Chaos, Solitons and Fractals 176 (2023) 114115

Contents lists available at ScienceDirect

Chaos, Solitons and Fractals

journal homepage: www.elsevier.com/locate/chaos

The reconstruction of flows from spatiotemporal data by autoencoders

Facundo Fainstein [a,b,1], Josefina Catoni [a,c,1], Coen P.H. Elemans [d], Gabriel B. Mindlin [a,b,e,*]

[a] Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física, Ciudad Universitaria, 1428 Buenos Aires, Argentina
[b] CONICET Universidad de Buenos Aires, Instituto de Física Interdisciplinaria y Aplicada (INFINA), Ciudad Universitaria, 1428 Buenos Aires, Argentina
[c] Research Institute for Signals, Systems and Computational Intelligence sinc(i) (FICH-UNL/CONICET), Santa Fe, Argentina
[d] Department of Biology, University of Southern Denmark, 5230 Odense M, Denmark
[e] Departamento de Matemática Aplicada, Universidad Rey Juan Carlos, Madrid, Spain

ARTICLE INFO

Keywords:
Dynamical systems
Chaos
Data driven analysis
Autoencoders
Spatiotemporal data

ABSTRACT

Artificial neural networks have become essential tools in data science for uncovering insights from complex data. However, they are usually seen as black boxes. In this work we explore how an autoencoder processes complex spatiotemporal information. We analyze the topological structure of reconstructed flows in the latent space of an autoencoder for two distinct test cases. The first case involves a synthetic spatiotemporal pattern for the temperature field in a convective problem, illustrating a classic extended system that exhibits low-dimensional chaos. The second case focuses on an experimental recording of the labial oscillations responsible for sound production in an avian vocal organ, as an example of periodic dynamics in a biological system. We find that the state representation in its latent space can be topologically equivalent to the phase space of the problem. Autoencoders thus retain phase space representations of the data hidden in its latent layer.

1. Introduction

It is a core approach in many disciplines of the natural sciences, to construct minimal models that account for the temporal evolution of variables that describe the system under study. The variables are interpretable magnitudes and, in the best cases, entirely measured. Time evolution rules seek to be written in terms of simple mechanisms, whenever possible based on previous theoretical frameworks (e.g., Newton's laws for a mechanical problem, or Maxwell's laws for an electromagnetic one). In recent years, a different strategy has emerged in the natural sciences, broadly known as “data driven”. This approach foregoes the need to pre-identify relevant variables and, in many cases, does not seek to elucidate minimal mechanisms. Typically, these strategies are applied when it is necessary to extract information from a massive amount of data, and the objective is focused on the ability to carry out predictions. However, these strategies should not be viewed as antagonistic. In fact, they can be utilized synergistically, as far as we understand the relationship between both perspectives. In this paper we propose to address this unification, within the framework of a particular strategy in data science: the way in which an autoencoder processes information obtained by filming the dynamics of an extended system. We intend to test the hypothesis that the predictive capacity of a neural network with autoencoder architecture stems from generating a representation of the dynamics in its latent space that is equivalent to the phase space of the problem.

A strategy to address this issue was presented by Champion et al. [1], who introduced a data driven method to discover low-dimensional models using an autoencoder. Autoencoders are a type of deep artificial neural networks that usually have a symmetrical layered structure and are designed to highly efficiently encode input data in a low-dimensional space [2,3]. The minimal architecture consists of two layers, with the same number of units (input and output), and a smaller layer in between. The network is trained to copy the input at the output. The space defined by the activation values of the units in the middle layer is called the latent space. The possibility to map each input to the output requires that, in the latent space, each state is uniquely represented. In fact, the dimension of the state representation is chosen by determining the minimum dimension of the latent space that still allows the reconstruction of all the inputs. If each point in this space corresponds to a state, each point will have a unique future. Therefore, the evolution of the latent variables of the problem gives rise to a flow which may be modeled by a dynamical system. In dynamical systems theory,

* Corresponding author at: Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Física, Ciudad Universitaria, 1428 Buenos Aires, Argentina.

E-mail address: gabo@df.uba.ar (G.B. Mindlin).

[1] These authors contributed equally to this work.

https://doi.org/10.1016/j.chaos.2023.114115
Received 28 August 2023; Accepted 25 September 2023
Available online 3 October 2023

the flow can be defined as the set of solutions of the system. In their work, Champion et al. [1] state that if restrictions are imposed on the vector field (sparseness), the dynamical system reconstructed from the trajectories in the latent space can constitute normal forms for the original equations, a strategy applicable to systems close to local bifurcations. However, studies in the field of topology of dynamical systems impose restrictions on representations of the data [4]. The topological organization of the flow acts as a fingerprint: if the representation fails to capture the topology of the phase space, any fitted model must be rejected. The question then naturally arises: is in general the topology of the reconstructed flow in the latent space equivalent to the original topology of the phase space in which the dynamics of the problem live?

