

Investigating echo state networks dynamics by means of recurrence analysis

Filippo Maria Bianchi^{*1}, Lorenzo Livi^{†‡2,3}, and Cesare Alippi^{§2,3}

¹Dept. of Physics and Technology, UiT the Arctic University of Norway, Hansine Hansens veg 18, 9019 Tromsø, Norway

²Dept. of Electronics, Information, and Bioengineering, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy

³ALaRI, Faculty of Informatics, Università della Svizzera Italiana, Via G. Buffi 13, 6904 Lugano, Switzerland

February 5, 2022

Abstract

In this paper, we elaborate over the well-known interpretability issue in echo state networks. The idea is to investigate the dynamics of reservoir neurons with time-series analysis techniques taken from research on complex systems. Notably, we analyze time-series of neuron activations with Recurrence Plots (RPs) and Recurrence Quantification Analysis (RQA), which permit to visualize and characterize high-dimensional dynamical systems. We show that this approach is useful in a number of ways. First, the two-dimensional representation offered by RPs provides a way for visualizing the high-dimensional dynamics of a reservoir. Our results suggest that, if the network is stable, reservoir and input denote similar line patterns in the respective RPs. Conversely, the more unstable the ESN, the more the RP of the reservoir presents instability patterns. As a second result, we show that the L_{\max} measure is highly correlated with the well-established maximal local Lyapunov exponent. This suggests that complexity measures based on RP diagonal lines distribution provide a valuable tool to quantify the degree of network stability. Finally, our analysis shows that all RQA measures fluctuate on the proximity of the so-called edge of stability, where an ESN typically achieves maximum computational capability. We verify that the determination of the edge of stability provided by such RQA measures is more accurate than two well-known criteria based on the Jacobian matrix of the reservoir. Therefore, we claim that RPs and RQA-based analyses can be used as valuable tools to design an effective network given a specific problem.

Keywords— Echo state network; Recurrence plot; Recurrence quantification analysis; Time-series; Dynamics.

1 Introduction

Since the very first recurrent neural network architectures, attempts have been made to describe and understand the internal dynamics of the system – see e.g. [1] and references therein. Nowadays, such efforts found renewed interest by those researchers trying to “open the black-box” [2–5]. This comes natural as recent advances in various fields, such as neurosciences and biophysical systems modeling, demand understanding of the inner mechanism that drives the inductive inference in order to produce novel scientific results [6].

Reservoir computing is a class of state-space models characterized by a fixed state transition structure, the *reservoir*, and a trainable, memory-less *readout* layer [7–9]. A reservoir must be sufficiently complex to

*filippo.m.bianchi@uit.no

†lorenzo.livi@usi.ch

‡Corresponding author

§cesare.alippi@polimi.it

capture all salient features of the inputs, behaving as a time dependent, non-linear kernel function, which maps the inputs into a higher dimensional space. The reservoir is typically generated with a pseudo-random procedure, not rarely driven by a set of rules-of-thumb and a trial-and-error approach. A popular reservoir computing architecture is the Echo State Network (ESN) [10], a recurrent neural network with a non-trainable, sparse recurrent reservoir and an adaptable (usually) linear readout mapping the reservoir to the output. ESN reservoir characterization and design attracted significant research efforts in the last decade [11–13]. This is mostly due to the puzzling behavior of the reservoir, which, although randomly initialized, has shown to be effective in modeling non-linear dynamical systems of various nature [14–19]. This also motivated researchers to better understand, and hence control, the dynamics of reservoirs given the fact they are driven by input signals. The many existing approaches to reservoir design can be roughly partitioned as either supervised or unsupervised [8]. In the former case, the output of the system is taken into account to analyze, and then accordingly design, the reservoir. In the latter case, the reservoir is designed by considering algebraic/topological properties of the weight matrix [20–22] or criteria based on statistics of the neuron activations [23, 24]. Recently, it was shown that topologies designed according to deterministic criteria [25, 26] produced state-of-the-art results in all major benchmarks, having a more contractive dynamics with respect to randomly generated reservoirs. Another interesting reservoir design was proposed in [27]. The authors showed that, for certain tasks, a large (randomly connected) reservoir could be replaced by a single non-linear neuron with delayed feedback. It is worth citing that also the nature of the neuron activation function (spiking vs analog) has been studied in this context [28], suggesting that networks operating with spiking neurons achieve inferior performance at the current state of research, possibly due to the adopted complex pre- and post-processing strategies. Input-dependent measures to characterize the reservoir dynamics have been proposed in [29], which considered the temporal profile of the Jacobian of the reservoir, rather than static quantities such as the spectral radius. Another input-dependent method for describing the dynamic profile of the reservoir is proposed in [30]. The authors demonstrated that both transfer entropy (the predictive information provided by a source about a destination that was not already contained in the destination history) and active information storage (the amount of past information that is relevant to predict the next state) calculated on the reservoir neurons are maximized right on the transition to an unstable regime.

The use of Poincaré recurrence of states provides fundamental information for the analysis of autonomous dynamical systems [31]. This follows from Poincaré’s theorem, which guarantees that the states of a dynamic system must recur during its evolution. In other terms, the system trajectory in phase space must return in an arbitrarily small neighbourhood of any of the previously visited states with probability one. Recurrences contain all relevant information regarding a system behavior in phase space and can be linked also with dynamical invariants (e.g., metric entropy) and features related to stability. However, especially for high-dimensional complex systems, the recurrence time, which is the time elapsed between recurring states, is difficult to calculate even when assuming full analytical knowledge of the system.

Recurrence Plots (RPs) [31–34], together with the computation of dynamical invariants and heuristic complexity measures called Recurrence Quantification Analysis (RQA), offer a simple yet effective tool to analyze such recurrences starting from a time-series derived from the system under analysis. An RP is a visual representation of recurrence time, which also provides information about the duration of the recurrence by inspecting line patterns [35]. RPs are constructed by considering a suitable distance in the phase (equivalently, state) space and a threshold used to determine the recurrence/similarity of states during the evolution of the system. Recently, RPs have been extended to study heterogeneous recurrences [36], i.e., qualitatively different types of recurrences in phase space. It is worth citing also the rapidly developing field of recurrence networks [37], whose main goal is to exploit complex network methods to analyze time-series related to dynamical systems.

In this paper, we address the interpretability issue of ESNs by analyzing the dynamics of the reservoir neuron activations with RPs and RQA complexity measures. To the best of our knowledge, recurrence analysis has never been used to investigate reservoirs or ESNs. RPs and RQA-based techniques allow the designer to visualize and characterize (high-dimensional) dynamical systems starting from a matrix encoding the recurrences of the system states over time. We show that RPs and RQA-based analyses allow to deduce important and consistent conclusions about the behavior of the network. Therefore, we suggest that they can be used as a valuable analysis tool to design a network for the problem at hand. The novelty content of what proposed can be summarized as:

1. We offer a method to visualize high-dimensional dynamics of input-driven ESNs through a two-dimensional representation encoding recurrences of states. Our results show that, if the network is stable, line patterns of RPs related to reservoir dynamics and those of the input signal are similar. This similarity is lost as soon as the network becomes unstable;
2. We comment that the degree of instability of the network can be quantified by means of a measure derived from the RP. This measure is named L_{\max} and is based on the diagonal lines of the RP only (hence it is computationally cost effective). To support this claim, we show that such a measure is correlated with the well-established maximal local Lyapunov exponent over the entire range of the control parameters;
3. We show that RQA measures can help the application designer to tune important network parameters (e.g., spectral radius and input scaling), in order to push the ESN in a critical state, where computational capability (defined in terms of prediction performance and short-term memory capacity) is maximized. Such a critical state (called edge of stability) depends on the dynamics of the network which, in turn, is influenced by the input signal. Therefore, this cannot be achieved by considering only static measures, e.g., taking into account only the reservoir topology. As we will show here, the proposed method based on RQA measures is sensitive also to the nature of the input signal driving the network.

This paper is structured as follows. In Sec. 2 we introduce a typical ESN model and discuss some important figures of merit that have been proposed in the literature to characterize a reservoir. In Sec. 3, we present the contribution of our work, i.e., we link recurrence analysis with ESNs and discuss how this new tool can be used in practice to design a network exposing both stability and forecasting accuracy properties. In Sec. 4, we introduce experiments to support our claims. Finally, Sec. 5 offers concluding remarks and future directions.

2 Dynamics of the ESN reservoir

A schematic representation of an ESN is shown in Fig. 1. An ESN consists of a reservoir of N_r nodes characterized by a non-linear transfer function $f(\cdot)$. At time instant k , the network is driven by the input $\mathbf{x}[k] \in \mathbb{R}^{N_i}$ and produces the output $\mathbf{y}[k] \in \mathbb{R}^{N_o}$, being N_i and N_o the dimensionalities of input and output, respectively. The weight matrices $\mathbf{W}_r^r \in \mathbb{R}^{N_r \times N_r}$ (reservoir internal connections), $\mathbf{W}_i^r \in \mathbb{R}^{N_i \times N_r}$ (input-to-reservoir connections), and $\mathbf{W}_o^r \in \mathbb{R}^{N_o \times N_r}$ (output-to-reservoir feedback connections) contain values in the $[-1, 1]$ interval drawn from a uniform distribution; a Gaussian distribution is another common choice.




Figure 1: The ESN architecture. Circles represent input \mathbf{x} , state, \mathbf{h} , and output, \mathbf{y} , respectively. Solid squares \mathbf{W}_r^o and \mathbf{W}_i^o , are the trainable matrices of the readout, while dashed squares, \mathbf{W}_r^r , \mathbf{W}_o^r , and \mathbf{W}_i^r , are randomly initialized matrices. The polygon represents the non-linear transformation performed by neurons and z^{-1} is the unit delay.

An ESN is a discrete-time non-linear system with feedback, whose model reads:

$$\mathbf{h}[k] = f(\mathbf{W}_r^r \mathbf{h}[k-1] + \mathbf{W}_i^r \mathbf{x}[k] + \mathbf{W}_o^r \mathbf{y}[k-1]); \quad (1)$$

$$\mathbf{y}[k] = g(\mathbf{W}_r^o \mathbf{h}[k] + \mathbf{W}_i^o \mathbf{x}[k]). \quad (2)$$

Activation functions $f(\cdot)$ and $g(\cdot)$, both applied component-wise, are typically implemented as a sigmoidal (*tanh*) and identity function, respectively. The output weight matrices $\mathbf{W}_r^o \in \mathbb{R}^{N_r \times N_o}$ and

$\mathbf{W}_i^o \in \mathbb{R}^{N_i \times N_o}$, which connect, respectively, reservoir and input to the output, represent the readout of the network. The standard training procedure for such matrices requires solving a straightforward regularized least-square problem [38].

Even though the three weight matrices \mathbf{W}_r^r , \mathbf{W}_o^r , and \mathbf{W}_i^r are generated randomly, they can be suitably designed and scaled to obtain desired properties. For instance, \mathbf{W}_o^r can scale with multiplicative constant ω_o ; in this work, we set $\omega_o = 0$, which has the effect of removing the output feedback connection. \mathbf{W}_i^r is controlled by scalar parameter ω_i . Since the gain of the sigmoid non-linearity in the neurons is largest around the origin, the scaling coefficient ω_i of \mathbf{W}_i^r determines the amount of non-linearity introduced by the processing units. In particular, inputs far from zero tend to drive the activation of the neurons towards saturation where they show more non-linearity. Finally, the spectral radius of \mathbf{W}_r^r , denoted as $\rho(\mathbf{W}_r^r)$ or ρ , controls important properties as discussed in the sequel.

