DETECTING INVARIANT MANIFOLDS IN RELU-BASED RNNS

Anonymous authorsPaper under double-blind review

000

001

002003004

006

008 009

010 011

012

013

014

016

017

018

019

021

024

025

026

027

028

031

033

034

037

038

040

041

042

043

044

045

047

048

051

052

ABSTRACT

Recurrent Neural Networks (RNNs) have found widespread applications in machine learning for time series prediction and dynamical systems reconstruction, and experienced a recent renaissance with improved training algorithms and architectural designs. Understanding why and how trained RNNs produce their behavior is important for scientific and medical applications, and explainable AI more generally. An RNN's dynamical repertoire depends on the topological and geometrical properties of its state space. Stable and unstable manifolds of periodic points play a particularly important role: They dissect a dynamical system's state space into different basins of attraction, and their intersections lead to chaotic dynamics with fractal geometry. Here we introduce a novel algorithm for detecting these manifolds, with a focus on piecewise-linear RNNs (PLRNNs) employing rectified linear units (ReLUs) as their activation function. We demonstrate how the algorithm can be used to trace the boundaries between different basins of attraction, and hence to characterize multistability, a computationally important property. We further show its utility in finding so-called homoclinic points, the intersections between stable and unstable manifolds, and thus establish the existence of chaos in PLRNNs. Finally we show for an empirical example, electrophysiological recordings from a cortical neuron, how insights into the underlying dynamics could be gained through our method.

1 Introduction

Recurrent neural networks (RNNs) are widely employed for time series forecasting (Park et al., 2018; Gu and Dao, 2023) and dynamical systems (DS) reconstruction (Hess et al., 2023; Brenner et al., 2022; 2024b; Platt et al., 2023), especially in scientific applications like climate modeling (Patel and Ott, 2023) or neuroscience (Durstewitz et al., 2023), as well as in medical domains. RNNs experienced a recent revival due to the advance of new powerful training algorithms that avoid vanishing or exploding gradients (Hess et al., 2023), and novel network architectures (Schmidt et al., 2021; Rusch and Mishra, 2020; Peng et al., 2023; Gu and Dao, 2023; Rusch et al.), in particular in the context of state space models (which are essentially linear RNNs with nonlinear readouts and input gating; (Gu and Dao, 2023; Orvieto et al., 2023)). What lags behind, like in many other areas of deep learning, is a thorough theoretical understanding of the behavior of these systems and how they achieve the tasks they were trained on. Yet, such an understanding is crucial especially in scientific and medical areas where we are often interested in using trained RNNs as surrogate models for the underlying DS, providing mechanistic insight into the dynamical processes that generated the observed time series.

Formally, RNNs are recursive maps, hence discrete-time dynamical systems (Mikhaeil et al., 2022). Dynamical systems theory (DST) offers a rich repertoire of mathematical tools for analyzing the behavior of such systems (Guckenheimer and Holmes, 2013). Yet, the exploration of high-dimensional RNN dynamics remains a challenge due to limitations in current numerical methods, which struggle to scale and tend to yield only approximate results (Katz and Reggia, 2018; Golub and Sussillo, 2018).

The dynamical behavior of a system is governed by its state space topology and geometry (Hasselblatt and Katok, 2002), most prominently topological objects like attractors, such as stable fixed points, cycles, or chaotic sets, which determine the system's long-term behavior. Similarly important are the

stable and unstable manifolds of fixed points and cycles, although they received much less attention in the scientific and ML communities. Stable manifolds are the sets of points in a system's state space that converge towards an equilibrium or periodic orbit in forward-time, while unstable manifolds, conversely, are the sets of points that converge to an equilibrium or periodic orbit in backward-time. Stable manifolds of saddle points delineate the boundaries between basins of attraction in multistable systems, that is, systems that harbor multiple attractor objects between which it can be driven back and forth by perturbations like external inputs or noise (Feudel et al., 2018). Multistability has been hypothesized to be an important computational property of RNNs, most prominently in computational neuroscience where it has been linked to working memory (Durstewitz et al., 2000) or decision making (Wang, 2008). Different attractors may, for instance, correspond to different active memory states, and the process of convergence to one of these attractors to memory retrieval and pattern completion, as in Hopfield networks (Hopfield, 1982).

Tracing out stable and unstable manifolds is also important for finding homo- and heteroclinic orbits, which connect a cyclic point to itself or to another such point, respectively. These orbits provide a skeleton for the dynamics, forming structures like separatrix cycles (Perko, 2001) and heteroclinic channels (Rabinovich et al., 2008) which have been implied in flexible sequence generation. Intersections between stable and unstable manifolds also give rise to so-called homoclinic points which create sensitive regions in a system's state space associated with chaos (Wiggins, 1988). Chaotic behavior, in turn, or regimes at the edge-of-chaos, have been associated with increased expressivity (larger function classes that can be emulated) and computational power in RNNs (Siegelmann and Sontag, 1992; Bertschinger and Natschläger, 2004; Pereira-Obilinovic et al., 2023). Chaos is also a practical problem in training RNNs, since it causes exploding gradients if not taken care of by special training techniques (Mikhaeil et al., 2022; Hess et al., 2023).