Recent studies on autoencoders and recurrent neural networks support this view. We recently analyzed the topology of a reconstructed flow using an autoencoder [5]. The Rössler system of differential equations was integrated, and the autoencoder was trained with segments of the time series of one of the scalar variables [6]. Each segment was represented by a point in the latent space, and consecutive segments gave rise to a flow in this space. In those numerical experiments, the success of reproducing each input segment to an output segment corresponded to cases in which the encoder was injective. Then, segments of the original system that constitute good approximations to periodic orbits were extracted. It was found that in the reconstructed flow, the segments showed the same topological organization. Moreover, another widely used architecture for temporal signal predictions, recurrent neural networks, was examined from the same perspective [7]. In that study, a network was trained with segments of scalar time series, and it was found that the reconstructed flows in its hidden units displayed the same topology as the original flow.

In this work, we aim to extend this line of research to the case in which the original data comes from a movie. This is one of the most direct experimental ways in which the dynamics of an extended problem can be presented to us. Moreover, in contrast to the case for a scalar time series, the processing of a raw spatiotemporal recording exposes a fundamental feature of data science: we do not propose a priori variables responsible for providing a minimal description of the problem's dynamics, instead we leave the task to the network. To address this general objective, we pose two problems from two different disciplines of the natural sciences where the use of this experimental approach is prevalent. First, we analyze a synthetic movie motivated by a classical problem in fluid dynamics: the evolution of the temperature field in the convection model studied by Lorenz [8]. Starting with the Lorenz equations and integrating them numerically, we obtain the variables that serve as amplitudes of a set of spatial modes. These structures are used to generate the movie to be analyzed. We aim to test the hypothesis that the reconstructed topology in the latent space corresponds to the known Lorenz flow [4,9]. Furthermore, autoencoders have already been used for many problems in this field [10]. Second, we analyze an experimental biological data set. We utilize experimental data capturing the labial dynamics in the vocal apparatus of a bird. The avian phonation problem was theoretically studied under the hypothesis that the syringeal labia's dynamics can be modeled by a relaxation oscillator. Processing the data from the experimental movie, we quantify the distance between successive images, revealing the presence of relaxation dynamics. Subsequently, we test the hypothesis that this dynamics is reflected in the latent space.

2. Analysis of a chaotic spatiotemporal pattern

2.1. Pattern synthesis

We generate a synthetic movie motivated by the atmospheric convection problem studied by Lorenz [8]. In this seminal work, Lorenz studies the flow in a thin layer of fluid between two horizontal surfaces, when a temperature difference is maintained between them. He assumes a complete translational symmetry in one of the horizontal directions (y), and periodicity in the other (x). He combines the Navier-Stokes equations, the continuity equation, and a diffusion-convection heat equation for the temperature profile [11]. Expressing the fields in terms of a stream function and θ, the departure of the temperature from the non-convective steady state, Lorenz proposes a modal decomposition for the fields. The expansion for temperature has the following form:

$$\theta(x,z,t) = \alpha_1 Y(t) cos\left(\frac{\pi a x}{h}\right) sin\left(\frac{\pi z}{h}\right) - \alpha_2 Z(t) sin\left(\frac{2\pi z}{h}\right) \tag{1}$$

Where z is the vertical coordinate, α_1 and α_2 are constants, h is the height of the fluid, a the aspect ratio, and $Y(t)$ and $Z(t)$ are functions of time that act as amplitudes for each spatial mode (which we call $\psi_1(x,z)$ and $\psi_2(x,z)$ respectively). In Fig. 1(a) we show these two spatial structures. Substituting these expressions into the partial differential equations and ignoring all high-order terms in the trigonometric functions, Lorenz writes his famous dynamical system for the amplitude of the modes:

$$\begin{aligned} \frac{dX}{dt} &= \sigma(Y - X) \\ \frac{dY}{dt} &= rX - Y - XZ \\ \frac{dZ}{dt} &= XY - bZ \end{aligned} \tag{2}$$

where X is the amplitude of the stream function's mode. By numerically integrating this dynamical system with $\sigma = 10$, $b = \frac{8}{3}$ and $r = 28$, we generate temporal series with 40,000 points (time step of 0.01). We discard the first 1000 points to avoid the transient. In Fig. 1(b) we show a segment of the time series. Using these series and Eq. (1), we obtain the spatiotemporal pattern for the temperature that would correspond to the convection problem if the formulation was exact (Fig. 1(c)). We use a spatial discretization of 0.025, so that each image has 40×40 pixels, we take $\alpha_1 = 2\alpha_2 = \frac{1}{20}$ and we add gaussian noise to each pixel (null mean value and a standard deviation of 0.01).