An ESN is typically designed so that the influence of past inputs on the network gradually fades away and the initial state of the reservoir is eventually washed out. This is granted by the Echo State Property (ESP), which ensures that, given any input sequence taken from a compact set, future trajectories of any two different initial states become indistinguishable. ESP was originally investigated in [38] and successively in [39]; we refer the interested reader to [40] for a more recent definition, where also the influence of input is explicitly accounted for. In ESNs with no output feedback, as in our case, the state update of equation (1) reduces to

$$\mathbf{h}[k] = f(\mathbf{W}_r^r \mathbf{h}[k-1] + \mathbf{W}_i^r \mathbf{x}[k]). \quad (3)$$

In order to study the stability of the network, we analyze the Jacobian of the state update (3) of the reservoir and generate a derived measure, the maximal local Lyapunov exponent (λ). Such a quantity is used to approximate (for an autonomous system) the separation rate in phase space of trajectories having very similar initial conditions. λ was proposed to characterize a reservoir, and demonstrated its efficacy in designing a suitable network configuration in several applications [7, 29, 41]. λ is calculated by considering the Jacobian at time k , which can be conveniently expressed if neurons are implemented with a *tanh* activation function as

$$\mathbf{J}(h[k]) = \begin{bmatrix} 1 - (h_1[k])^2 & 0 & \dots & 0 \\ 0 & 1 - (h_2[k])^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 - (h_{N_r}[k])^2 \end{bmatrix} \mathbf{W}_r^r, \quad (4)$$

where $h_l[k]$ is the activation of the l -th neuron, with $l = 1, 2, \dots, N_r$. λ is then computed as

$$\lambda = \max_{n=1, \dots, N_r} \frac{1}{K} \sum_{k=1}^K \log(r_n[k]), \quad (5)$$

where $r_n[k]$ is the module of n -th eigenvalue of $\mathbf{J}(h[k])$ and K is the total number of time-steps in the considered trajectory. Typically, the stable-unstable transition is detected numerically by considering the sign of λ (5). In autonomous systems, $\lambda > 0$ indicates that the dynamics is chaotic.

Local, first-order approximations provided by Eq. 4 are useful also for studying the stability of a (simplified) reservoir operating around the zero state, $\mathbf{0}$. In fact, implementing $f(\cdot)$ as a *tanh* assures $f(\mathbf{0}) = \mathbf{0}$, i.e., $\mathbf{0}$ is a fixed point of the ESN dynamics. Therefore, by linearizing (3) around $\mathbf{0}$ and assuming a zero-input, we obtain from (4)

$$\mathbf{h}[k] = \mathbf{J}(\mathbf{0})\mathbf{h}[k-1] = \mathbf{W}_r^r \mathbf{h}[k-1]. \quad (6)$$

Linear stability analysis of (6) suggests that, if $\rho(\mathbf{W}_r^r) < 1$, then the dynamic around $\mathbf{0}$ is stable. In the more general case, the non-linearity of the sigmoid functions in (3) forces the norm of the state vector of the reservoir to remain bounded. Therefore, the condition $\rho(\mathbf{W}_r^r) < 1$ loses its significance and does not guarantee stability when the system deviates from a small region around $\mathbf{0}$ [41]. This means that, actually, it is possible to find reservoirs (3) having $\rho(\mathbf{W}_r^r) > 1$ while still possessing the ESP with stability. In fact, the effective local gain decreases when the operating point of the neurons shifts toward

the positive/negative branch of the sigmoid, where stabilizing saturation effects start to influence the excitability of reservoir dynamics [39]. In the more realistic and useful scenario where the input driving the network is a generic (non-zero) signal, a sufficient condition for the ESP is met if \mathbf{W}_r^r is diagonally Schur-stable, i.e., if there exists a positive definite diagonal matrix, \mathbf{P} , such that $(\mathbf{W}_r^r)^T \mathbf{P} \mathbf{W}_r^r - \mathbf{P}$ is negative definite [39]. However, this recipe is fairly restrictive in practice. In many cases, such a sufficient condition might generate reservoirs that are not rich enough in terms of provided dynamics, since the use of a conservative scaling factor might compromise the amount of memory in the network and thus the ability to accurately model a given problem. Therefore, for most practical purposes, the necessary condition $\rho(\mathbf{W}_r^r) < 1$ is considered “sufficient in practice”, because if the spectral radius is less than 1, the state update map is contractive with high probability, regardless of the input and given a sufficiently large reservoir [22].

The number of reservoir neurons and the bounds on ρ can be used for a naïve quantification of the computational capability of a reservoir [39]. However, those are static measures that only consider the algebraic properties of \mathbf{W}_r^r , without taking into account other factors, such as the input scaling ω_i and the particular properties of the given input signals. Moreover, it is still not clear how, in a mathematical sense, these stability bounds relate to the actual ESN dynamics when processing non-trivial input signals [40]. In this context, the idea of pushing the system toward the so-called “edge of stability” (also called edge of criticality) has been explored. [42–44] show that several dynamical systems, among which randomly connected recurrent neural networks, achieved the highest computational capabilities when moving toward the unstable (sometime even chaotic) regime, where the ESP is lost and the system enters into an oscillatory behavior. This justifies the use of spectral radii above the unity in some practical applications. It is important to remark that, although the critical state might maximize the computational capability, there are also tasks that either require very little in terms of computation or reservoirs with fast reaction times and fast locking into attractor states (e.g., multistable switching circuits [45]).

Further descriptors used for characterizing the dynamics of a reservoir are based on information-theoretic quantities, such as the (average) transfer entropy and the active information storage one [30]. The authors have shown that such quantities peak right when $\lambda > 0$. In addition, also the minimal singular value of the Jacobian (4), denoted as η , was demonstrated to be an accurate predictor of ESN performance, providing more accurate information regarding the ESN dynamics than both λ and ρ [29]. Finding hyper-parameters that maximize η results in a dynamical system that is far from singularity, it has many degrees of freedom, a good excitability, and it separates well the input signals in phase space [29].

Complementary to λ , η , and the aforementioned information-theoretic descriptors, in this study we propose to use RPs and the related RQA measures for characterizing the dynamics of an ESN when driven by an input signal.

3 Analyzing ESN dynamics by investigating recurrences of neuron activations

In this section, we discuss how the input-driven dynamics of an ESN reservoir (3) can be fruitfully analyzed by means of RP-based techniques.

The sequence of ESN states can be seen as a multi-variate time-series \mathbf{h} , consisting of N_r state variables coming from the reservoir neuron activations. An RP is constructed by calculating a $K \times K$ binary matrix \mathbf{R} . The generic element R_{ij} is defined as

$$R_{ij} = \Theta(\tau_{\text{RP}} - d(\mathbf{h}[i], \mathbf{h}[j])), \quad 1 \leq i, j \leq K, \quad (7)$$

where K is the length of the time-series, $d(\cdot, \cdot)$ is a dissimilarity measure operating in phase space, and $\Theta(\cdot)$ the Heaviside function: $\Theta(x) = 0$ if $x < 0$; $\Theta(x) = 1$ otherwise. $\tau_{\text{RP}} > 0$ is a user-defined threshold used to identify recurrences. τ_{RP} can be defined in different ways, but typically it represents a percentage of the average or the maximum phase space distance between the states. The selection of an “optimal” threshold τ_{RP} is a problem-dependent issue. However, we comment that, usually, τ_{RP} has an impact only on the quantitative information derived from an RP but does not affect the general properties of the system as seen through a recurrence analysis. \mathbf{R} is constructed by considering also a nonnegative

dissimilarity measure $d(\cdot, \cdot)$, which evaluates the distance between states. Typical examples include the Euclidean, Manhattan, or max-norm distances. Several options have been exploited in the literature for τ_{RP} and $d(\cdot, \cdot)$; we refer the reader to [31] for further detailed discussions. Fig. 2 depicts the algorithmic steps required to generate an RP on ESN states.

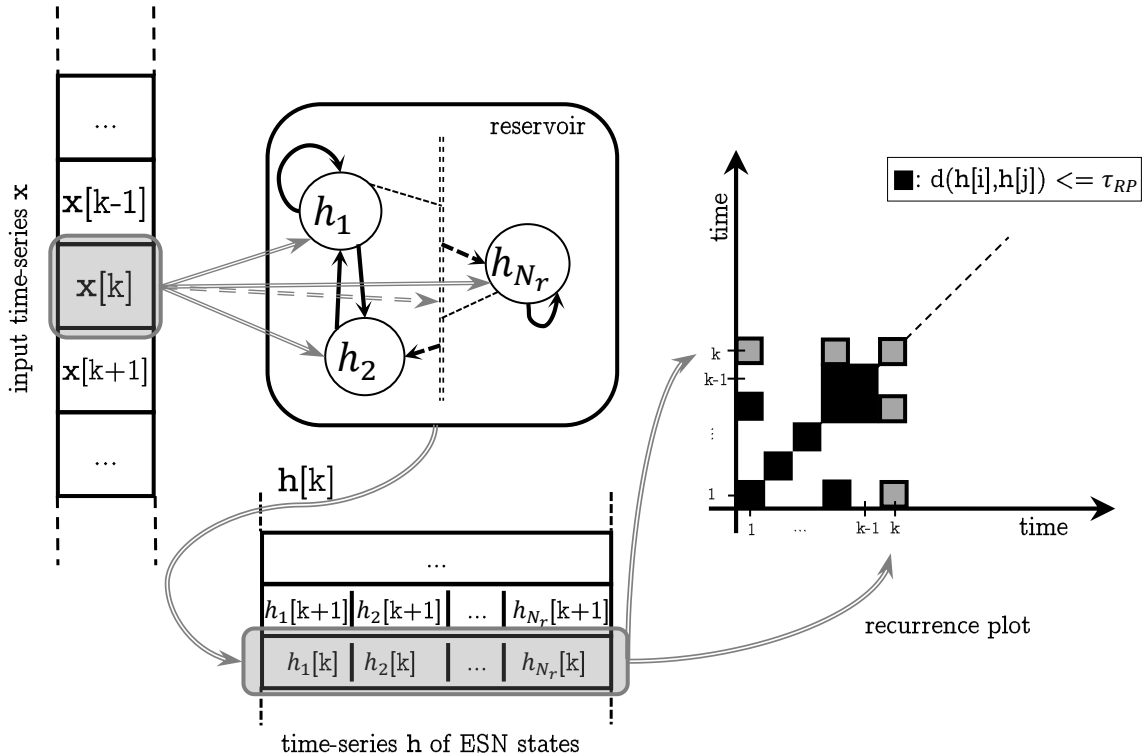


Figure 2: When $\mathbf{x}[k]$ is fed as input to the N_r neurons of the ESN reservoir, the internal state is updated to $\mathbf{h}[k] = [h_1[k], h_2[k], \dots, h_{N_r}[k]]^T$, where $h_n[k]$ is the output of the n -th neuron. Once the time-series \mathbf{h} is generated, the RP is constructed by using a threshold τ_{RP} and a dissimilarity measure $d(\cdot, \cdot)$. If $d(\mathbf{h}[k], \mathbf{h}[i]) \leq \tau_{RP}$, the cell of the RP in position (k, i) is colored in black, otherwise it is left white. The elements in gray highlight the operations performed at time-step k .

Depending on the properties of the analyzed time-series, different line patterns emerge in an RP [35]. Besides providing an immediate visualization of the system properties, starting from \mathbf{R} , it is possible to compute also several complexity measures, those associated with an RQA. Such measures are defined by the distribution of both vertical/horizontal and diagonal line structures present in the RP and provide a numerical characterization of the properties associated with the underlying dynamics.

In Sec. 3.1, we provide details about the RQA measures considered in this study. In Sec. 3.2, we discuss how a network designer should take into account the two important issues of stability and computational capability by using the tools here introduced. Finally, in Sec. 3.3 we explain how to interpret the RP and related RQA measures relative to ESN state sequences driven by different classes of inputs.