Here we introduce a novel algorithm for computing the stable and unstable manifolds of fixed and cyclic points for the class of piecewise-linear (PL) RNNs (PLRNNs), which use piecewise nonlinearities like the rectified-linear unit (ReLU) as their activation function. Unlike traditional techniques designed for smooth dynamical systems—such as numerical continuation methods (Krauskopf et al., 2007) —the proposed method explicitly exploits the piecewise-linear structure of ReLU-based PLRNNs to enable the exact location of stable and unstable manifolds. In contrast to smooth systems, methods for discrete-time systems are much more scarce, especially if these involve discontinuities in their Jacobians like ReLU-based RNNs. PLRNNs have emerged as one of the most powerful classes of models for dynamical systems reconstruction (Hess et al., 2023; Brenner et al., 2022; 2024b; Nassar et al., 2019), partly because the linear subspaces of such models support the indefinite online retention of memory contents without exploding or vanishing gradients (Schmidt et al., 2021; Orvieto et al., 2023; Gu and Dao, 2023), and partly because their PL structure makes them more tractable (Coombes et al., 2024). For instance, efficient algorithms for exactly localizing fixed and cyclic points in polynomial time exist (Eisenmann et al., 2023), on which we will partly build here. We illustrate how the algorithm can be employed to delineate the boundaries of basins of attraction and to detect homoclinic intersections, thus establishing the existence of chaos, and show how it can be used on empirical data to gain insight into dynamical mechanisms.

2 RELATED WORK

While there is an extensive literature by now on dynamical systems reconstruction with RNNs (Brenner et al., 2022; 2024a; Hess et al., 2023; Patel and Ott, 2023; Platt et al., 2023; Cestnik and Abel, 2019), work on algorithms for dissecting topological properties of trained systems is much more scarce. Existing research is almost exclusively focused on algorithms for locating stable and unstable fixed points and cycles (Golub and Sussillo, 2018; Eisenmann et al., 2023), but not the invariant manifolds associated with them, despite their crucial role in structuring the state space and dynamics. More generally in the DST literature, a related class of methods are so-called continuation methods which numerically trace back important curves and manifolds in dynamical systems, mostly for systems of ordinary differential equations (Krauskopf et al., 2007; Osinga, 2014). Recently methods have been developed specifically for PL maps (Simpson, 2023), but, similar to continuation methods, they seriously suffer from the curse of dimensionality. Thus, current methods effectively work only for very low-dimensional ($\leq 5d$) systems, in contrast to the size of many modern RNNs used for time series forecasting or dynamical systems reconstruction. The specific structure of PL maps, of which all types of PLRNNs are specific examples, considerably eases the computation of certain topological properties. This also provided the basis for the SCYFI algorithm for locating fixed points and cycles,

which scales polynomially, in fact often linearly, with the dimensionality of the PLRNN's latent space (Eisenmann et al., 2023). To our knowledge, no existing method efficiently detects stable and unstable manifolds in discrete-time RNNs. Our approach fills this gap by leveraging the piecewise-linear nature of ReLU-based RNNs to construct manifolds directly, bypassing the limitations of traditional techniques.

3 METHODS

3.1 PLRNN ARCHITECTURES

A PLRNN (Durstewitz, 2017) is simply a ReLU-based RNN with an additional linear term, which takes the basic form

$$\boldsymbol{z}_{t} = \boldsymbol{A}\boldsymbol{z}_{t-1} + \boldsymbol{W}\Phi(\boldsymbol{z}_{t-1}) + \boldsymbol{h}, \tag{1}$$

where $A \in \mathbb{R}^{M \times M}$ is a diagonal matrix, $W \in \mathbb{R}^{M \times M}$ an off-diagonal matrix, the element-wise ReLU non-linearity $\Phi(\cdot) = \max(0, \cdot)$, and $h \in \mathbb{R}^M$ a bias term. Several variants of the PLRNN have been introduced to enhance its expressivity or reduce its dimensionality (Brenner et al., 2022), of which we picked for demonstration here the shallow PLRNN (shPLRNN; (Hess et al., 2023)), $z_t = Az_{t-1} + W_1\Phi(W_2z_{t-1} + h_2) + h_1$ with $W_1 \in \mathbb{R}^{M \times H}$, $W_2 \in \mathbb{R}^{H \times M}$, and the recently proposed almost-linear RNN (ALRNN; (Brenner et al., 2024a)), which attempts to use as few nonlinearities as needed for the problem at hand, $\Phi(z_t) = [z_{1,t}, \dots, z_{M-P,t}, \max(0, z_{M-P+1,t}), \dots, \max(0, z_{M,t})]$.