2.2. The neural network

Fig. 1(d) shows the typical structure of an autoencoder as the one used in this work. We analyze the movie using individual frames as inputs. Through an encoder, the network maps each frame to a point in the latent space, and then decodes it to an output frame. During training, the network's weights are optimized to minimize the mean squared error (MSE) between the input and output frames. If the minimization procedure finds parameters that achieved a null difference, the encoder is an injective function, since different inputs would otherwise be decoded to the same output. In other words, a minimization of the MSE implies a representation in latent space without self-intersections of the trajectory that represents the consecutive states of the data set. We seek to assess whether this latent space representation preserves the topological structure of the underlying flow.

To analyze the movie, we partition the frames into two sets. The first 30,000 data frames are used for training the network, while the remaining 9000 for testing it. We compute the temporal average frame and subtract it from each sample in the data set. Each frame consists of 40×40 pixels, resulting in initial and final fully connected layers of the network with 1600 units. Intermediate fully connected layers containing 64, 32, and 16 units encode the frame to a three-dimensional latent space. ReLu activation functions are used for all layers except the middle and output layers, where no activation functions are used. The network is implemented in Keras version 2.6.0, backened by Tensorflow 2.6.0. The training process consists of batches of 512 samples and 600 epochs. The Adam algorithm with learning rate 0.001, $\beta_1 = 0.9$ and $\beta_2 = 0.999$, is used to minimize the MSE. We do not perform any systematic search in

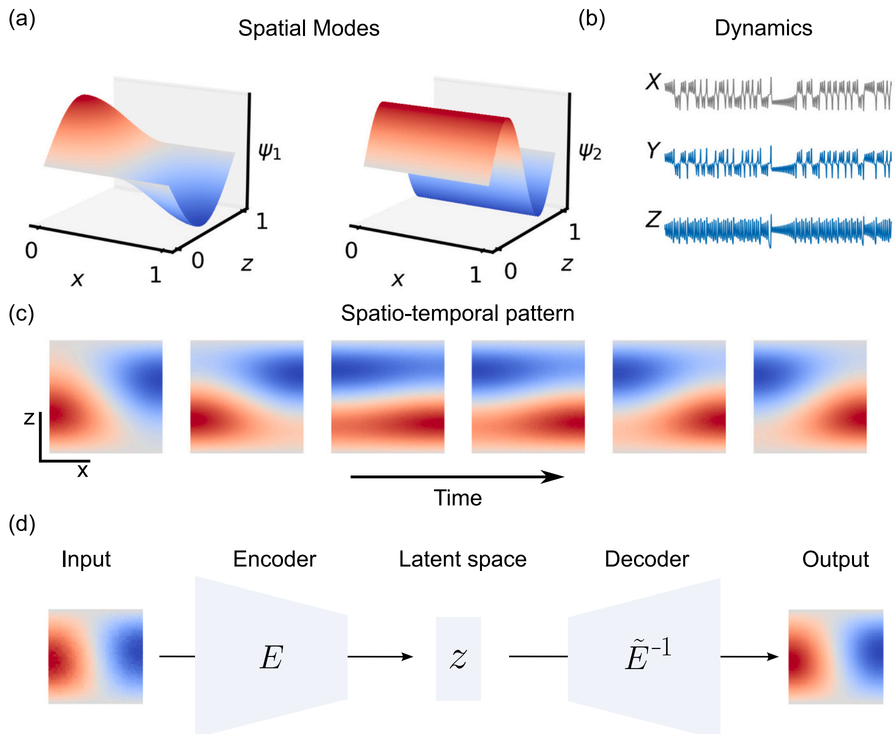

Fig. 1. Chaos in a spatiotemporal pattern describing an atmospheric convection problem. (a) The spatial structures that expand the temperature field. We use $h = 1$, $a = 1$. (b) The dynamics of the variables that serve as amplitudes of the modes. Blue time series are used, as indicated in the text. (c) Six frames of the synthetic movie for the temperature field. (d) Schematic of the architecture of an autoencoder. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the hyperparameter space to optimize the MSE in the test set.

