii) we infer the sensitivities of an objective functional to the system parameters from the adjoint of parameter-aware echo state network.

2 Parameter-aware Echo State Network

Recurrent neural networks (RNNs) are specialized neural network architectures for sequential data processing and have been successfully employed for time series prediction. RNNs can be treated as discrete differential equations and can be studied with dynamical systems theory. In this work, we develop an RNN that learns the parametrized dynamics of a system such that we can make predictions in regimes that have not been seen during training. We further show that such a network can also infer the sensitivity of the dynamics to the parameters. We utilize the echo state network (ESN) [8], which whilst being a universal approximator [5] has the advantage over other RNNs of not needing backpropagation during training, and thus requiring less computational effort. ESNs can make time-accurate short-term predictions [11, 2, 14], and infer long-term statistics and invariant properties, such as Lyapunov exponents, of chaotic flows [12, 10]. A parameter-aware extension of the ESN has been formulated, and been applied to predict amplitude death, i.e., when the system bifurcates to fixed point solutions, [16], and multi-stable regimes including chaos [15]. The parameter-aware ESN is a nonlinear discrete map from reservoir state at time step i to reservoir state at time step $i + 1$

$$\mathbf{r}(i + 1) = (1 - \alpha)\mathbf{r}(i) + \alpha \tanh(\mathbf{W}_{in}[\mathbf{y}_{in}(i); \text{diag}(\boldsymbol{\sigma}_p)(\mathbf{p} - \mathbf{k}_p)] + \mathbf{W}\mathbf{r}(i)), \quad (1)$$

where $\mathbf{y}_{in}(i) \in \mathbb{R}^{N_y}$ is the input vector, $\mathbf{p} \in \mathbb{R}^{N_p}$ is the parameter vector, $\mathbf{r}(i) \in \mathbb{R}^{N_r}$ is the reservoir state, $\mathbf{W}_{in} \in \mathbb{R}^{N_r \times (N_y + N_p)}$ is the input matrix, and

$\mathbf{W} \in \mathbb{R}^{N_r \times N_r}$ is the state matrix, and $\text{diag}(\cdot)$ denotes a diagonal matrix that has (\cdot) as its diagonal. The output is predicted from the reservoir state with a linear read-out layer

$$\hat{\mathbf{y}}(i+1) = \mathbf{W}_{out} \mathbf{r}(i+1), \quad (2)$$

$\hat{\mathbf{y}}(i+1) \in \mathbb{R}^{N_y}$ is the output vector, the dynamics of which we aim to model, and $\mathbf{W}_{out} \in \mathbb{R}^{N_y \times N_r}$ is the output matrix. The matrices \mathbf{W}_{in} and \mathbf{W} are sparse, randomly generated, and not trained, whilst \mathbf{W}_{out} is trained via ridge regression [8]. Training an ESN involves searching for optimal hyperparameters that consist of input matrix scaling σ_{in} , the spectral radius of the state matrix ρ , the leak rate α , the Tikhonov regularizer λ , and finally two additional hyperparameters per parameter p_i ; k_p that shifts the parameter and σ_p that scales it.

3 Adjoint of Echo State Network

We mathematically derive the adjoint of the ESN, which is an autonomous dynamical system when it runs in closed-loop. The dynamics (1) can be expressed as a constraint $\mathbf{F}(i) = \mathbf{F}(\mathbf{r}(i), \mathbf{r}(i-1), \mathbf{p}) = 0$. We consider the time-averaged objective functional given as a sum over N discrete time steps

$$\mathcal{J} = \frac{1}{N} \sum_{i=1}^N \tilde{\mathcal{J}}(\mathbf{r}(i)), \quad (3)$$

in which the reservoir state \mathbf{r} depends on the system's parameters \mathbf{p} . The goal of sensitivity analysis is to determine the gradient of the objective functional with respect to the system's parameters. For the parameter-aware ESN, the sensitivity is then expressed as