3.1 RQA complexity measures

We assume here to have generated \mathbf{R} according to (7) on a time-series of length K . Many of the following measures are based on the histograms $P(l)$ and $P(v)$, counting, respectively, the number of diagonal and

vertical lines having specific lengths l and v ,

$$P(l) = \sum_{i,j=1}^{K-l} (1 - R_{i-1,j-1})(1 - R_{i+l,j+l}) \prod_{k=0}^{l-1} R_{i+k,j+k};$$

$$P(v) = \sum_{i,j=1}^{K-v} (1 - R_{i,j})(1 - R_{i,j+v}) \prod_{k=0}^{v-1} R_{i,j+k}.$$

Such RQA measures cover different aspects of the dynamics (e.g., features related to the stability and complexity of the system evolution). Abbreviations and notation are kept consistent with [31].

RR. The recurrence rate is a measure of the density of recurrences in \mathbf{R} ,

$$\text{RR} = \frac{1}{K^2} \sum_{i,j=1}^K R_{ij}. \quad (8)$$

Given Eq. 7, RR corresponds to the correlation sum, an important concept used in chaos theory. RR could help also with the selection of τ_{RP} when performing multiple tests on different conditions, e.g., by preserving the rate.

DET. A measure of the determinism level of the system, based on the percentage of diagonal lines of minimum length l_{\min} ,

$$\text{DET} = \frac{\sum_{l=l_{\min}}^K lP(l)}{\sum_{l=1}^K lP(l)} \in [0, 1]. \quad (9)$$

A periodic system would have DET close to unity and close to zero for a signal with no time-dependency. The threshold l_{\min} can be used also to discriminate different forms of determinism, such as periodic and chaotic motions. In fact, it is known that in chaotic dynamics diagonal lines are very short – due to exponential divergence in phase space.

L_{max}. Maximum diagonal line length,

$$L_{\max} = \max\{l_i\}_{i=1}^{N_l}, \quad N_l = \sum_{l \geq l_{\min}} P(l), \quad (10)$$

where l_i is i th diagonal line length, N_l is the total number of diagonal lines, and $1 \leq L_{\max} \leq \sqrt{2}K$. This measure is related to the mean exponential divergence in phase space. A measure of divergence can be obtained as

$$\text{DIV} = 1/L_{\max} \in (0, 1]. \quad (11)$$

Chaotic systems do not present long diagonal lines, since trajectories diverge exponentially fast. As formally discussed in [31], it is possible to relate L_{\max} with the correlation entropy of the system – that is a dynamical invariant measure. Notably, correlation entropy can be estimated by using the diagonal lines distribution. The correlation entropy is a lower-bound for the sum of the positive Lyapunov exponents, providing thus a formal connection with the analysis of sensitivity on initial conditions.

LAM. Laminarity is a descriptor of the presence of laminar phases in the system,

$$\text{LAM} = \frac{\sum_{v=v_{\min}}^K vP(v)}{\sum_{v=1}^K vP(v)} \in [0, 1], \quad (12)$$

where v_{\min} is a threshold for the minimal vertical line length to be considered. Laminar phases denote states of the system that do not change or change very slowly for a number of consecutive time-steps.

ENTR. A complexity measure of an RP with respect to the diagonal lines distribution,

$$\text{ENTR} = - \sum_{l=l_{\min}}^K p(l) \ln(p(l)), \quad (13)$$

where $p(l) = P(l)/N_l$. As usual, $0 \leq \text{ENTR} \leq \log(K)$. Signals with no time-dependence present $\text{ENTR} \simeq 0$, i.e., the diagonal lines distribution is fully concentrated on very short lines (e.g., single dots). Conversely, such a measure assumes high values when the diagonal lines distribution become heterogeneous.

SWRP. Entropy of a weighted RP. Such a measure has been recently introduced in [33] and provides an alternative to ENTR (13). It quantifies the complexity of scalar distributions of *strengths* for each time-step, providing a solution for the border effects that might appear in the computation of ENTR. The computation of SWRP considers the similarity matrix $\tilde{\mathbf{S}}$, whose elements are defined as $\tilde{S}_{ij} = \exp(-d(\mathbf{h}[i], \mathbf{h}[j]))$, where $d(\cdot, \cdot)$ is the same as in Eq. (7). The results do not depend on τ_{RP} and are not influenced by the length K of the time-series. The strengths are used to measure the heterogeneity of the density of a given point in phase space. The i -th strength is defined as: $s_i = \sum_{j=1}^K \tilde{S}_{ij}$. The distribution of the strengths is represented with histograms using a user-defined number B of bins. SWRP measure is then defined as

$$\text{SWRP} = - \sum_{b=1}^B p(b) \ln p(b), \quad (14)$$

where $P(b)$ is the number of elements in the b -th bin and $p(b) = P(b)/K$ is the respective probability; as usual, $0 \leq \text{SWRP} \leq \log(B)$.

3.2 Design of a stable and effective network

Recurrence analysis can be used to (i) guarantee stability for a given configuration and (ii) tune the network, so that the highest computational capability can be achieved. These aspects are treated in the sequel.

Stability issue

Constructing an RP from \mathbf{h} allows to investigate the network stability issue. As mentioned before, the degree of stability can be calculated with L_{max} (10): the higher L_{max} , the more stable the system. In addition to L_{max} , the designer can have an immediate visual interpretation of stability by inspecting RP: short and erratic diagonal lines in RPs denote instability/chaoticity, while long diagonal lines denote regularity (e.g., a periodic motion). We show in the experimental section that L_{max} is anticorrelated with λ and hence it can be considered as a reliable indicator for the (input-dependent) degree of network stability. In the experiments we highlight also that, if the network is stable, the dynamics of the input and reservoir produce very similar line patterns in the respective RPs. If the network is unstable, then the aforementioned similarity is lost. Therefore, RPs can be used by the designer as visual tools to analyze the response of the network to a specific input.

Computational capability issue

A widespread criterion to determine the (signal-dependent) edge of stability relies on the sign of λ , which becomes positive when the system enters into a (globally) unstable regime. We propose, instead, to use information derived from RR (8), DET (9), LAM (12), ENTR (13), and SWRP (14) for this purpose. We observed that, when such indexes start to fluctuate, the network achieves high prediction accuracy. This provides a guideline to the network designer for setting critical parameters, such as ρ and ω_i , in presence of input. RQA measures, when evaluated on ESNs initialized with different weight matrices, assume very different values given ρ and ω_i when the network is unstable. Accordingly, it emerges that the standard deviation for such RQA measures is very high in such unstable regimes. We propose to identify the configurations of ρ and ω_i which bring the system to the edge of stability with those for which we observe a sudden increment in the standard deviation. Let P and Ω be the domains for ρ and ω_i , respectively. In this study, we calculate such configurations by checking when the standard deviations of RQA measures assume values greater than their mean values computed over $P \times \Omega$. For a given RQA measure, say q , we identify the edge of stability as a set of pairs $\{(\rho^1, \omega_i^1), (\rho^2, \omega_i^2), \dots, (\rho^{|\Omega|}, \omega_i^{|\Omega|})\} \subset P \times \Omega$. For every $\omega_i^j \in \Omega$, we select the largest $\rho^j \in P$ such that

$$\sigma_q(\rho^j, \omega_i^j) \leq \bar{\sigma}_q, \quad (15)$$

is satisfied. $\sigma_q(\rho^j, \omega_i^j)$ is the standard deviation of q for the configuration (ρ^j, ω_i^j) and $\bar{\sigma}_q$ is the average value of the standard deviation computed over $P \times \Omega$. Fig. 3 shows an illustrative example. We stress that other criteria based on the fluctuations of RQA measures could be conceived. However, as we will show in the experiments, such a criterion offers a more accurate description for the edge of stability than λ .




Figure 3: Standard deviations of a RQA measure (RR) for different values of ρ and ω_i , relative to two different input signals. The edge of stability is identified in correspondence of the parameters configurations for which the standard deviation increases abruptly (shown as a red line), assuming values greater than its mean value.

3.3 Visualization and classification of reservoir dynamics

In the following, we show how RPs permit to visualize, and hence classify, the reservoir dynamics when fed with inputs possessing important characteristics. In doing so, we assume given a stable ESN described by (3); RPs are constructed following the procedure depicted in Fig. 2. Although many classes of important signals/systems exist (with related sub-classes) [31], here we focus on the ability to discriminate between important classes for the input signals: (i) with/without time-dependence; (ii) periodic/non-periodic motions; (iii) laminar behaviours; (iv) chaotic dynamics; finally, (v) non-stationary processes.

We refer to examples in order to observe the different behaviours shown by an RP in correspondence with the aforementioned classes. We recall that RPs offer visual information and the validity of the following comments are general and not application-specific. In addition, we stress that RP-based analyses are applicable also to short time-series (with tens/hundreds of records).

Time-dependency

A uniformly distributed RP is a clear sign of the absence of a time-dependency in the time-series (e.g., an uncorrelated signal). It is possible to rely on specific RQA measures in order to numerically investigate the presence of time-dependency. A signature for this is observed by checking the outcomes of DET (9), ENTR (13), and SWRP (14). All three indexes would yield very low values (close to zero) when there is no time-dependence in the signal. A periodic signal offers a counter-example having instead a strong time-dependency; DET would be very high for a periodic signal, yet ENTR and SWRP would still be low. In fact, ENTR and SWRP are conceived to highlight the complexity aspects: signals with low complexity include those with no temporal structure. As an example, in Fig. 4(a) we examine the RP generated by feeding the ESN with white Gaussian noise, a typical example of signal with no