2.3. Topological analysis of the underlying dynamics

The flux generated by integrating the Lorenz equations evokes a mask: the trajectories seem constrained to explore, for long periods of

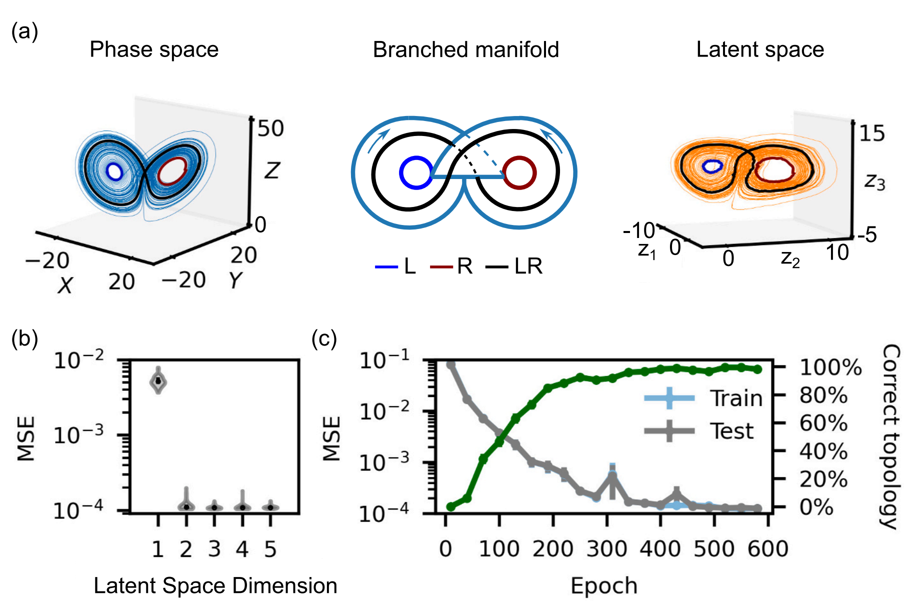

Fig. 2. Autoencoders preserve the topological organization of the flow. (a) Phase space of the Lorenz dynamical system (left). Branched manifold supporting all the unstable periodic orbits coexisting with the attractor (middle). Latent space representation of the spatiotemporal pattern for the test data set (right). (b) MSE of the test data set for training sessions with different number of units in the middle layer. We performed 80, 72, 170, 67 and 56 training sessions for the five dimensions shown, respectively. (c) Evolution of the MSE for the test data set using a three-dimensional latent space, and percentage of autoencoders with correct topology as functions of the training epochs. The mean and standard error are shown.

time, a two-dimensional manifold, avoiding hole-shaped areas of phase space (Fig. 2(a)). Formally, the Lorenz flow cannot live in two dimensions: it is not possible to have attractors more complex than limit cycles in a two-dimensional space, as we know from Poincare Bendixon's theorem [12]. However, Birman and Williams showed that all the periodic orbits of a flow generated by the Lorenz equations can be isotopically mapped to a branched manifold preserving its topological structure [9]. The branched manifolds are a generalization of the differential manifolds, with singularities in bounded areas, which admit a tangent space at each point. The manifold capable of supporting all the orbits of a Lorenz system is illustrated in the middle panel of Fig. 2(a).

This qualitative description can be transformed into a quantitative one, since a branched manifold like the one illustrated in Fig. 2(a) can be described algebraically. This description is carried out using two algebraic objects: the so-called topological matrix T, and the joining array A [4]. The elements T_{ij} describe how the branches are crossed if $i \neq j$, and T_{ii} describes the torsion of the $i-$ branch. The joining array describes the order in which the branches are joined at branch lines. This algebraic description provides an enormous power of synthesis: little information accounts for the topological organization of the branches capable of sustaining the set of all the periodic orbits of the flow. In this way, two inequivalent flows, such as Smale's and Lorenz's, will exhibit sets of knotted orbits in branched manifolds described by different invariants. Therefore, reconstructing these objects constitutes the first step in the topological description of a flow.

The way in which two periodic orbits are linked in three-dimensional phase space can be characterized by a topological invariant known as the linking number. Given two oriented curves (C_1 and C_2), the invariant can be computed algorithmically from a two-dimensional projection of these [11]. The algorithm consists of finding the crossings between both curves. For each crossing, the tangent vector $\overline{v_i}$ to each curve C_i in the direction of the flow is computed. The sign function is evaluated at $\overline{v_1} \times \overline{v_2}$, if C_1 is above C_2 at the crossing point (otherwise the product is inverted). Each crossing contributes with +½ o −½ depending on the sign obtained and the invariant is the resulting sum for all the crossings. To compute the linking number, we use a function developed for Python [5]. For our analysis, we separate segments of the movie that correspond to good approximations of the unstable orbits that coexist with the attractor. We do this by taking two Poincaré sections ($z = 27$ and $y < -10$; $z = 27$ and $y > 10$) and finding the best approximations for the fixed points of the map. In this way, we find the two period one orbits (called L and R) and the period two orbit (called LR). The approximations of the orbits correspond to segments of the movie with frames of the test data set. These three orbits are not linked, so the linking number between each of these pairs is zero.