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \frac{1}{N} \sum_{i=1}^N \frac{d\tilde{\mathcal{J}}(\mathbf{r}(i))}{d\mathbf{r}(i)} \mathbf{Q}(i), \quad (4)$$

where we define $\mathbf{Q}(i) := d\mathbf{r}(i)/d\mathbf{p}$, $\mathbf{Q}(i) \in \mathbb{R}^{N_r \times N_p}$. Notice that the dimension of $\mathbf{Q}(i)$ grows with the number of parameters. This increasing computational cost motivates solving the adjoint problem instead. The Lagrangian of the objective functional, \mathcal{L} , subjected to the system dynamics, $\mathbf{F} = 0$, is

$$\mathcal{L} := \mathcal{J} - \langle \mathbf{q}^+, \mathbf{F} \rangle, \quad (5)$$

where the dot product $\langle \mathbf{q}^+, \mathbf{F} \rangle$ is defined as

$$\langle \mathbf{q}^+, \mathbf{F} \rangle := \sum_{i=1}^N \mathbf{q}^{+T}(i) \mathbf{F}(i), \quad (6)$$

and $\mathbf{q}^+ \in \mathbb{R}^{N_r}$ are the Lagrange multipliers, or the adjoint variables. We solve for the gradient $d\mathcal{L}/d\mathbf{p} = d\mathcal{J}/d\mathbf{p}$. After expanding the summation and rearranging the terms, we choose the Lagrange multipliers such that we can eliminate the

terms $d\mathbf{r}(i)/d\mathbf{p}$, e.g., [9]. This provides us with the evolution equations of the adjoint variables

$$\frac{d\mathcal{J}}{d\mathbf{p}} = \sum_{i=1}^N \mathbf{q}^{+T}(i) \frac{\partial \mathbf{r}(i)}{\partial \mathbf{p}}, \quad (7a)$$

$$\mathbf{q}^+(i) = \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{r}(i))}{\partial \mathbf{r}(i)}^T + \frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i)}^T \mathbf{q}^+(i+1), \quad (7b)$$

$$\mathbf{q}^+(N) = \frac{1}{N} \frac{\partial \tilde{\mathcal{J}}(\mathbf{r}(N))}{\partial \mathbf{r}(N)}^T. \quad (7c)$$

Practically, we first let the ESN run autonomously for the given time window and save the direct solution, which serves as the base trajectory. We then solve the adjoint equations backwards in time starting from the terminal condition $\mathbf{q}^+(N)$ (7c). The solution of the adjoint equations requires the computation of the Jacobian evaluated at the base trajectory, i.e., the gradient of the reservoir state at time step $i+1$ with respect to the reservoir state at time step i

$$\frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{r}(i)} = (1-\alpha) \mathbf{I}_{N_r \times N_r} + \alpha \text{diag}(1 - \tilde{\mathbf{r}}^2(i)) (\mathbf{W}_{in}^y \mathbf{W}_{out} + \mathbf{W}), \quad (8)$$

and the gradient with respect to the parameters

$$\frac{\partial \mathbf{r}(i+1)}{\partial \mathbf{p}} = \alpha \text{diag}(1 - \tilde{\mathbf{r}}^2(i)) \mathbf{W}_{in}^p \text{diag}(\boldsymbol{\sigma}_p), \quad (9)$$

where $\tilde{\mathbf{r}}(i) = (\mathbf{r}(i+1) - (1-\alpha)\mathbf{r}(i))/\alpha$, $\mathbf{W}_{in} = [\mathbf{W}_{in}^y \ \mathbf{W}_{in}^p]$, and \mathbf{I} denotes the identity matrix.

4 Computation of data-driven chaotic sensitivities

We demonstrate the data-driven computation of adjoint sensitivities on the Lorenz 63 system, which is a reduced-order model to study atmospheric convection

$$\frac{dx}{dt} = s(y-x), \quad \frac{dy}{dt} = x(r-z) - y, \quad \frac{dz}{dt} = xy - bz, \quad (10)$$

where x, y, z are the state variables, and s, r, b are the system's parameters. We generate the dataset for the training and validation of the ESN by time-marching the ODEs (10) for different sets of parameters. The parameters ($s = 10, r = 28, b = 8/3 \approx 2.667$) lead to a chaotic solution [7]. We randomly choose 20 regimes for training, and 5 regimes for validation from a grid of parameters, $s = \{8, 10, 12, 14, 16\}$, $r = \{30, 35, 40, 45, 50\}$, and $b = \{1, 1.5, 2, 2.5, 3\}$. We observe that it is important to validate the ESN on regimes unseen during the training in order to choose a model that generalizes well over a range of different parameters. The regimes in the training and validation datasets display chaotic behaviour with varying Lyapunov times (LTs), i.e., the time-scale of divergence

of two nearby trajectories in a chaotic system, (between 0.77 and 4.80 time units). The numerical integration is performed with a fourth-order Runge-Kutta scheme with a time-step of $\Delta t = 0.01$. After a transient, we select the first 4 time units for washout and 10 time units as training data. The hyperparameters for the ESN are determined via Bayesian optimization [14], in which we evaluate the short-term closed-loop performance of random realisations of the model on the validation dataset. The optimal hyperparameters for an ESN of reservoir size, $N_r = 1200$ and a connectivity, $N_{conn} = 3$ between the reservoir state variables are found as $\rho = 0.2201$, $\sigma_{in} = 0.0679$, $\sigma_s = 0.0028$, $k_s = 68.73$, $\sigma_r = 0.0015$, $k_r = 84.81$, $\sigma_b = 0.0393$, $k_b = 74.46$, $\alpha = 0.8853$, and $\lambda = 10^{-10}$.

