
Inferring Stability Properties of Chaotic Systems on Autoencoders’ Latent Spaces

Elise Özalp

Department of Aeronautics
Imperial College London
elise.ozalp@imperial.ac.uk

Luca Magri

Department of Aeronautics
Imperial College London
The Alan Turing Institute
l.magri@imperial.ac.uk

Abstract

The data-driven learning of solutions of partial differential equations can be based on a divide-and-conquer strategy. First, the high dimensional data is compressed to a latent space with an autoencoder; and, second, the temporal dynamics are inferred on the latent space with a form of recurrent neural network. In chaotic systems and turbulence, convolutional autoencoders and echo state networks (CAE-ESN) successfully forecast the dynamics, but little is known about whether the stability properties can also be inferred. We show that the CAE-ESN model infers the invariant stability properties and the geometry of the tangent space in the low-dimensional manifold (i.e. the latent space) through Lyapunov exponents and covariant Lyapunov vectors. This work opens up new opportunities for inferring the stability of high-dimensional chaotic systems in latent spaces.




Figure 1: Illustration of the CAE-ESN with the Kuramoto-Sivashinsky data as an example.

1 Introduction

Convolutional autoencoder recurrent neural networks have become a promising approach for forecasting chaotic and turbulent systems [1–4]. This framework addresses the challenge of modelling high-dimensional chaotic systems by compressing the data into a nonlinear latent representation through a convolutional autoencoder (CAE) and a recurrent neural network to propagate the lower-dimensional temporal dynamics on the latent manifold. This addresses a challenge in chaotic partial differential equations (PDEs) and turbulent systems: although the state might live in a high dimensional space; in the asymptotic limit, the solution converges to a lower dimensional attractor.

The chaotic properties of the attractor are characterized by invariant measures, such as Lyapunov exponents, the Kaplan-Yorke dimension, and the geometry of the tangent space. Although it is possible to infer these measures through echo state networks (ESNs) [5] and long short-term memory networks [6], these methods for data-driven stability analysis do not effectively apply to high-

dimensional data. Furthermore, the stability properties of hybrid CAE-RNN models, specifically of the latent manifolds generated by autoencoders, remain unexplored.

In this paper, we address this gap by employing ESNs, a reservoir computer, as recurrent neural networks [7] and applying stability analysis to investigate and characterize latent manifolds. We thus extend the CAE-ESN methodology [3] by performing data-driven stability analysis on the latent space of a convolutional autoencoder of a spatiotemporal PDE, i.e., the Kuramoto-Sivashinsky system.

2 Learning stability properties in latent spaces

The Kuramoto-Sivashinsky (KS) equation is a prototypical partial differential equation (PDE) to study chaos in spatial-dissipative systems [8]. The PDE models instabilities and flame fronts and is given by

$$\mathbf{u}_t + \mathbf{u}_{xx} + \mathbf{u}_{xxxx} + \mathbf{u}\mathbf{u}_x = 0 \quad (1)$$

with periodic boundary conditions on the spatial domain $[0, L]$. To solve the PDE, we spatially discretize the domain into N_x grid points. This discretization allows us to interpret Eq. (1) as a dynamical system in the form of $\mathbf{u}_t = f(\mathbf{u})$ with f being a smooth function and $\mathbf{u} \in \mathbb{R}^{N_x}$.

As the domain length L increases, the solution of Eq. (1) transitions from a quasiperiodic system to a chaotic solution [9]. This behaviour can be analysed via stability analysis, which focuses on the tangent space defined by the Jacobian $J = \partial f / \partial \mathbf{u}$ to examine the chaotic stability. By imposing an infinitesimal perturbation $\delta \mathbf{u}$ to the trajectory of the system and evolving it after Eq. (1), we obtain the linearized tangent equation

$$\delta \mathbf{u}_t = J(\mathbf{u}(t))\delta \mathbf{u}. \quad (2)$$

For a detailed explanation of tangent space methods, we refer the reader to [10, 11] and briefly summarize the key properties that can be derived from Eq. (2) here.

Chaotic systems are characterized by exponential sensitivity to perturbations, a property which can be measured through Lyapunov exponents (LEs). Each Lyapunov exponent $\lambda_1, \dots, \lambda_{N_x}$ measures the average rates of convergence and divergence of nearby trajectories in the system's phase space \mathbb{R}^{N_x} . When the system is chaotic, the solution \mathbf{u} occupies only a small subset \mathcal{A} of the phase space in \mathbb{R}^{N_x} . This subset, known as the attractor, has a fractal structure, the dimension of which can be estimated through the Kaplan-Yorke dimension D_{KY} [12], among other dimensionality estimates [13, 14].