2.4. The dimensionality of the latent space and the topology of the reconstructed flow

In the panel on the right of Fig. 2(a), we show the flow obtained in the latent space from the encoding of each frame of the test set, using one of the trained networks. Remarkably, we observe that even when training the network with individual frames, they are mapped into a set that exhibits significant similarities with phase space. Moreover, upon analyzing the periodic orbits we find that none of them are knotted, so the organization is equivalent to that of the original system.

We explore the network's performance depending on the dimension of the latent space by performing numerous training sessions while varying the number of units in the latent layer. In Fig. 2(b), we show the MSE distributions for the test data set at the best epoch. We find a limited predictive power of the autoencoder when the dimension of the latent space is one. However, with two units (and beyond), we achieve a reasonable reconstruction. This is because although the dynamical system governing the evolution of the spatiotemporal pattern is three-dimensional, the flow spends extended periods exploring a two-dimensional manifold. In fact, the branched manifold consists of a set of two two-dimensional manifolds that are joined together in a one-dimensional curve (Fig. 2(a) middle panel). This is reflected in the fact that measures of the dimension of the attractor are close to two, for example the correlation dimension is 2.05 [13]. Therefore, the predictive ability of the autoencoder depends on the dimensionality of the underlying dynamical system.

We assess the topological structure of the reconstructed flow in the latent space for 170 training sessions using networks with a three-dimensional latent space. For each training epoch, we compute the linking number between the three periodic orbits. In Fig. 2(c) we show the MSE and the percentage of autoencoders with correct topology, as functions of the training epoch. Initially, the MSE is relatively high, and the topological structure is often not correct. In fact, in 38 of the autoencoders we find self-intersections of the flow in this stage. However, at the end of the training process 169 out of the 170 autoencoders presented a correct topology. Thus, our analysis shows that in a synthetic, well-studied chaotic extended system, the latent space of the autoencoder retains the topology of the dynamical system.

3. Low dimensional relaxation dynamics in avian phonation

In the previous section, we delved into the structure of reconstructed flows in the latent space of an autoencoder when analyzing synthetic chaotic spatiotemporal data from a fluid dynamical problem. In this section, we extend our study for an experimental recording of a biological system. We use this approach because the analysis of movies is widespread within experimental biology, spanning from microscopy to non-invasive techniques in animal behavior studies. Furthermore, this is a problem for which the dynamics are hypothesized to be simple, encompassing just periodic solutions [14]. The study of this problem thus serves as a representative example, ubiquitous in the natural sciences, of an extended system exhibiting oscillatory dynamics. Beyond testing the reconstruction of a periodic orbit in the latent space of an autoencoder, we are seeking to identify the type of oscillator. Specifically, we investigate the presence of temporal heterogeneities in the representation linked to the dynamics of a relaxation oscillator, as posited by earlier theoretical studies of this problem [14].

3.1. Direct observations of labial oscillations in the avian vocal organ

We analyze a movie of oscillating structures during sound production in the avian vocal organ, the syrinx, in an ex vivo preparation [15]. The syrinx produces sound primarily through oscillations of tissue folds called labia that open and close the air passage. To describe the transfer of kinetic energy of air to labial oscillations, one of the simplest models assumes that the labia support both lateral oscillations and an upward propagating surface wave [14]. Moreover, the assumption that these two modes participate in a relaxation oscillator has played a crucial role in explaining several acoustic features of birdsong, such as a departure of tonality due to harmonics in the sound [16,17]. The movie is part of an experimental program aimed at directly validating these hypotheses.

In the experiment, air is blown throughout the excised syrinx of a pigeon inducing membrane oscillations [15,18]. The details of the experimental procedure can be found in Ref. 18. In Fig. 3(a) we show six frames from the movie. The complete record has 400 frames, and one period of the syringeal dynamics is approximately covered every 29 frames.

3.2. Relaxation dynamics in the latent space

To project the dynamics to a low-dimensional latent space, we design an autoencoder with the same architecture and loss function as described in the previous section. The differences are that now the frames have 90×110 pixels, so the first and final layers have 9900 units, and we use a two-dimensional latent space. The training process consists of batches of 16 samples and 200 epochs. The first 75 % of the data is