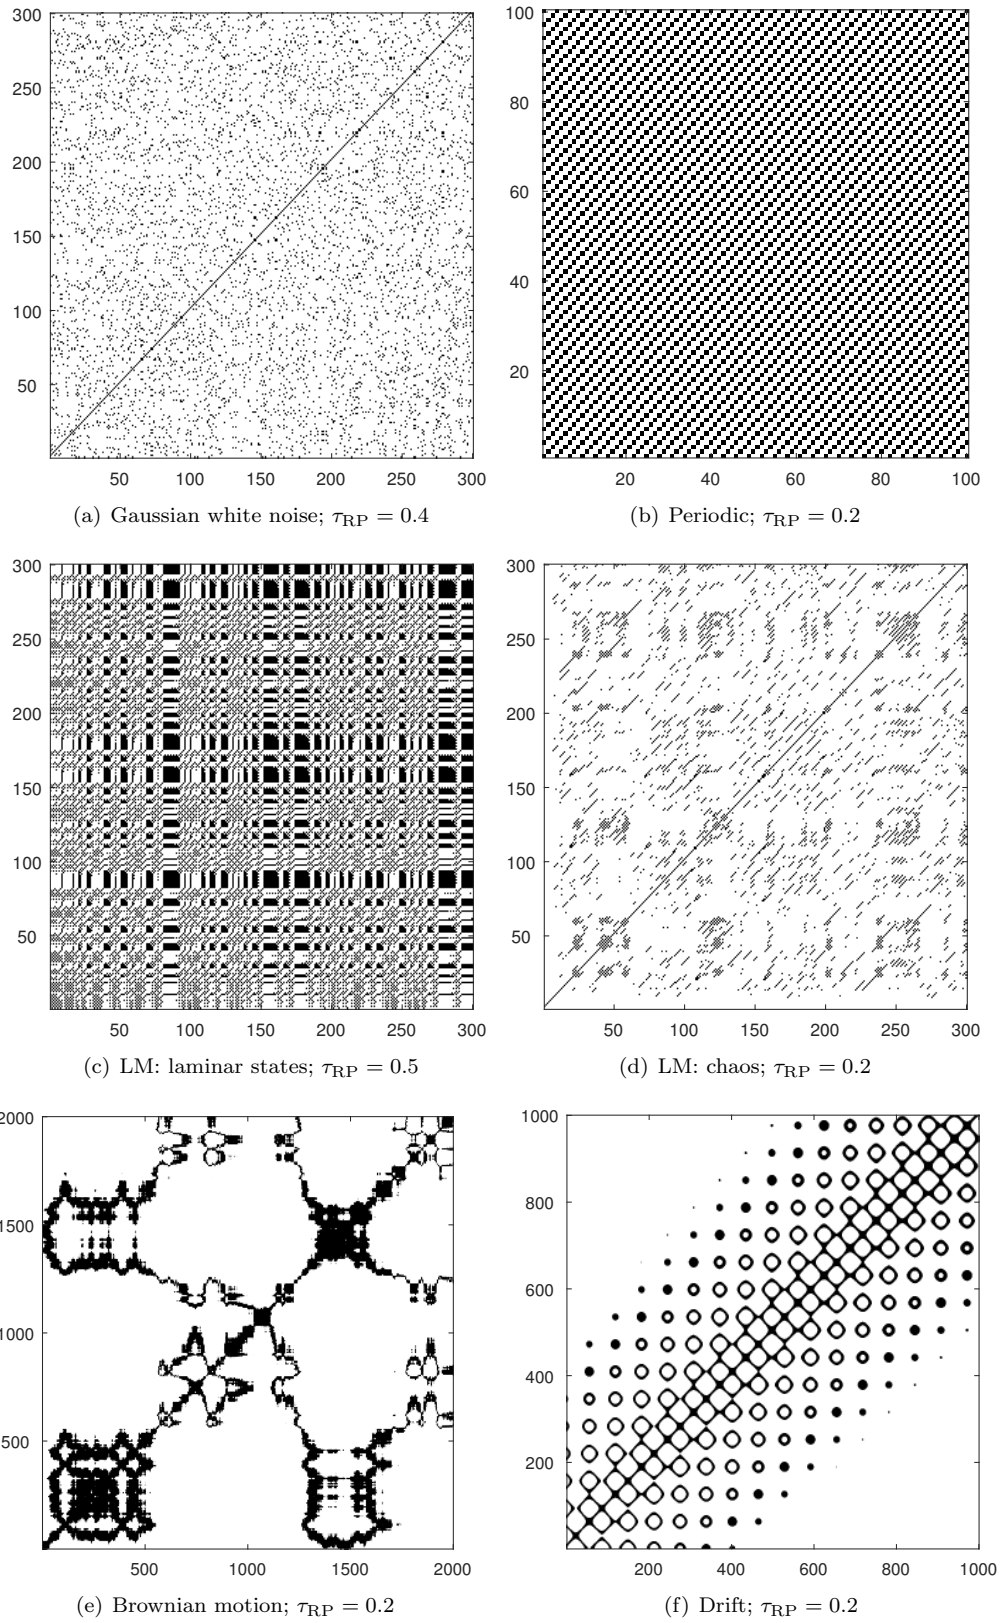


Figure 4: RPs generated by state sequences \mathbf{h} of ESN fed with input signals associated with the considered classes. Both axes represent time.

time-dependency. We observe a uniform RP for the reservoir states, which is peculiar for all signals composed by realizations of statistically independent variables.

Periodicity

Every periodic system would induce long diagonal lines and the vertical spacing provides the characteristic period of the oscillation. A periodic system is typically accompanied by high values for DET (9) and L_{\max} (10). In addition, as stressed before, it has low complexity as expressed by ENTR and SWRP. In Fig. 4(b), we show an example of periodic motion generated with a sinusoid having a single dominating frequency. The regularity of the diagonal lines can be immediately recognized from the figure.

Laminarity

A system presents laminar phases if its state does not change or change very slowly over a number of successive time-steps. Laminar phases can be visually recognized in an RP by the presence of (fairly) large black rectangles. Every system possessing laminar phases is characterized by high values for LAM (12). In order to provide an example (but results are independently valid), we exploit the logistic map (LM),

$$\mathbf{x}[n + 1] = \tau_{\text{LM}}\mathbf{x}[n](1 - \mathbf{x}[n]), \quad (16)$$

where usually $\tau_{\text{LM}} \in (0, 4]$; here LM (16) is used with initial condition $\mathbf{x}[0] = 0.5$. In Fig. 4(c) we show the RP when $\tau_{\text{LM}} = 3.679$. The system exhibits chaos-chaos transitions in that configuration. In fact, the related RP is compatible with the one of a (mildly) chaotic system, showing also the presence of laminar phases visualized as large black rectangles.

Chaoticity

RPs offer a particularly useful visual tool in the case of chaotic dynamics. In fact, such dynamics can be recognized by the presence of erratic and very short diagonal lines. As a consequence, RR (8) would be very low. ENTR and SWRP are two measures of complexity that are useful in characterizing also the degree of chaoticity: the higher the measures, the more chaotic/complex the system. Chaos is characterized by trajectories diverging exponentially fast. This can be quantified with L_{\max} (10) and DIV (11), whose values would be respectively very low and close to one for systems with a high degree of chaoticity.

As an example, we consider a chaotic system obtained through (16) with $\tau_{\text{LM}} = 4$. The reservoir dynamics, as shown in the RP in Fig. 4(d), denotes fully developed chaos, as indicated by the presence of short and erratic diagonal lines only.

Non-stationarity

Peculiar line patterns observed for all nonstationary signals include large white areas with irregular patterns denoting abrupt changes in the dynamics. Drift is another typical form of nonstationarity, which is visually recognized in an RP by the fading of recurrences in the upper-left and lower-right corners. In Fig. 4(e), we show an example by feeding the ESN with a well-known nonstationary signal: Brownian motion, a random walk resulting in a nonstationary stochastic process; whose increments correspond to Gaussian white noise, a stationary process. In Fig. 4(f), we show an example of drift by adding a linear trend to a sinusoid. Nonstationarity can be numerically detected by considering an RQA measure called TREND (not used in our study) and by analyzing the variation of RQA measures when time-delay is applied to the signal (see [31] for technical details).

4 Experiments

In this section, we consider two time-series generated respectively by an oscillatory and by the Mackey-Glass (MG) dynamical system. We chose these two signals since both of them are often considered as benchmarks for prediction in the ESN literature [7, 38, 46] and they exemplify a very regular and a mildly chaotic system, respectively. We perform two experiments. In the first one – Sec. 4.1 – we use RPs to visualize the dynamics of reservoirs when driven by a given input signal. When the reservoir operates in a stable regime, RPs of reservoir and input show similar line patterns.

In the second experiment – Sec. 4.2 – we use RPs as an analysis tool to tune two critical hyperparameters, the spectral radius (ρ) and the input scaling (ω_i), in an ESN driven by input signal. We

study the derived RQA measures for different settings of ρ and ω_i , with the aim to (i) assess the network stability and (ii) determine a suitable setting for these two parameters, where the predictive power of the network is maximized. Note that we focus on the prediction of the time-series generated by MG system [29], rather than on the test for learning its chaotic attractor from data [38]. The latter test requires to enable the output feedback, which leads to an internal dynamic and to a Jacobian matrix, different from the ones considered in this work.

In both experiments, we consider an ESN with no output feedback ($\omega_o = 0$), configured with a standard setting: uniformly distributed weights in $[-1, 1]$ for W_i^r and W_r^r , percentage of non-zero connections in W_r^r of 25%. The readout is trained by setting the regularization parameter in the linear regression to 0.1. According to the standard drop-out procedure, we discarded the first 100 elements of \mathbf{h} in order to get rid of the ESN transient states. In the first experiment, we set the number of reservoir neurons to $N_r = 50$. In preliminary tests, we noted that the number of reservoir neurons, even if very small (two neurons), does not affect the results obtained for the analysis of the recurrences. Therefore, we do not provide here a study in these terms. We used the Manhattan distance for evaluating the dissimilarity in the phase space,

$$d(\mathbf{h}[j], \mathbf{h}[i]) = \sum_{n=1}^{N_r} |h_n[j] - h_n[i]|. \quad (17)$$

The threshold τ_{RP} has been calculated by using a percentage of the average dissimilarity value between the states in \mathbf{h} . For simplifying the notation, we directly indicate the percentage referring to τ_{RP} , e.g., $\tau_{\text{RP}} = 0.1$ indicates 10% of the average distance between all states. In few cases, to better visualize the RPs, we increase τ_{RP} beyond values typically used in the literature. However, as stated before, our main aim here is to show the agreement between RQA measures and the stability–instability transition together with the corresponding accuracy of the network, rather than providing an *exact* quantification of its characteristics.

Our results are easily reproducible by using the ESN¹ and RP² toolboxes available online.

4.1 Visualization of ESN dynamics

The first signal is a sinusoid defined as $\mathbf{x}[k] = \sin(\psi k)$, where $k = 1, 2, \dots, 5000$ and $\psi = 3/50$.

In Fig. 5 we show the RPs relative to the input signal and the reservoir states; input signal has been embedded into a 2-dimensional phase space. In Fig. 5(b) we used a conservative setting for ρ . The resulting RP is compatible with a periodic system having a single frequency. This suggests that, although the (randomly initialized) reservoir performs a non-linear mapping, its resulting dynamics preserves the one of the input system. In fact, the vertical spacing between the diagonal lines is almost the same, which is 98 in Fig. 5(a) and 101 in Fig. 5(b). This small difference is due to a slight discrepancy in the period of the two systems (i.e., input and reservoir), whose cause was formally explained in [40]. There the authors demonstrate that when an ESN is driven by a periodic signal with period ϑ , the induced network dynamics asymptotically becomes periodic as well, with period $r\vartheta$, $r > 0$. On the other hand, when ρ is pushed beyond one, the RPs are comparable with the ones of unstable/chaotic systems with a degree of instability monotonically related to the value of ρ . Please note that, while in Figs. 5(a) and 5(b) we used $\tau_{\text{RP}} = 0.2$ and $\tau_{\text{RP}} = 0.1$, respectively, in Figs. 5(c) and 5(d) we used a large threshold, $\tau_{\text{RP}} = 0.8$, to improve readability of the plots.