Further insight into the system's geometry is provided by the covariant Lyapunov vectors (CLVs) $\mathbf{v}_1, \dots, \mathbf{v}_{N_x}$ [10, 15], which offer an invariant splitting of the tangent space into unstable, neutral, and stable subspaces, corresponding to positive, zero, and negative Lyapunov exponents, respectively. In hyperbolic systems, the angles between CLVs from different subspaces are bounded away from zero. The angle $\theta_{\mathbf{v}_i, \mathbf{v}_j}$ between a pair of CLVs \mathbf{v}_i and \mathbf{v}_j from different subspaces is calculated as

$$\theta_{\mathbf{v}_i, \mathbf{v}_j} = \frac{180^\circ}{\pi} \cos^{-1}(|\mathbf{v}_i \cdot \mathbf{v}_j|), \quad (3)$$

where \cdot denotes the dot product. For spatial-dissipative systems such as the KS system, the CLV angles also allow for the isolation of physically relevant modes that capture the essential dynamics of the system. The number of these physically relevant modes is typically larger than the Kaplan-Yorke dimension but still significantly smaller than the number of grid points. From a theoretical perspective, it has been conjectured that capturing the physical modes is sufficient to faithfully integrate PDEs [16]. This provides a guideline for the design of an optimal reduced order model, which should capture all the time-invariant properties of the attractors, i.e. the LEs and the Kaplan-Yorke dimension, but also reproduce the geometric structure of the attractor through the CLVs.

3 Methods

For dissipative chaotic systems such as in Eq. (1), the state space may be high-dimensional, but in the asymptotic limit, the solution converges to a lower-dimensional chaotic attractor. This work follows a two-step approach: first, we compute the low-dimensional manifold with a convolutional autoencoder (CAE) for a nonlinear reduced-order representation, and second, we propagate the latent dynamics using an echo state network (ESN), see Fig. 1.¹

¹A tutorial is available on GitHub: github.com/MagriLab/LatentStability. The training is performed on a single Quadro RTX 8000.

3.1 Convolutional autoencoder

Let $\{\mathbf{u}(t_i, \mathbf{x})\}_{i=0,1,\dots}$ be the discretized solution to Eq. (1). We first aim to find a lower-dimensional manifold representation of the system using a CAE. The encoder \mathcal{E} maps the physical input $\mathbf{u}(t_i) \in \mathbb{R}^{N_x}$ to a latent representation $\mathbf{y}(t_i) \in \mathbb{R}^{N_{lat}}$. The decoder \mathcal{D} then maps the physical state from the latent representation to the phase space

$$\hat{\mathbf{u}}(t_i) \approx \mathbf{u}(t_i), \quad \text{where } \hat{\mathbf{u}}(t_i) = \mathcal{D}(\mathbf{y}(t_i)), \quad \mathbf{y}(t_i) = \mathcal{E}(\mathbf{u}(t_i)), \quad (4)$$

and $\hat{\mathbf{u}}(t_i)$ is the autoencoder reconstruction. The network is trained by minimizing the loss $\mathcal{L}(\mathbf{u}, \hat{\mathbf{u}}) = \frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} \|\mathbf{u}(t_i) - \hat{\mathbf{u}}(t_i)\|_2^2$, where N_{tr} is the number of training samples. By first applying an autoencoder to each snapshot, we obtain a spatial compression of the input, with the convolutional layers filtering the spatial multiscale structure of the data. The objective is to find a latent space that approximately retains the geometry of the attractor while providing a minimal representation.

3.2 Echo state network

The trained autoencoder provides a transformation from the physical snapshot $\mathbf{u}(t_i)$ to the latent representation $\mathbf{z}(t_i)$ and vice versa. To learn the temporal dynamics of the training data, we employ the ESN [7] on the time-ordered latent representations $\{\mathbf{y}(t_i)\}_{i=1,\dots,N_{tr}}$, resulting in the CAE-ESN. The ESN evolves the latent representations $\mathbf{y}(t_i)$ to the next-time step prediction $\hat{\mathbf{y}}(t_{i+1})$ according to

$$\mathbf{r}(t_{i+1}) = \tanh([\mathbf{y}(t_i); 1]^T \mathbf{W}_{in} + \mathbf{r}(t_i)^T \mathbf{W}), \quad (5)$$

$$\hat{\mathbf{y}}(t_{i+1}) = [\mathbf{r}(t_{i+1}); 1]^T \mathbf{W}_{out}. \quad (6)$$

Here, the matrices \mathbf{W}_{in} and \mathbf{W} are pseudorandomly generated and fixed [17] and σ_{in} is found by Bayesian optimization [18]. By formulating the loss as $\mathcal{L}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} \|\mathbf{y}(t_i) - \hat{\mathbf{y}}(t_i)\|_2^2$, the ESN training is performed by solving for \mathbf{W}_{out} using ridge regression, speeding up the training by avoiding backpropagation. The ESN can be employed in two modes, open-loop and closed-loop. Whilst training, the network operates in open-loop mode, predicting the next time step based on a reference input. Once \mathbf{W}_{out} has been fixed, the network can be employed in the closed-loop configuration where the network prediction is used as input in Eq. (5), allowing the network to autonomously evolve without additional input data.