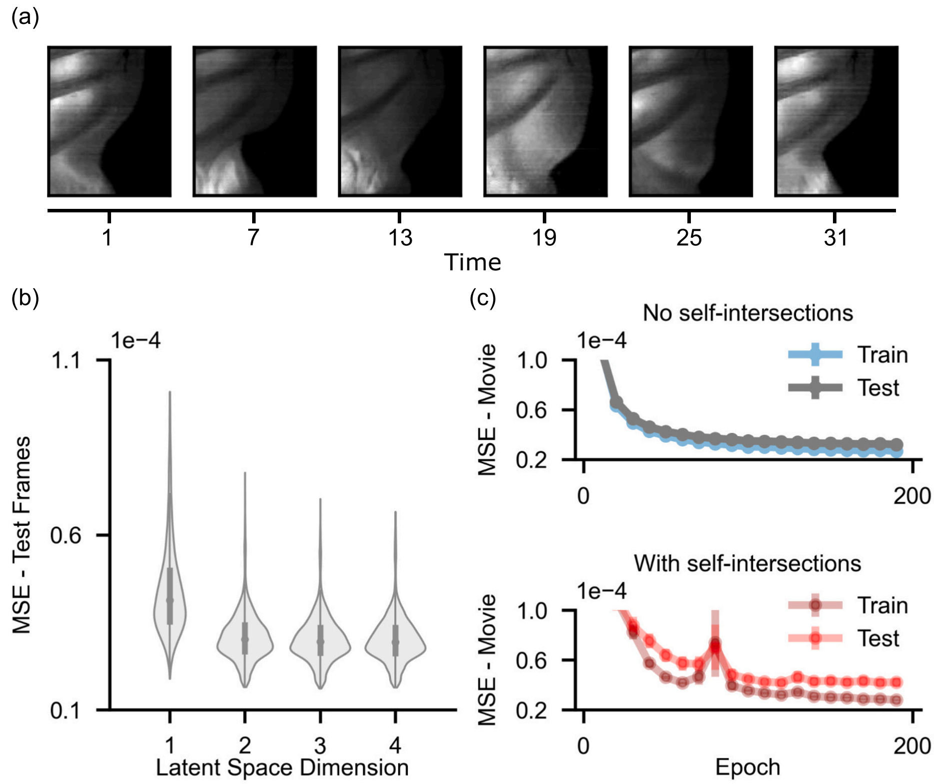

Fig. 3. The experimental movie showing labial oscillations in the avian vocal organ. (a) Six frames are displayed. (b) The mean squared error does not diminish significantly for latent spaces of dimension larger than two. (c) We fitted 100 models, out of which 94 runs lead to non-self-intersecting trajectories in the latent space. The convergence of the MSE, in average, as a function of the epoch is displayed in the upper panel. The average of the MSE for the six cases leading to self-intersecting trajectories is shown in the bottom panel. The error bars show the standard errors.

used to train the network, and the last 25 % to test it. We compute the temporal average frame and subtract it from all samples in the set.

The dimensionality of the latent space is chosen based on the calculations displayed in Fig. 3(b). We conduct 100 fittings for networks with different numbers of units in the middle layer. Interestingly, increasing the dimensionality of the latent space beyond two does not significantly reduce the MSE. Using this architecture, a closed curve is consistently obtained in the latent space. The evolution of the MSE during the training procedure is displayed in Fig. 3(c). We show fittings both with and without self-intersections of the trajectory in the latent space. A comparison of these fittings reveals a similar mean MSE in the last epoch for the train data set but not for the test data set, showing approximately 5 % and 40 % percentage difference, respectively. Encodings with self-intersections of the trajectory lead to overfitting.

Syringeal dynamics has been modeled in the literature as a relaxation oscillator [14]. The most outstanding characteristic of this class of oscillators is that its attracting limit cycle does not have a single time scale. As a result, an inhomogeneity in the density of points along the cycle is obtained in phase space.

It is not obvious that a heterogeneity in the distance between successive frames corresponds to one in the encoded points in the latent space. To carry out this validation, we compute the Frobenius norm between consecutive frames (d_g) and the Euclidean distance between consecutive points of the reconstructed flow in the latent space (d_{lat}, see the two panels of Fig. 4(a)). By calculating the correlation between these two signals we can test the hypothesis that heterogeneities in the latent space reflect dynamic aspects of the problem. In Fig. 4(b) we present the calculation of the correlation between the distance of consecutive frames of the movie, and the distances between consecutive points in latent space for each of the 94 trained networks which did not present self-intersections of the trajectory in the latent space.

Linear approaches to the problem of dimensional reduction of a spatiotemporal pattern are based on singular value decomposition (SVD), a technique capable of providing the linearly optimal way of truncating the dynamics into correlation eigenstates [19]. If we use this procedure to analyze our problem, we find two principal modes, and the density of points in this two-dimensional modal space also correlates with the series of distances between successive frames. However, the correlation achieved using SVD is lower than the highest correlation obtained using an autoencoder, as illustrated in Fig. 4(b). The selected autoencoder outperformed the optimal linear procedure, showing the power of using these networks in dynamics as a nonlinear generalization of SVD.