The second input signal is given by a time-series generated from the Mackey-Glass (MG) system,

$$\frac{dx}{dt} = \frac{\alpha x(t - \tau_{\text{MG}})}{1 + x(t - \tau_{\text{MG}})^{10}} - \beta x(t). \quad (18)$$

We obtained a time-series of 150000 time-steps using $\tau_{\text{MG}} = 17$, $\alpha = 0.2$, $\beta = 0.1$, initial condition $x(0) = 1.2$, and 0.1 as integration step for (18). In Fig. 6, we report RPs of the input signal and those related to the three different settings of the ESN with an increasing value of ρ . As in the previous case, it is worth stressing the similarity between RPs on input in Fig. 6(a) and reservoir in Fig. 6(b). As the reservoir is pushed toward instability (by increasing ρ), it is possible to observe the usual incremental transition toward a chaotic regime in the related RPs. In Figs. 6(c) and 6(d) ρ is set to 1.5 and 2,

¹<http://www.reservoir-computing.org/node/129>

²<http://www.recurrence-plot.tk/>

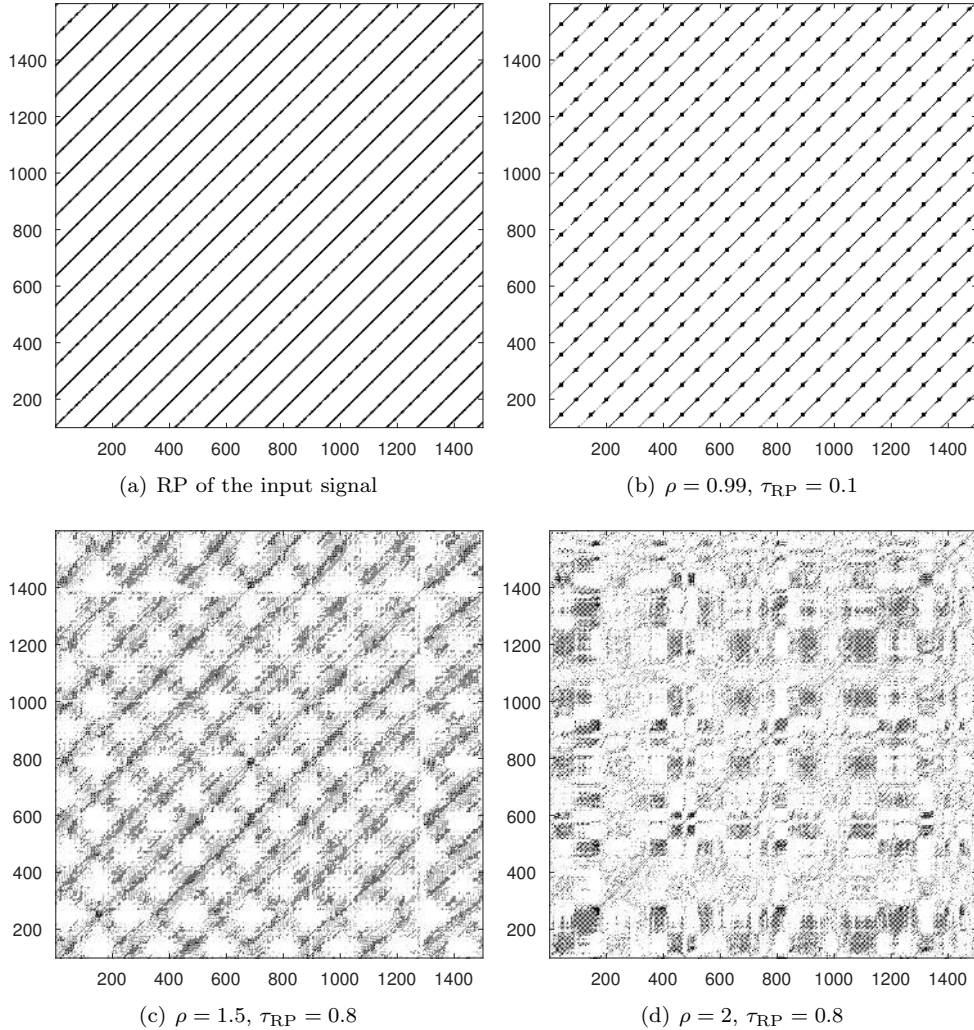


Figure 5: RP for the input signal and the sequence of states of the reservoir. When $\rho = 0.99$ the activations are compatible with the input dynamics. When ρ exceeds one, the activations denote instability.

respectively. In both cases we observe typical line patterns of unstable systems. However, we note that when $\rho = 2$ the system becomes fully chaotic (e.g., see 4(d) for a visual comparison), with very short and erratic diagonal lines.

4.2 Characterization of reservoirs with RQA

Here we show how RQA measures can be exploited to identify the configuration of the ESN hyper-parameters with maximum computational capability. We evaluate the accuracy of our methodology by considering sinusoidal and MG time-series. We compare the results with λ (5) and the minimal singular value of the Jacobian, η , whose average value is computed over time; both are widely used in the literature for identifying an optimal network setting in an unsupervised way. In particular, we consider the regions of the hyper-parameter space where λ crosses zero and where η is maximized as discriminators to determine the (input-dependent) edge of stability of the network

The error measure for prediction is the Normalized Root Mean Squared Error (NRMSE),

$$\text{NRMSE} = \sqrt{\frac{\langle \|\mathbf{y} - \mathbf{d}\|^2 \rangle}{\langle \|\mathbf{y} - \langle \mathbf{d} \rangle\|^2 \rangle}}, \quad (19)$$

being \mathbf{y} the ESN output (2) and \mathbf{d} the desired one. The prediction accuracy, γ , is defined as $\gamma =$

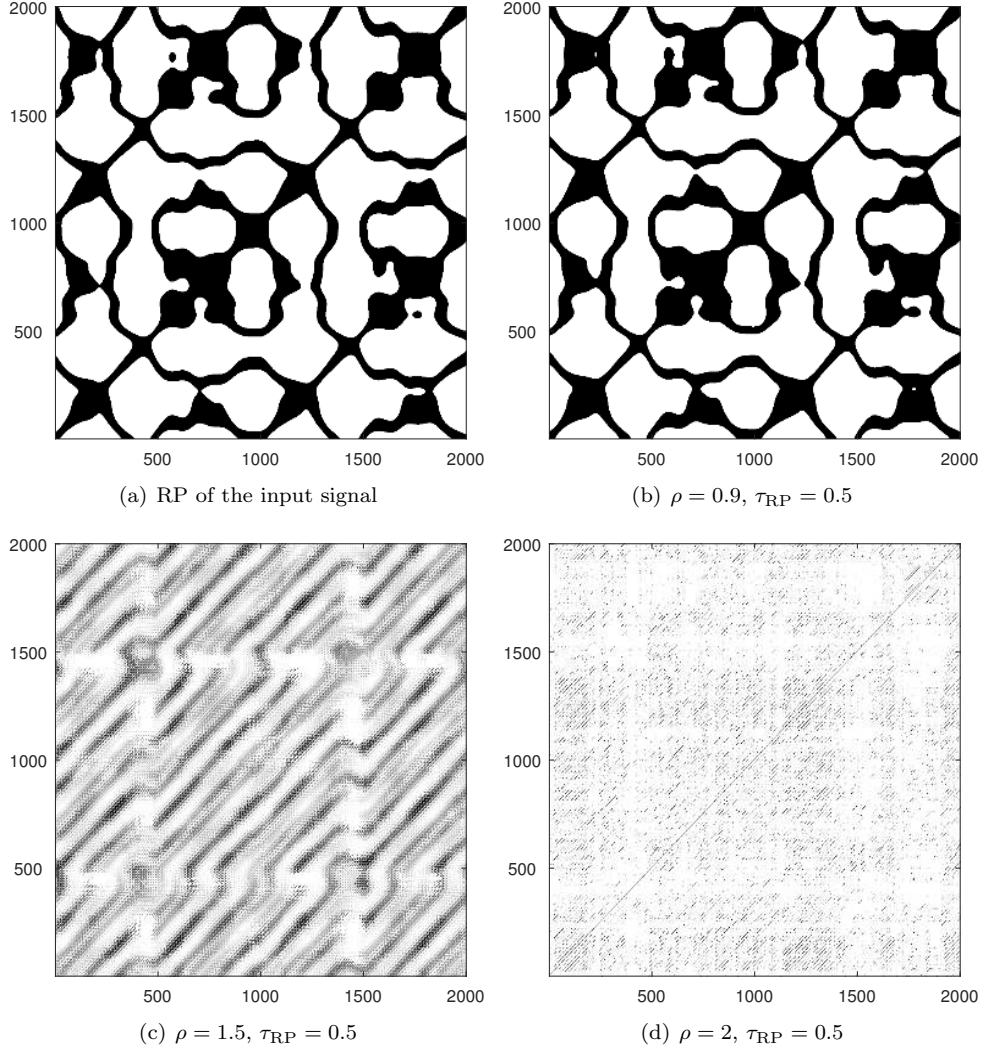


Figure 6: RPs of the MG time-series and of the state of the reservoir.

$\max\{0, 1 - \text{NRMSE}\}$.

In the following, we monitor how γ , λ , η , and RQA measures vary as the values of ρ and ω_i are changed. First, we compare λ with L_{\max} , the value of the longest diagonal line in an RP – see Eq. 10 and related discussion. L_{\max} is a global indicator of stability that we show here to be highly correlated with λ . Successively, we show that the edge of stability determined by using Eq. 15 produces results that are significantly better (in a statistical sense) than those derived by using either λ or η . In order to obtain more interpretable RQA values varying in the unit interval, we normalize all RQA measures by using a unity-based normalization separately for each measure, which considers every setting of ρ and ω_i for each signal. It is important to note that this operation neglects the possibility to perform a more in-depth analysis on the RQA measures – which we remind is not the scope of this paper. However, this choice allows the network designer to immediately recognize areas of interest and transitions in the related plots. We always use 50 bins for computing the SWRP measure (14). All results are always reported as average of 15 different simulations with independent initialization of \mathbf{W}_i^r and \mathbf{W}_r^r in the ESN. For this second experiment, since the number of reservoir neurons has an impact on the prediction performance, we adopted two different configurations for N_r .

Sinusoidal signal. The ESN is trained to perform a 25-step ahead prediction and it is tested on a time-series of 1500 time-steps; we used a reservoir size $N_r = 75$. We test the network performance by varying $\rho \in P = [0.01, 2]$ and $\omega_i \in \Omega = [0.01, 1]$, both discretized with resolution 0.1.

In Fig. 7, we show the RQA measures calculated by varying ρ and ω_i . The ESN, when $\rho \leq 1$, shows high levels of regularity and stability, as expressed by RR, DET, LAM, and L_{\max} . Notably, when $\rho \simeq 1$, we note that L_{\max} starts to decrease, suggesting that we are approaching the unstable phase. ENTR and SWRP behave similarly, although they are maximized for different configurations of ρ and ω_i . In fact, ENTR assumes the relative maximum value for larger input scaling when $\rho \simeq 1$. On the other hand, SWRP assumes high values for smaller values of ω_i and ρ .

Fig. 8 shows a 2D section of Fig. 7, obtained by selecting a specific input scaling $\omega_i = 0.8$. In this way, we can visually assess the agreement between L_{\max} and λ , and also the increasing RQA variability (i.e., standard deviation) on the edge of stability. In Fig. 8(a), we reported λ , L_{\max} , and DIV (11) as we vary ρ . In order to improve the visualization, λ has been rescaled by dividing it for its maximum (absolute) value. λ and L_{\max} are anticorrelated with (Pearson) correlation equal to -0.78 : the value of L_{\max} decreases as ρ increases, while λ , as expected, increases with ρ . Additionally, we can observe that there exists a positive correlation (0.79) between λ and DIV. We stress that the agreement between λ and L_{\max} is consistent for the entire range of ω_i , as represented in Fig. 9(d) and quantified in Tab. 1, confirming that statistics of the RP diagonal lines offer consistent and solid complexity measures characterizing the network stability. Please, notice that the correlations in Tab. 1 are computed on two matrices in $\mathbb{R}^{|P|} \times \mathbb{R}^{|\Omega|}$. Since the matrix elements are not interrelated, such matrices can be represented as vectors having dimension $|P| \cdot |\Omega|$. The correlation is then computed on such vectors.

Let us comment the results for RR, DET, LAM, ENTR, and SWRP with respect to determination of the edge of stability. All these RQA measures denote low fluctuations until λ becomes positive – see related panels in Fig. 8 – which is achieved for $\rho \simeq 1.2$. γ s are shown in Fig. 9(a), while in Fig. 9(e) we present different lines in the ρ - ω_i plane, each one denoting a specific determination of the edge of stability (by using either λ and the RQA measures). In the same plot, with a solid red line we indicate the configurations where the best performance is achieved. First, we observe that, in general, for low values of ω_i we obtain a slightly inferior accuracy. In addition, the gray area in Fig. 9(a) is larger for higher values of ω_i and it includes also several configurations with $\rho = 1.2$. This is justified by the fact that high-amplitude signals tend to saturate the nonlinear activation function and cause the poles to shrink toward the origin. This results in a system with a larger stability margin and a more contractive dynamics, hence justifying the use of larger values of ρ .