This setup effectively defines a dynamical system, and the stability properties of the network can be calculated [5]. The first step is to calculate the Jacobian of the ESN, $\mathbf{J}_{esn}(\mathbf{r}(t_{i+1})) = (1 - \mathbf{r}(t_i)^2)(\mathbf{W}_{in}^T \mathbf{W}_{out}^T + \mathbf{W}^T)$. The Jacobian \mathbf{J}_{esn} is then employed in the tangent equation, Eq. (2), to calculate the LEs and CLVs [5]. The ESN predicts the temporal dynamics on the latent manifold and the Jacobian, therefore, defines the tangent space of the latent manifold. Consequently, the calculated stability properties are inferred on the latent manifold, not in the full physical space.

4 Numerical tests

We consider the KS equation with $L = 22$ for which the system is chaotic with an attractor dimension of $D_{KY} = 6.007$ and a largest Lyapunov exponent of 0.045 [19]. In related work [20, 21], the number of physical modes is determined to be 8, providing an estimate for the minimal autoencoder latent dimension. We generate the reference solution of the Kuramoto-Sivashinsky equation with $N_x = 512$ for a duration of $T = 10^5$ and $\Delta t = 0.05$.

First, we train a convolutional autoencoder to map the full state $\mathbf{u}(t_i)$ with $N_x = 512$ to the latent representation $\mathbf{z}(t_i)$ with dimension $N_{lat} = 8$. After training the autoencoder, we train the ESN on the latent space. In Fig. 2(a), we present the autonomous prediction of the CAE-ESN based on the reduced representation. Despite using less than 2% of degrees of freedom, the CAE-ESN accurately forecasts the evolution for about $2LT$ ². After $2LT$ the trajectories slowly diverge, reflecting the inherent sensitivity of chaotic systems to perturbations.

More insight into the physical behaviour of the CAE-ESN is obtained by investigating the invariant properties of the network. In Fig. 2(b), we show the convergence of the LEs of the CAE-ESN over

²One Lyapunov time (LT) is defined by λ_1^{-1} and denotes a characteristic time scale at which two trajectories of a chaotic system diverge.




Figure 2: (a) Autonomous prediction of the CAE-ESN on the latent manifold of dimension $N_{lat} = 8$. (b) Convergence of the Lyapunov exponents of the CAE-ESN over $100LT$.

$100LT$. To assess the robustness, we take an ensemble of 10 ESNs to analyse the stability properties in the latent space. In Fig. 3(a), we compare the reference LEs (in black squares) with the mean LEs obtained from 10 ESNs (in red circles). The mean Kaplan-Yorke Dimension of the CAE-ESN is given by $6.008 \pm 5 \cdot 10^4$ compared to the reference of $D_{KY} = 6.007$. The CAE-ESNs correctly reproduce the first 10 LEs, which correspond to the physical behaviour of the system, indicating that the forecast based on the reduced representation captures all essential chaotic dynamics.




Figure 3: (a) The first 10 Lyapunov exponents of the dynamical system (black squares) compared to the Lyapunov exponents across 10 different ESNs (red dots) after $100LT$. The shaded area indicates the standard deviation. (b - d) The angle distribution of the Kuramoto-Sivashinsky system for leading covariant Lyapunov vectors of the (b) unstable-neutral, (c) unstable-stable and (d) neutral-stable Lyapunov exponents.

We further examine the CAE-ESN abilities to capture the geometric structure by studying the angles between CLVs in Fig. 3(b - d). The agreement of angle statistics between the reference and the CAE-ESN is within a negligible numerical error with a Wasserstein distance of 0.001, demonstrating that the geometrical structure of the attractor has been effectively captured in the latent manifold.