The model proposed in the literature postulates the existence of two nullclines: one linear, and the other cubic [20]. The dynamics is organized by the slow exploration of two of the branches of the cubic nullcline, and the rapid alternation between both branches. Agreement with such a model would imply that for each period of the oscillations there are two maxima and two minima in the distance between states along the trajectory. Since the analyzed movie is recorded at a constant frame rate, and there is a correlation between the frames' representation in latent space and the distances between consecutive frames, we expect to have two local maxima and two local minima in the distance between successive points in latent space. This expectation is positively confirmed in Fig. 4(a). In Fig. 4(c) we represent the distance of points in the latent space using a colour code, where dark areas correspond to regions of high density.

The dynamics of this problem was hypothesized over two decades ago, based on energetic considerations about the modes involved and aimed at explaining spectral features of the vocalizations. Reconstructing the phase dynamics in this problem is crucial to account for the behavioral output of this system: the sound. The spectral features of a

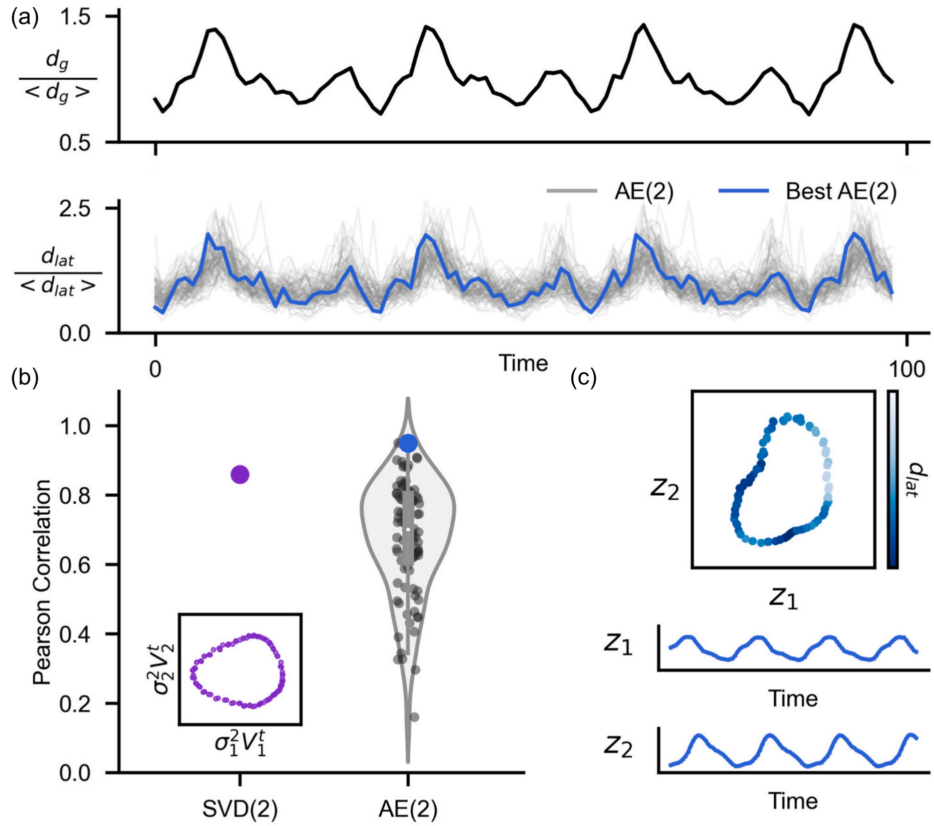

Fig. 4. Local similarity between the frames' dynamics and the reconstructed flow in the latent space. (a) The temporal evolution of the distances i n frame and latent space for the test set, divided by its mean values, computed as indicated in the text. Each fitting is truncated after 200 epochs. In blue, we show the simulation that led to a d_{lat} presenting the highest correlation with d_g. (b) The values of the correlation. AE(2) stands for autoencoder with a two-dimensional latent space. We show the projection into a space consisting of the first two modes of a singular value decomposition analysis (denoted by SVD(2)), and the correlation of its distance with d_g. (c) The projection of the data onto the latent space for the optimal autoencoder. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

relaxation oscillator alternating between two fast regions and two slow regions in phase space is completely different from an oscillator that, for instance, has a uniformly evolving phase. This result constitutes the first direct validation of this hypothesis. It is important to note that while the reconstruction of a normal form from the data could account for the existence of a limit cycle, it may not necessarily explain a phase heterogeneity. Autoencoders thus provide a topologically correct and dynamically faithful low-dimensional representation of an experimental spatiotemporal recording.

4. Conclusions

In this work we explored the topology of reconstructed flows from spatiotemporal data using autoencoders. We addressed two distinct problems. In the first one, the dynamics was known, since the movie was generated through the numerical integration of the dynamical system that governed the behavior of the modal amplitudes of the problem. In the second one, the movie was experimental in nature, and corresponds to the first direct observations of the labial oscillations in the vocal apparatus of a bird. In both cases, we find that the topology recovered in the latent space corresponds to the expected topology, which accounts for the success of the technique in terms of its ability to recover the dynamics in a low-dimensional representation space.