To quantify the accuracy in identifying the edge of stability, we report in Tab. 2 the average distances of the edges determined through RQA measures with respect to the maximum values of γ (solid red line in Fig. 9(e)). Our results show that, in all cases except for SWRP, RQA measures determine a more accurate edge of stability with respect to the one obtainable by relying on λ and η . Notably, RR produces differences with γ that are significantly smaller than those of λ and η ($p < 0.0001$); the same holds for DET, LAM, and ENTR ($p < 0.0001$), but not for SWRP, whose outcome is significantly worse than λ and η ($p < 0.0001$). Such results have been obtained by using the t -test with 0.05 as significance threshold. The important differences with respect to λ might be related to its low sensitivity on ω_i , as can be noticed by looking at Fig. 9(b). On the other hand, η is more sensitive to ω_i than λ , as can be noted in Fig. 9(c). In fact, differences in Tab. 2 between η and λ are statistically significant ($p < 0.0075$).

Let us discuss the results in Tab. 3 in terms of correlation between RQA measures and γ . For the sinusoidal input, correlations of RQA measures with γ are not very satisfactory (although all correlations are significantly different from zero), with the only exception of SWRP. However, we note that the estimated (normalized) mutual information indicates, in general, the presence of relevant nonlinear dependencies.

MG time-series. The prediction of the MG time-series is a common benchmark where ESNs achieved good performance in terms of prediction accuracy [47, 48]. For this test, we train the ESN to perform a 20-step ahead prediction on a test time-series of 2000 time-steps; we set the size of the reservoir to $N_r = 100$; intervals P and Ω are defined as before.

In Fig. 10, we show the outcomes for six RQA measures as both ρ and ω_i change. With respect to the sinusoidal input, we immediately observe a higher sensitivity with respect to ω_i . This can be explained by the need to use a higher degree of non-linearity for dealing with the forecast of the MG time-series. The dynamics of the reservoir appear as highly deterministic with laminar phases (as expressed by DET and LAM, respectively) for the entire range of ω_i . We note a consistent maximization for SWRP and ENTR, which assume high values for a large set of configurations that include also values for $\rho \simeq 2$ when $\omega_i > 0.5$. For what concerns RR, instead, we notice that it is maximized only for very small values of ρ and high values of ω_i , suggesting that density of recurrences (i.e., RR) is very sensitive to ρ for the MG




Figure 7: RQA measures for sinusoid while changing ρ (vertical axis) and ω_i (horizontal axis).

time-series.

In Fig. 11, we show the mean and standard deviation of the RQA values when $\omega_i = 0.5$ is fixed. This is performed, as before, in order to assess the network stability and locate the edge of stability in terms of increase in the fluctuations of the RQA measures. In Fig. 11(a) we show the mean value of λ , which has been graphically rescaled also in this case, L_{\max} , and DIV as we vary ρ . As for the sinusoidal input, λ and L_{\max} show a good anticorrelation, with a value of -0.64 . Analogously, λ and DIV are correlated with a slightly lower value of 0.60 . The agreement between L_{\max} and λ is confirmed for all values of ω_i , as depicted in Fig. 12(d) and numerically quantified in the second row of Tab. 1.

Let us now take into account the issue of determining the edge of stability. In Fig. 12(a), we notice that high prediction accuracy γ is obtained for configurations located in the center of the plot toward the right-hand side, corresponding to $\rho \simeq 1.5$ and $\omega_i \geq 0.8$. For instance, when $\omega_i = 0.5$ the best accuracy is obtained for $1 \leq \rho \leq 1.3$. When considering the specific case of $\omega_i = 0.5$, λ becomes greater than 0 when $\rho \geq 0.7$, as shown in Fig. 11(a). However, we note that all RQA measures start to fluctuate for much higher values of ρ , close to 1.5. This is graphically represented in Fig. 12(e). As clearly shown by the plot, λ significantly underestimates the location of the edge of stability for almost all values of $\omega_i > 0.2$. On the other hand, we note that all RQA measures show a very good correlation with the red line denoting the best performance, γ , even though they slightly overestimate the edge. This comparison is made formal in Tab. 2, where we show that all RQA measures denote smaller distances with γ than both λ and η , with differences that are statistically significant. Notably, p -values are as follows: for RR, LAM, ENTR, and SWRP $p < 0.0001$, and for DET $p < 0.0010$. Therefore, in the case of the MG time-series, the information provided by RQA measures regarding the edge of stability is significantly more accurate than both λ and η . As in the previous case, we note a low sensitivity of λ with respect to the input scaling; see Fig. 12(b). The results obtained for η (Fig. 12(c)) show that, instead, this criterion significantly overestimate the edge of stability, especially for values of $\omega_i \in [0.3, 0.5]$. Finally, it is worth to notice that differences between λ and η are statistically significant ($p < 0.0001$), showing that η performs better also in the MG case with respect to λ .

Let us now take into account the correlations between prediction accuracy and RQA measures. As shown in Tab. 3, we experience a good linear agreement between SWRP and γ , denoting a correlation

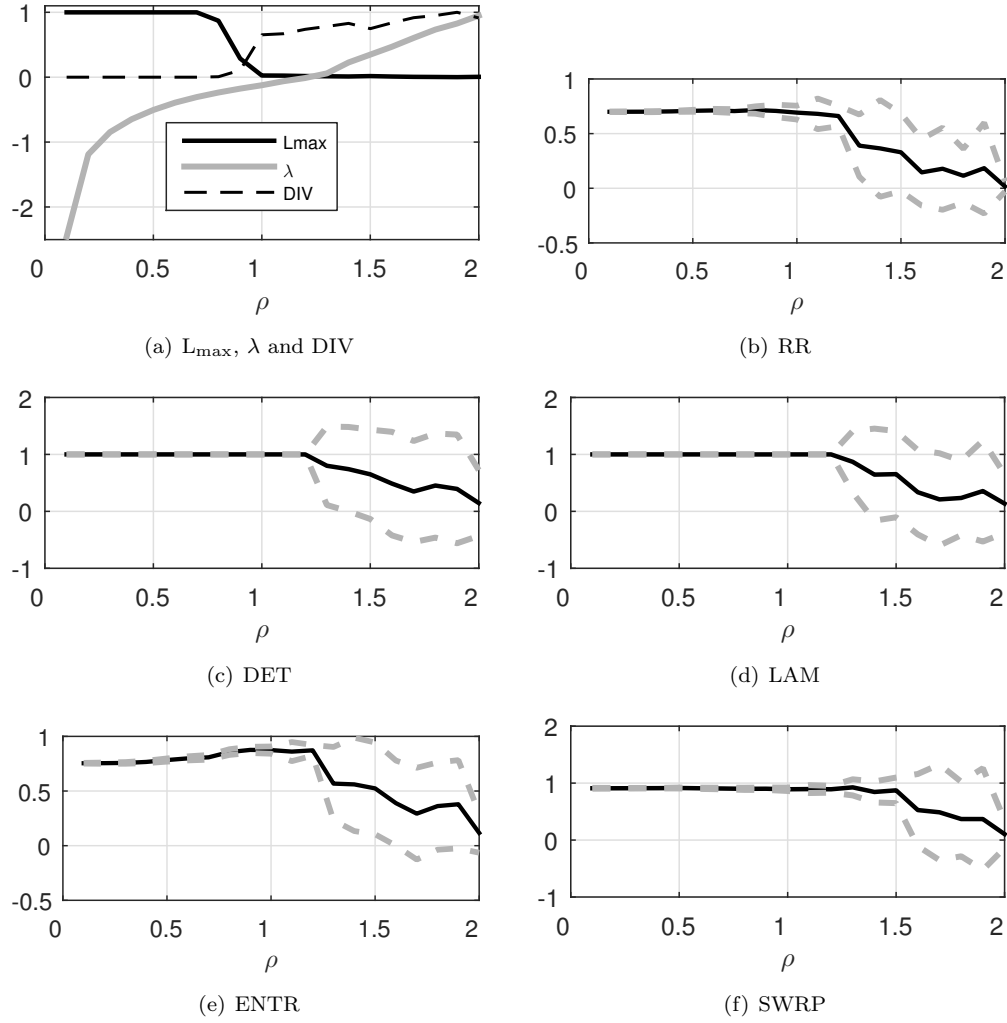


Figure 8: RQA measures of sinusoidal input for a fixed input scaling, $\omega_i = 0.8$. The mean value of each measure is drawn with a solid black line and the standard deviation with gray dashed lines. All RQA measures are stable around their mean values until ρ is pushed beyond 1. Fig. 8(a) depicts the rescaled, mean value of λ (gray solid line), the mean value of L_{\max} (solid black line), and the mean value of DIV (dashed black line).

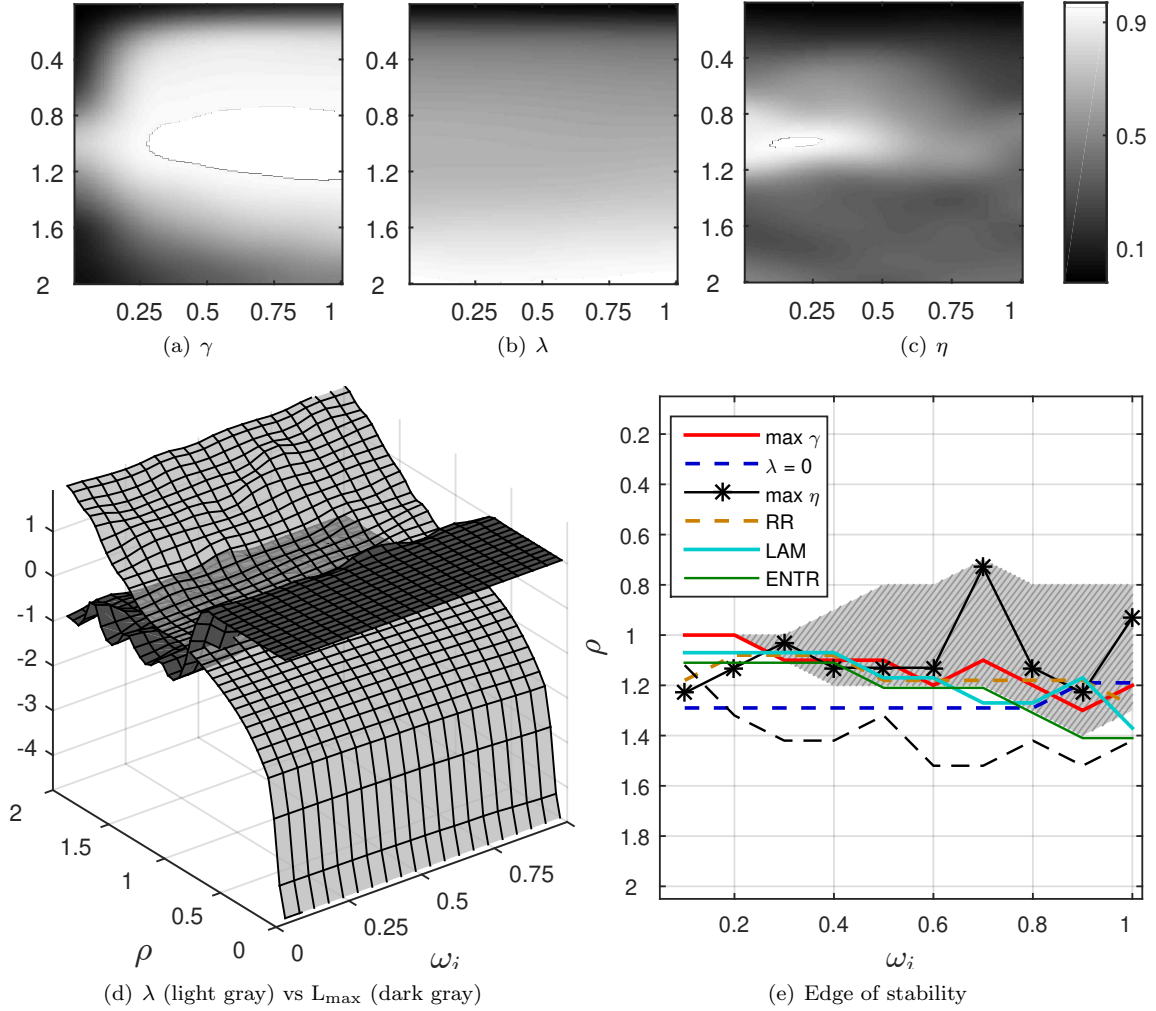


Figure 9: Sinusoidal input. Prediction accuracy γ in Fig. 9(a), λ in Fig. 9(b), and η in Fig. 9(c), calculated for different values of ρ and ω_i . In Fig. 9(d) λ (light gray surface) assumes positive values in correspondence of the gap where L_{\max} (dark gray surface) goes to 0. In Fig. 9(e), the edge of stability is identified as the configurations where λ crosses 0, η is maximized, and with RQA measures according to Eq. 15). Please notice that, for the sake of readability, here we show only three RQA measures; see Tab. 2 for detailed results. The gray area in Fig. 9(e) shows configurations with high accuracy and it corresponds to the white area in Fig. 9(a) where γ assumes larger values; the red line where γ is maximum.

of 0.68. Also ENTR achieves a good correlation of 0.62, as in fact ENTR and SWRP behave similarly for the MG time-series (see Fig. 10). Nonlinear dependency, as indicated by mutual information, shows high values for RR, ENTR, and SWRP.