Previous works have applied and analysed the ensemble adjoint method to compute the sensitivity of the time-averaged z , \bar{z} , to the parameter ρ in the Lorenz 63 system due to their nearly linear relationship e.g., [6, 3]. Therefore, we choose \bar{z} as the quantity of interest. First, the short-term prediction and the long-term inference of statistics for z are illustrated in Fig. 1 for two regimes; (i) the reference configuration, ($s = 10, r = 28, b = 8/3 \approx 2.667$) with a Lyapunov time of $LT = 1.1$, and (ii) a configuration with a shorter Lyapunov time $LT = 0.8$ ($s = 13, r = 52, b = 2.0$). Neither of these regimes were seen during training or validation. Nonetheless, time-accurate short-term forecasting can be performed with a predictability horizon e.g., [14] of $4.8 LT$ for the regime (i), and $5.8 LT$ for regime (ii) (average over 100 initial conditions on the attractor), and the long-term statistics are captured for a wide range of seen and unseen regimes with a single instance of the parameter-aware ESN.




Fig. 1. Short-term prediction (a,c) and long-term inference of the statistics (b,d) of z for two regimes with different Lyapunov times (LTs); (a,b) ($s = 10, r = 28, b = 8/3 \approx 2.667$), and (c,d) ($s = 13, r = 52, b = 1.75$). The statistics are calculated over 5000 LTs, after a washout stage where we repeatedly feed the same initial condition, and a transient time that is discarded. The parameter-aware echo state network can successfully infer the dynamics and long-term statistics of different chaotic regimes even when they were not seen during training.

Next, we compute the “climate” sensitivity of \bar{z} to all parameters using an ensemble adjoint method [6]. Climate sensitivity refers to the long-term behaviour

of the system, where the sensitivity is independent of initial condition. In the ensemble adjoint method, the sensitivities over an ensemble of short time-series with different initial conditions on the attractor are computed and the climate sensitivity is estimated by the mean of the ensemble. The sensitivity estimator is associated with a bias and a variance depending on the integration time of each trajectory [1]. While the bias decreases with increasing integration time, the variance increases and the so-called Lévy flights, i.e., long jumps in the mean estimation, appear, requiring more ensemble members for convergence [3]. We opt for an integration time of half a Lyapunov time and an ensemble of 10000 trajectories, which provide an acceptable trade-off between accuracy and computational cost. Figure 2 compares the objective and the adjoint sensitivity estimates of the ESN with the original system as well as a direct estimate using a polynomial fit of the objective values. Sensitivities pertaining to the parameters s , r , and b , are shown while each of them is varied and the other two are fixed at the given values for the regime (i). The ESN estimates closely match with the true estimates obtained via ensemble adjoint method. The difference between the direct estimate and the adjoint estimates is expected due to the above mentioned bias. We only show the component of the sensitivity associated with the parameter we vary, but in fact the adjoint method has output the sensitivities to all parameters, which similarly match with the true estimates.




Fig. 2. Inference of sensitivities of \bar{z} to parameters, s , r , and b . Top row (a,b,c) shows the change of \bar{z} for varying s , r , and b , for the true system, echo state network (ESN), and a polynomial fit on the true system's values. Bottom row (d,e,f) shows the derivative of \bar{z} with respect to the respective parameter obtained by the ensemble adjoint method using the true system's adjoint, ESN's adjoint, and by differentiating the polynomial fit.