To compare the dynamics of the original problem in phase space and in latent space, in the first example we focused on the topological organization of segments of the movie that are good approximations to the lowest periodic orbits of the system. Specifically, we explored how approximations to period one and period two orbits are linked around each other. Recent works in the field of topology of dynamical systems indicate that their organization imposes restrictions on branched manifolds that support the attractor, so we focus on verifying that in the latent space, this organization is fully reflected. For the second problem, we not only verify that a periodic orbit was reproduced in latent space, but also that the distribution of densities of the points representing states reflected the dynamics of a relaxation oscillator.

Data driven methods are commonly used to carry out predictions, after a training process with dynamic records [21]. Often presented as black boxes, these techniques were not originally expected to provide insights into our understanding of the underlying dynamic mechanisms. However, we show that deep aspects of the dynamics are faithfully represented in the latent space of an autoencoder [5]. In particular, we show that autoencoders trained with spatiotemporal data enable the reconstruction of the topological structure of the underlying flows.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Codes and the results for the analysis are openly available in Zenodo at https://doi.org/10.5281/zenodo.8252884, Ref. [22]. Experimental recordings of the avian vocal organ are available from C. P. H. E. upon reasonable request.

Acknowledgments

We acknowledge partial financial support from the ANPCyT-FON-CyT (Argentina) under Grant PICT-2018-00619, and UBACyT (Argentina). G. B. M. acknowledges the hospitality of U. R. J. C. during his stays.

References

[1] Champion K, Lusch B, Kutz JN, Brunton SL. Data-driven discovery of coordinates and governing equations. Proc Natl Acad Sci U S A 2019;116:224455.
[2] Goodfellow I, Bengio Y, Courville A. Deep learning. MIT Press; 2016.
[3] Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. Science 2006;313:504.
[4] Gilmore R, Lefranc M. The topology of chaos. Wiley-VCH; 2003.
[5] Uribarri G, Mindlin GB. The structure of reconstructed flows in latent spaces. Chaos 2020;30:093109.
[6] Rössler OE. An equation for continuous chaos. Phys Lett A 1976;57(5):397–8.
[7] Uribarri G, Mindlin GB. Dynamical time series embeddings in recurrent neural networks. Chaos, Solitons Fractals 2022;154:111612.
[8] Lorenz EN. Deterministic nonperiodic flow. J Atmos Sci 1963;20(2):130–41.
[9] Birman JS, Williams RF. Knotted periodic orbits in dynamical systems I: Lorenz's equations. Topology 1983;22(1):47–82.
[10] Agostini L. Exploration and prediction of fluid dynamical systems using auto-encoder technology. Phys Fluids 2020;32:067103.
[11] Solari H, Natiello M, Mindlin GB. Nonlinear dynamics: a two-way trip from physics to math. CRC Press; 1996.
[12] Hirsch MW, Smale S, Devaney RL. Differential equations, dynamical systems, and an introduction to Chaos. Academic Press; 2013.
[13] Grassberger P, Procaccia I. Characterization of strange attractors. Phys Rev Lett 1983;50(5):346.
[14] Gardner T, Cecchi G, Magnasco M, Laje R, Mindlin GB. Simple motor gestures for birdsongs. Phys Rev Lett 2001;87:208101.
[15] Elemans CPH, et al. Universal mechanisms of sound production and control in birds and mammals. Nat Commun 2015;6:9978.
[16] Amador A, Perl YS, Mindlin GB, Margoliash D. Elemental gesture dynamics are encoded by song premotor cortical neurons. Nature 2013;495:59.
[17] Sitt JD, Amador A, Goller F, Mindlin GB. Dynamical origin of spectrally rich vocalizations in birdsong. Phys Rev E 2008;78(1):011905.
[18] Jiang W, Rasmussen JH, Xue Q, Ding M, Zheng X, Elemans CPH. High fidelity continuum modeling predicts avian voiced sound production. Proc Natl Acad Sci U S A 2020;117:4718.
[19] Holmes P, Lumley JL, Berkooz G, Rowley CW. Turbulence, coherent structures, dynamical systems and symmetry. Cambridge University Press; 2012.
[20] Laje R, Gardner TJ, Mindlin GB. Neuromuscular control of vocalizations in birdsong: a model. Phys Rev E 2002;65(5):051921.
[21] Brunton SL, Kutz JN. Data-driven science and engineering: machine learning, dynamical systems, and control. Cambridge University Press; 2022.
[22] Fainstein F, Catoni J, Elemans CPH, Mindlin GB. Data and code for 'the reconstruction of flows from spatiotemporal data by autoencoders'. Zenodo 2023. https://doi.org/10.5281/zenodo.8252884.