	$\lambda \setminus L_{\max}$	$\lambda \setminus \text{DIV}$
Sin	-0.74	0.53
MG	-0.65	0.57

Table 1: Correlations between λ , DIV, and L_{\max} .

Table 2: Average distances and standard deviations between the configurations where γ is maximized and the indicators used to determine the edge of stability: η is maximum, λ crosses 0, and $\sigma_q \geq \bar{\sigma}_q$ for RQA measures. Results for RQA shown in bold are statistically significant.

	η	λ	RR	DET	LAM	ENTR	SWRP
Sin	1.4±1.26	1.7±0.95	0.7±0.67	1.0±0.82	0.9±0.74	0.8±0.63	2.5±0.85
MG	3.7±2.21	7.3±4.78	2.0±0.94	3.1±1.29	2.6±1.35	2.9±1.29	3.0±0.94

Table 3: Dependency between γ and RQA measures. First row shows correlations; second row estimated mutual information.

	RR	DET	LAM	ENTR	SWRP
Sin	0.35	0.42	0.40	0.43	0.55
Sin	0.66	0.69	0.67	0.68	0.81
MG	0.57	0.40	0.38	0.62	0.68
MG	0.77	0.55	0.54	0.78	0.80




Figure 10: RQA measures of MG while changing ρ (vertical axis) and ω_i (horizontal axis).

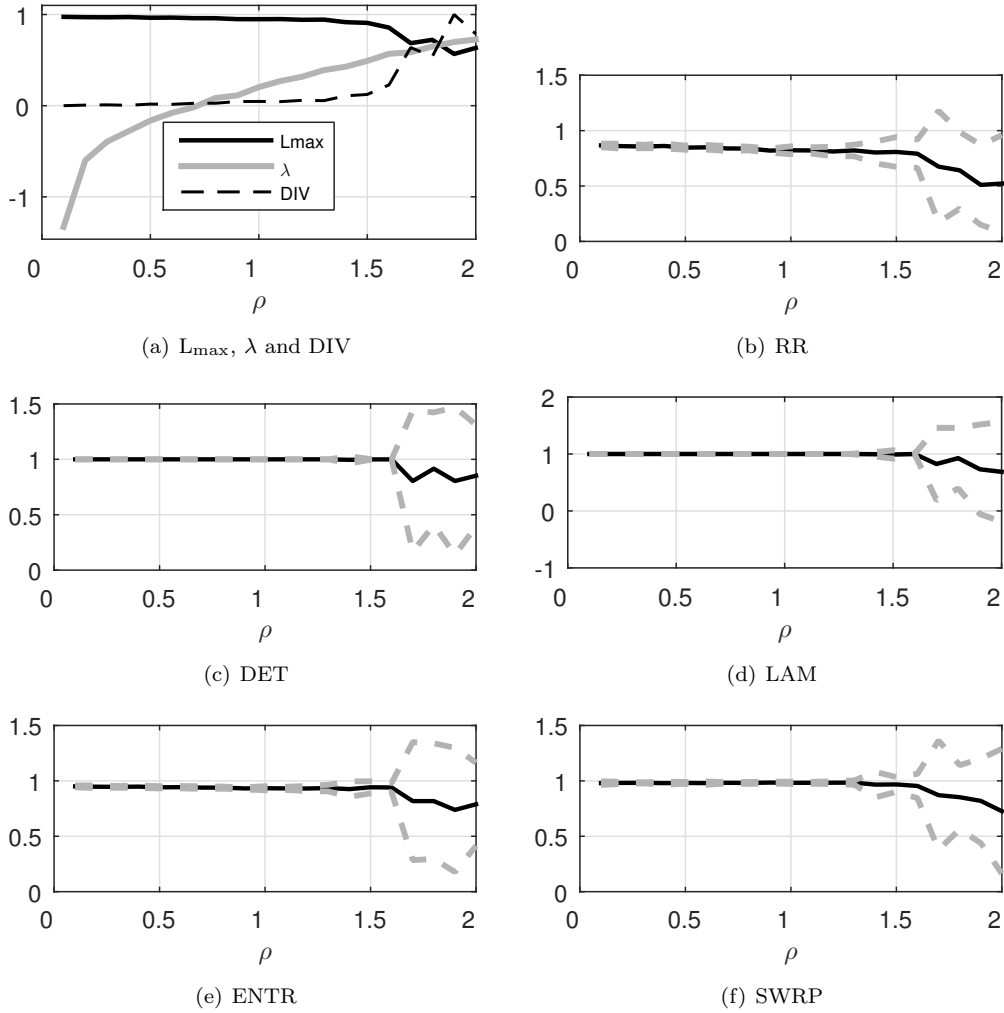


Figure 11: RQA measures for MG by considering a fixed input scaling, $\omega_i = 0.5$. Solid black lines represent the mean value of each measure; gray dashed lines the standard deviation. The standard deviation of every RQA increases significantly as soon as the system enters into unstable regime. In Fig. 11(a), we report the rescaled, mean value of λ (gray solid line), the mean value of L_{\max} (solid black line), and the mean value of DIV (dashed black line).

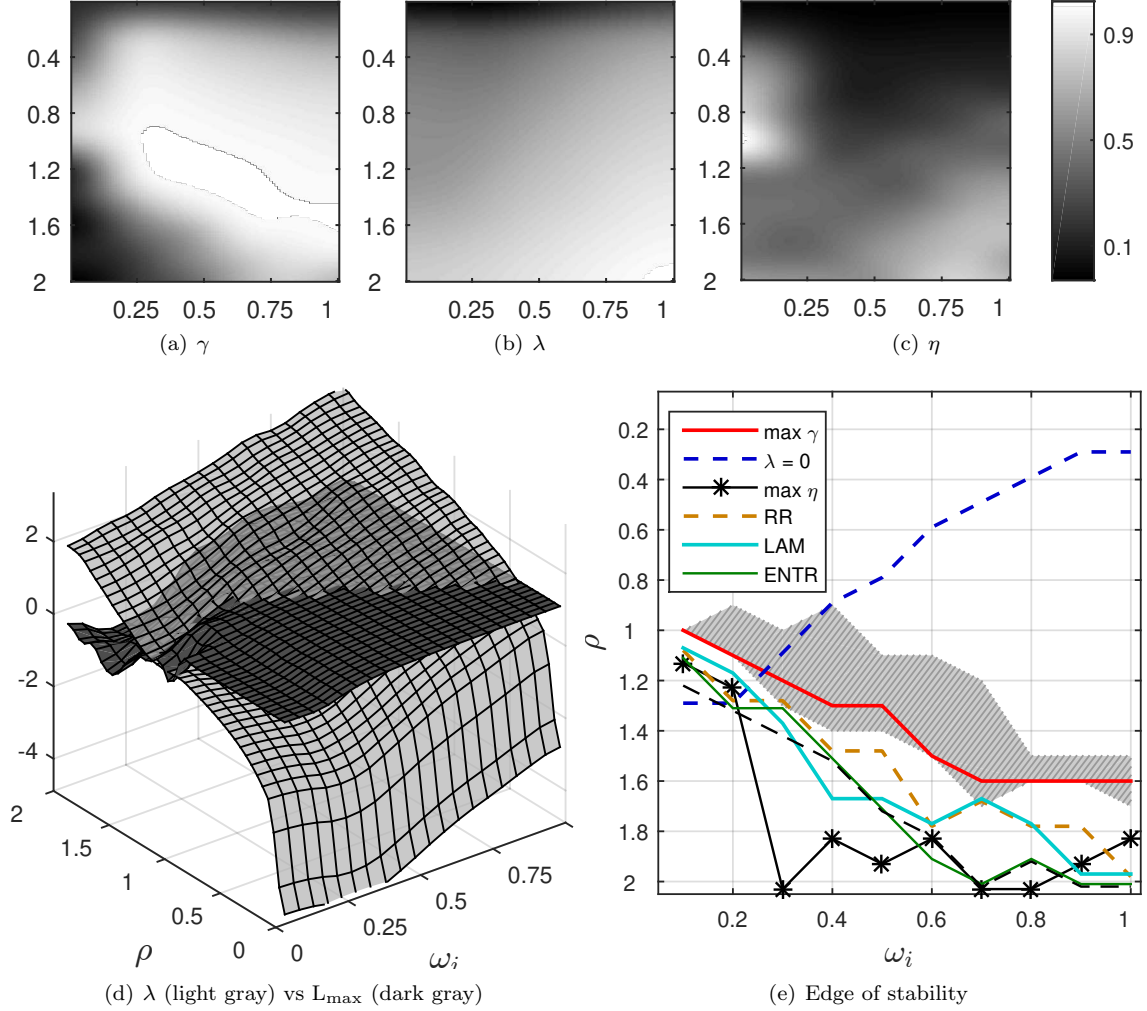


Figure 12: MG time-series. Prediction accuracy γ in Fig. 12(a), λ in Fig. 12(b), and η in Fig. 12(c), calculated for different values of ρ and ω_i . In Fig. 12(d) λ (light gray surface) becomes positive when the value of L_{\max} (dark gray surface) is drastically reduced to a value close to zero. In Fig. 12(e), the edge of stability is identified as the configurations where λ crosses 0, η is maximized, and for RQA measures according to Eq. 15. Please notice that, for the sake of readability, here we show only three RQA measures; see Tab. 2 for detailed results. The gray area in Fig. 12(e) shows configurations with high accuracy and it corresponds to the white area in Fig. 12(a) where γ assumes larger values; the red line the values for which γ is maximum.

5 Conclusions

In this paper, we have investigated the possibility to analyze the dynamics of echo state networks by means of recurrence plots and related RQA complexity measures. Being a non-autonomous system, its dynamical properties depend not only on the spectral radius of the reservoir matrix, but also on the particular input driving the network and the related scaling coefficients. We have shown how RPs can be used for reliably analyzing the state evolution of the network when qualitatively different signals are presented as input. Our results indicate that the recurrences elaborated from the reservoir dynamics clearly denoted an instability level monotonically related to the instability of the network. This coherent behavior suggests that RPs could be used as a visual tool to design a network.