5 Conclusions and future directions

We propose a data-driven method to obtain adjoint sensitivities of a system with respect to its parameters using a parameter-aware extension of echo state network. The fact that an echo state network (ESN) is a dynamical system, for which an adjoint system exists, and the success of ESNs to not only forecast predictions and but also infer invariant properties, motivate this approach. We demonstrate this approach on a prototypical chaotic system with three system's parameters. First, we train a parameter-aware ESN on data from different chaotic regimes of Lorenz system, and show its performance when running autonomously to make short-term time-accurate predictions and replicate long-term statistics of the chaotic attractor at regimes that have not been seen during training or validation. Second, using the derived adjoint equations of the ESN, we compute the sensitivity of a time-averaged state variable, \bar{z} , at varying parameter values to all parameters, s , r , and b . We find that the estimated sensitivities match closely with the sensitivities obtained by the adjoint of the original system. We observe a known bias [1] between the ensemble adjoint sensitivities and direct estimates calculated by fitting a polynomial to the objective function. This work opens possibilities for data-driven sensitivity analysis of chaotic flows without adjoint solvers. Future directions include analysis of the scalability of the method to higher-dimensional systems.

Acknowledgments. This research has received financial support from the ERC Starting Grant No. PhyCo 949388 and UKRI AI for Net Zero grant EP/Y005619/1. L.M. is also grateful for the support from the grant EU-PNRR YoungResearcher TWIN ERC-PI_0000005.

References

1. Chandramoorthy, N., Fernandez, P., Talnikar, C., Wang, Q.: An Analysis of the Ensemble Adjoint Approach to Sensitivity Analysis in Chaotic Systems. In: 23rd AIAA Computational Fluid Dynamics Conference. American Institute of Aeronautics and Astronautics, Denver, Colorado (Jun 2017). <https://doi.org/10.2514/6.2017-3799>
2. Doan, N., Polifke, W., Magri, L.: Physics-informed echo state networks. *Journal of Computational Science* **47**, 101237 (Nov 2020). <https://doi.org/10.1016/j.jocs.2020.101237>
3. Eyink, G.L., Haine, T.W.N., Lea, D.J.: Ruelle's linear response formula, ensemble adjoint schemes and Lévy flights. *Nonlinearity* **17**(5), 1867–1889 (Sep 2004). <https://doi.org/10.1088/0951-7715/17/5/016>
4. Giles, M.B., Pierce, N.A.: An Introduction to the Adjoint Approach to Design. *Flow, Turbulence and Combustion* **65**(3/4), 393–415 (2000). <https://doi.org/10.1023/A:1011430410075>
5. Grigoryeva, L., Ortega, J.P.: Echo state networks are universal. *Neural Networks* **108**, 495–508 (Dec 2018). <https://doi.org/10.1016/j.neunet.2018.08.025>
6. Lea, D.J., Allen, M.R., Haine, T.W.N.: Sensitivity analysis of the climate of a chaotic system. *Tellus A* **52**(5), 523–532 (Oct 2000). <https://doi.org/10.1034/j.1600-0870.2000.01137.x>

7. Lorenz, E.N.: Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences* **20**(2), 130–141 (Mar 1963). [https://doi.org/10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2)
8. 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*, vol. 7700, pp. 659–686. Springer Berlin Heidelberg, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35289-8_36
9. Magri, L.: Adjoint Methods as Design Tools in Thermoacoustics. *Applied Mechanics Reviews* **71**(2), 020801 (Mar 2019). <https://doi.org/10.1115/1.4042821>
10. Margazoglou, G., Magri, L.: Stability analysis of chaotic systems from data. *Nonlinear Dynamics* **111**(9), 8799–8819 (May 2023). <https://doi.org/10.1007/s11071-023-08285-1>
11. Pathak, J., Hunt, B., Girvan, M., Lu, Z., Ott, E.: Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach. *Physical Review Letters* **120**(2), 024102 (Jan 2018). <https://doi.org/10.1103/PhysRevLett.120.024102>
12. 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 (Dec 2017). <https://doi.org/10.1063/1.5010300>
13. Peter, J.E., Dwight, R.P.: Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches. *Computers & Fluids* **39**(3), 373–391 (Mar 2010). <https://doi.org/10.1016/j.compfluid.2009.09.013>
14. Racca, A., Magri, L.: Robust Optimization and Validation of Echo State Networks for learning chaotic dynamics. *Neural Networks* **142**, 252–268 (Oct 2021). <https://doi.org/10.1016/j.neunet.2021.05.004>
15. Roy, M., Mandal, S., Hens, C., Prasad, A., Kuznetsov, N.V., Dev Shrimali, M.: Model-free prediction of multistability using echo state network. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **32**(10), 101104 (Oct 2022). <https://doi.org/10.1063/5.0119963>
16. Xiao, R., Kong, L.W., Sun, Z.K., Lai, Y.C.: Predicting amplitude death with machine learning. *Physical Review E* **104**(1), 014205 (Jul 2021). <https://doi.org/10.1103/PhysRevE.104.014205>