5 Conclusions

We propose a data-driven method to infer the stability properties of chaotic partial differential equations. This is based on the convolutional autoencoder echo state network (CAE-ESN), where the CAE discovers a nonlinear latent representation of minimal dimension, on which the ESN propagates the temporal dynamics. By applying the CAE-ESN to the chaotic Kuramoto-Sivashinsky equation, we show that (i) invariant measures, such as Lyapunov exponents and the Kaplan-Yorke dimension, are accurately and robustly inferred in the latent space, (ii) the geometric properties of the attractor, including the angles of the covariant Lyapunov vectors (CLVs), are accurately learned. This approach enables both nonlinear forecasting and stability analysis of chaotic partial differential equations. Inferring stability properties of latent manifolds generated by a CAE offers a perspective on understanding autoencoder manifolds using dynamical systems tools.

Acknowledgments

This research has received financial support from the EPSRC Grant No. EP/W026686/1 and from the ERC Starting Grant No. PhyCo 949388.

References

- ¹A. J. Linot and M. D. Graham, “Data-driven reduced-order modeling of spatiotemporal chaos with neural ordinary differential equations”, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **32** (2022).
- ²P. R. Vlachas, G. Arampatzis, C. Uhler, and P. Koumoutsakos, “Multiscale simulations of complex systems by learning their effective dynamics”, *Nature Machine Intelligence* **4**, 359–366 (2022).
- ³A. Racca, N. A. K. Doan, and L. Magri, “Predicting turbulent dynamics with the convolutional autoencoder echo state network”, *Journal of Fluid Mechanics* **975**, A2 (2023).
- ⁴K. Fukami, T. Nakamura, and K. Fukagata, “Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data”, *Physics of Fluids* **32** (2020).
- ⁵G. Margazoglou and L. Magri, “Stability analysis of chaotic systems from data”, *Nonlinear Dynamics* **111**, 8799–8819 (2023).
- ⁶E. Özalp, G. Margazoglou, and L. Magri, “Reconstruction, forecasting, and stability of chaotic dynamics from partial data”, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **33** (2023).
- ⁷H. Jaeger, “The “echo state” approach to analysing and training recurrent neural networks—with an erratum note”, Bonn, Germany: German National Research Center for Information Technology GMD Technical Report **148**, 13 (2001).
- ⁸Y. Kuramoto, “Diffusion-Induced Chaos in Reaction Systems”, *Progress of Theoretical Physics Supplement* **64**, 346–367 (1978).
- ⁹D. T. Papageorgiou and Y. S. Smyrlis, “The route to chaos for the kuramoto-sivashinsky equation”, *Theoretical and Computational Fluid Dynamics* **3**, 15–42 (1991).
- ¹⁰F. Ginelli, P. Poggi, A. Turchi, H. Chaté, R. Livi, and A. Politi, “Characterizing dynamics with covariant lyapunov vectors”, *Physical Review Letters* **99**, 130601 (2007).
- ¹¹M. Sandri, “Numerical calculation of lyapunov exponents”, *Mathematica Journal* **6** (1996).
- ¹²P. Frederickson, J. L. Kaplan, E. D. Yorke, and J. A. Yorke, “The liapunov dimension of strange attractors”, *Journal of differential equations* **49**, 185–207 (1983).
- ¹³S. H. Strogatz, *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering* (CRC press, 2018).
- ¹⁴P. V. Kuptsov and U. Parlitz, “Theory and computation of covariant lyapunov vectors”, *Journal of nonlinear science* **22**, 727–762 (2012).
- ¹⁵F. Huhn and L. Magri, “Stability, sensitivity and optimisation of chaotic acoustic oscillations”, *Journal of Fluid Mechanics* **882**, A24 (2020).
- ¹⁶K. A. Takeuchi, H.-I. Yang, F. Ginelli, G. Radons, and H. Chaté, “Hyperbolic decoupling of tangent space and effective dimension of dissipative systems”, *Physical Review E* **84**, 046214 (2011).
- ¹⁷M. Lukoševičius, “A practical guide to applying echo state networks”, in *Neural networks: tricks of the trade: second edition* (Springer, 2012), pp. 659–686.
- ¹⁸A. Racca and L. Magri, “Robust optimization and validation of echo state networks for learning chaotic dynamics”, *Neural Networks* **142**, 252–268 (2021).
- ¹⁹P. Gupta, P. J. Schmid, D. Sipp, T. Sayadi, and G. Rigas, “Mori-zwanzig latent space koopman closure for nonlinear autoencoder”, arXiv preprint arXiv:2310.10745 (2023).
- ²⁰P. Cvitanović, R. L. Davidchack, and E. Siminos, “On the state space geometry of the kuramoto-sivashinsky flow in a periodic domain”, *SIAM Journal on Applied Dynamical Systems* **9**, 1–33 (2010).
- ²¹A. J. Linot and M. D. Graham, “Deep learning to discover and predict dynamics on an inertial manifold”, *Physical Review E* **101**, 062209 (2020).