Our results highlighted a strong statistical agreement between the maximal local Lyapunov exponent and L_{\max} , an RQA complexity measure based on the diagonal lines in an RP. This fact could be exploited to define novel, input-dependent stability criteria for reservoirs based on statistics of state recurrences. We have also examined prediction problems on two signals: a sinusoidal waveform and the Mackey-Glass time-series. The smooth variation of λ was mainly driven by ρ but not very sensitive to the different values used for ω_i . This fact prevented to accurately identify the input-dependent edge of stability of the network, i.e., the region of the control parameter space where the network departs from a stable regime to reach and operate in an unstable, yet computationally effective phase. We have considered also a related criterion based on the minimum singular value (η) of the Jacobian matrix. This criterion offered, in general, better solutions than λ for determining the edge of stability. However, both criteria were significantly outperformed by the criterion proposed in this paper based on the fluctuations of RQA measures. We observed that, when the network approached the edge of stability, all RQA measures abruptly changed their average value, showing also a rapid increase of fluctuations.

The results presented in this paper suggest to allocate future research effort in a number of ways. First, it would be interesting to study reservoirs initialized according to different (either statistical or deterministic) rules. In addition, it could be interesting to derive a connection between the network short-term memory capacity and RQA measures. It is worth considering invariant measures elaborated from RPs: the correlation entropy provides a natural lower-bound to the sum of positive Lyapunov exponents. It might be interesting to study the joint recurrences, calculated between the input system and the reservoir. This might be useful to discover (hidden) mechanisms of synchronization. Finally, we suggest that such RP-based analyses could be extended to fully trainable recurrent neural networks. The recurrent layer would not be modeled as a multi-dimensional time-series, but instead as a time-varying network. To this end, we hypothesize a possible interplay between graph matching [49] and recurrence analysis.

References

- [1] Y. Zhang and J. Wang, "Global exponential stability of recurrent neural networks for synthesizing linear feedback control systems via pole assignment," *IEEE Transactions on Neural Networks*, vol. 13, no. 3, pp. 633–644, May 2002.
- [2] U. D. Schiller and J. J. Steil, "Analyzing the weight dynamics of recurrent learning algorithms," *Neurocomputing*, vol. 63, pp. 5–23, 2005.
- [3] R. L. Marichal and J. D. Piñeiro, "Analysis of multiple quasi-periodic orbits in recurrent neural networks," *Neurocomputing*, vol. 162, pp. 85–95, 2015.
- [4] D. Sussillo and O. Barak, "Opening the black box: Low-dimensional dynamics in high-dimensional recurrent neural networks," *Neural Computation*, vol. 25, no. 3, pp. 626–649, 2013.
- [5] Y. Shen and J. Wang, "An improved algebraic criterion for global exponential stability of recurrent neural networks with time-varying delays," *IEEE Transactions on Neural Networks*, vol. 19, no. 3, pp. 528–531, Mar. 2008.
- [6] D. Sussillo, "Neural circuits as computational dynamical systems," *Current Opinion in Neurobiology*, vol. 25, pp. 156–163, 2014.
- [7] D. Verstraeten, B. Schrauwen, M. d’Haene, and D. Stroobandt, "An experimental unification of reservoir computing methods," *Neural Networks*, vol. 20, no. 3, pp. 391–403, 2007.
- [8] M. Lukoševičius and H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Computer Science Review*, vol. 3, no. 3, pp. 127–149, 2009.
- [9] J. Dambre, D. Verstraeten, B. Schrauwen, and S. Massar, "Information processing capacity of dynamical systems," *Scientific Reports*, vol. 2, 2012.
- [10] H. Jaeger, "Adaptive nonlinear system identification with echo state networks," in *Advances in Neural Information Processing Systems*, S. Becker, S. Thrun, and K. Obermayer, Eds. MIT Press, 2003, pp. 593–600.
- [11] T. Strauss, W. Wustlich, and R. Labahn, "Design strategies for weight matrices of echo state networks," *Neural Computation*, vol. 24, no. 12, pp. 3246–3276, 2012.

- [12] C. Gallicchio and A. Micheli, “Architectural and Markovian factors of echo state networks,” *Neural Networks*, vol. 24, no. 5, pp. 440–456, 2011.
- [13] Q. Ma, W. Chen, J. Wei, and Z. Yu, “Direct model of memory properties and the linear reservoir topologies in echo state networks,” *Applied Soft Computing*, vol. 22, pp. 622–628, 2014.
- [14] S. Scardapane, D. Wang, and M. Panella, “A decentralized training algorithm for echo state networks in distributed big data applications,” *Neural Networks*, 2015.
- [15] F. M. Bianchi, S. Scardapane, A. Uncini, A. Rizzi, and A. Sadeghian, “Prediction of telephone calls load using Echo State Network with exogenous variables,” *Neural Networks*, vol. 71, pp. 204–213, 2015.
- [16] Y. Liu, Q. Liu, W. Wang, J. Zhao, and H. Leung, “Data-driven based model for flow prediction of steam system in steel industry,” *Information Sciences*, vol. 193, pp. 104–114, 2012.
- [17] F. M. Bianchi, E. De Santis, A. Rizzi, and A. Sadeghian, “Short-term electric load forecasting using echo state networks and PCA decomposition,” *IEEE Access*, vol. 3, pp. 1931–1943, Oct. 2015.
- [18] S. I. Han and J. M. Lee, “Fuzzy echo state neural networks and funnel dynamic surface control for prescribed performance of a nonlinear dynamic system,” *IEEE Transactions on Industrial Electronics*, vol. 61, no. 2, pp. 1099–1112, Feb. 2014.
- [19] M. D. Skowronski and J. G. Harris, “Automatic speech recognition using a predictive echo state network classifier,” *Neural Networks*, vol. 20, no. 3, pp. 414–423, 2007.
- [20] Y. Xue, L. Yang, and S. Haykin, “Decoupled echo state networks with lateral inhibition,” *Neural Networks*, vol. 20, no. 3, pp. 365–376, 2007.
- [21] L. Boccatto, R. Attux, and F. J. Von Zuben, “Self-organization and lateral interaction in echo state network reservoirs,” *Neurocomputing*, vol. 138, pp. 297–309, 2014.
- [22] B. Zhang, D. J. Miller, and Y. Wang, “Nonlinear system modeling with random matrices: Echo state networks revisited,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 1, pp. 175–182, Jan. 2012.
- [23] M. C. Ozturk, D. Xu, and J. C. Principe, “Analysis and design of echo state networks,” *Neural Computation*, vol. 19, no. 1, pp. 111–138, 2007.
- [24] B. Schrauwen, M. Wardermann, D. Verstraeten, J. J. Steil, and D. Stroobandt, “Improving reservoirs using intrinsic plasticity,” *Neurocomputing*, vol. 71, no. 7, pp. 1159–1171, 2008.
- [25] A. Rodan and P. Tiño, “Simple deterministically constructed cycle reservoirs with regular jumps,” *Neural Computation*, vol. 24, no. 7, pp. 1822–1852, 2012.
- [26] —, “Minimum complexity echo state network,” *IEEE Transactions on Neural Networks*, vol. 22, no. 1, pp. 131–144, Nov. 2011.
- [27] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, “Information processing using a single dynamical node as complex system,” *Nature Communications*, vol. 2, p. 468, 2011.
- [28] L. Büsing, B. Schrauwen, and R. Legenstein, “Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons,” *Neural Computation*, vol. 22, no. 5, pp. 1272–1311, 2010.
- [29] D. Verstraeten and B. Schrauwen, “On the quantification of dynamics in reservoir computing,” in *Artificial Neural Networks – ICANN 2009*, C. Alippi, M. Polycarpou, C. Panayiotou, and G. Ellinas, Eds. Springer Berlin, Heidelberg, 2009, vol. 5768, pp. 985–994.
- [30] J. Boedecker, O. Obst, J. T. Lizier, N. M. Mayer, and M. Asada, “Information processing in echo state networks at the edge of chaos,” *Theory in Biosciences*, vol. 131, no. 3, pp. 205–213, 2012.
- [31] N. Marwan, M. Carmen Romano, M. Thiel, and J. Kurths, “Recurrence plots for the analysis of complex systems,” *Physics Reports*, vol. 438, no. 5, pp. 237–329, 2007.
- [32] N. Marwan, “How to avoid potential pitfalls in recurrence plot based data analysis,” *International Journal of Bifurcation and Chaos*, vol. 21, no. 4, pp. 1003–1017, 2011.
- [33] D. Eroglu, T. K. D. M. Peron, N. Marwan, F. A. Rodrigues, L. da F. Costa, M. Sebek, I. Z. Kiss, and J. Kurths, “Entropy of weighted recurrence plots,” *Physical Review E*, vol. 90, no. 4, p. 042919, 2014.
- [34] N. Marwan, S. Schinkel, and J. Kurths, “Recurrence plots 25 years later—gaining confidence in dynamical transitions,” *EPL (Europhysics Letters)*, vol. 101, no. 2, p. 20007, 2013.
- [35] N. Marwan and J. Kurths, “Line structures in recurrence plots,” *Physics Letters A*, vol. 336, no. 4, pp. 349–357, 2005.
- [36] H. Yang and Y. Chen, “Heterogeneous recurrence monitoring and control of nonlinear stochastic processes,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 24, no. 1, p. 013138, 2014.
- [37] R. V. Donner, Y. Zou, J. F. Donges, N. Marwan, and J. Kurths, “Recurrence networks—a novel paradigm for nonlinear time series analysis,” *New Journal of Physics*, vol. 12, no. 3, p. 033025, 2010.
- [38] H. 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*, vol. 148, p. 34, 2001.
- [39] I. B. Yildiz, H. Jaeger, and S. J. Kiebel, “Re-visiting the echo state property,” *Neural Networks*, vol. 35, pp. 1–9, 2012.
- [40] G. Manjunath and H. Jaeger, “Echo state property linked to an input: Exploring a fundamental characteristic of recurrent neural networks,” *Neural Computation*, vol. 25, no. 3, pp. 671–696, 2013.
- [41] D. Verstraeten, B. Schrauwen, and D. Stroobandt, “Reservoir-based techniques for speech recognition,” in *Proceedings of the IEEE International Joint Conference on Neural Networks*, Vancouver, BC, Canada, Jul. 2006, pp. 1050–1053.
- [42] C. G. Langton, “Computation at the edge of chaos: Phase transitions and emergent computation,” *Physica D: Nonlinear Phenomena*, vol. 42, no. 1, pp. 12–37, 1990.
- [43] N. Bertschinger and T. Natschläger, “Real-time computation at the edge of chaos in recurrent neural networks,” *Neural Computation*, vol. 16, no. 7, pp. 1413–1436, 2004.
- [44] R. Legenstein and W. Maass, “Edge of chaos and prediction of computational performance for neural circuit models,” *Neural Networks*, vol. 20, no. 3, pp. 323–334, 2007.

- [45] H. Jaeger, *Short term memory in echo state networks*. GMD-Forschungszentrum Informationstechnik, 2001.
- [46] J. J. Steil, "Memory in backpropagation-decorrelation o(n) efficient online recurrent learning," in *Artificial Neural Networks: Formal Models and Their Applications-ICANN 2005*, W. Duch, J. Kacprzyk, E. Oja, and S. Zadrożny, Eds. Berlin, Heidelberg: Springer, 2005, pp. 649–654.
- [47] Z. Shi and M. Han, "Support vector echo-state machine for chaotic time-series prediction," *IEEE Transactions on Neural Networks*, vol. 18, no. 2, pp. 359–372, Mar. 2007.
- [48] H. Jaeger and H. Haas, "Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication," *Science*, vol. 304, no. 5667, pp. 78–80, 2004.
- [49] L. Livi and A. Rizzi, "The graph matching problem," *Pattern Analysis and Applications*, vol. 16, no. 3, pp. 253–283, 2013.