The piecewise linear structure of these PLRNNs can be exposed by rewriting the ReLU in Equation (1) as

$$z_t = F_{\theta}(z_{t-1}) = (A + WD_{t-1})z_{t-1} + h,$$
 (2)

 $D_t := \operatorname{diag}(d_t)$ with $d_t = (d_{1,t}, d_{2,t}, \dots, d_{M,t})$ and $d_{i,t} = 0$ if $z_{i,t} \leq 0$ and $d_{i,t} = 1$ otherwise (Monfared and Durstewitz, 2020b). For the ALRNN, we have $d_{1:M-P,t} = 1 \forall t$. Hence, we have 2^M linear subregions for the standard PLRNN, $\leq \sum_{k=0}^M \binom{H}{k}$ for the shPLRNN (?), and 2^P for the ALRNN.

3.2 MATHEMATICAL PRELIMINARIES

Recall that a fixed point of a recursive map $x_t = F(x_{t-1})$ is a point x^* for which $x^* = F(x^*)$, and a cyclic point with period m is a fixed point of the m-times iterated map, i.e. such that $x_{t+m} = F^m(x_t) = x_t$ (hence, a fixed point is a period-1 point). An m-cycle is a periodic sequence $\{x_1^* \dots x_m^*\}$ of such points with all points distinct, $x_i^* \neq x_j^* \ \forall 1 \leq i < j \leq m$.

Definition 1 (Un-/stable manifold). Let $F : \mathbb{R}^M \to \mathbb{R}^M$ be a map and p be a hyperbolic period-m cyclic point of F. The *local stable manifold* of p, $W^s_{loc}(p)$, is defined as

$$W_{loc}^s(\mathbf{p}) := \{ \mathbf{x} \in \mathbb{R}^M : F^{nm}(\mathbf{x}) \to \mathbf{p} \text{ as } n \to \infty \}.$$

The local unstable manifold of p, $W_{loc}^{u}(p)$, is defined as

$$W_{loc}^{u}(\mathbf{p}) := \{ \mathbf{x} \in \mathbb{R}^{M} : F^{-nm}(\mathbf{x}) \to \mathbf{p} \text{ as } n \to \infty \}.$$

The *global stable manifold* is the union of all preimages of the local stable manifold, and the *global unstable manifold* is created by the union of all images (forward iterations) of the local unstable manifold (Patra and Banerjee, 2018):

$$W^{s}(\boldsymbol{p}) := \bigcup_{n=1}^{\infty} F^{-n} \Big(W^{s}_{loc}(\boldsymbol{p}) \Big), \qquad W^{u}(\boldsymbol{p}) := \bigcup_{n=1}^{\infty} F^{n} \Big(W^{u}_{loc}(\boldsymbol{p}) \Big). \tag{3}$$

If the map F is noninvertible (i.e., does not have a unique inverse), the (global) stable manifold could be disconnected, making its computation hard as we need to trace back disconnected sets of points in time to determine it. However, most commonly we aim to approximate smooth continuous-time dynamical systems $\dot{\boldsymbol{x}} = f(\boldsymbol{x})$ in dynamical systems reconstruction through our RNN map F_{θ} . For such systems, the Picard-Lindelöf theorem guarantees uniqueness of solutions and the flow (solution) operator $\phi(t, \boldsymbol{x}_0) = \boldsymbol{x}_0 + \int_0^t f(\boldsymbol{u}) \, \mathrm{d}\boldsymbol{u}$ will be a diffeomorphism (Perko, 2001), hence invertible, such that it is reasonable to assume (or enforce) invertibility for F_{θ} as well.

In each locally linear region of a PL dynamical system, the state space decomposes into invariant subspaces for the sets of positive and negative eigenvalues, respectively. For real eigenvalues, trajectories lie on hyperplanes on which the motion is directed towards or away from the fixed point in a straight line. More generally, complex eigenvectors can occur and cause curvature, or eigenvalues may repeat, and the general dynamics in each affine subregion $(z_t = \tilde{W} z_{t-1} + h, \tilde{W} := A + WD(z_{t-1}))$ is given by

$$z_{t} = \sum_{\substack{j=1\\j\neq i}}^{n} c_{j} \lambda_{j}^{t} v_{j} + \lambda_{i}^{t} \sum_{r=1}^{m} c_{r} \frac{t^{r-1}}{(r-1)!} w_{r} + \sum_{k=0}^{t-1} \tilde{W}^{k} h.$$
 (4)

where the $c_i \in \mathbb{R}$ are determined from the initial condition, v_i are the eigenvectors, w_i are generalized eigenvectors (associated with eigenvalues with geometric multiplicity less than the algebraic, and for multiplicity of m > 1). This formula describes the exponential evolution for all non-defective eigenvalues λ_j and polynomially modulated exponentials for the defective eigenvalue λ_i , see Appx. C for details and special cases (Perko, 2001; Hirsch et al., 2013).

Stable manifolds of saddle objects segregate the state space into different basins of attraction (Ganatra and Banerjee, 2022), which are sets of points from which the state evolves toward one or the other attractor. Formally, they are given by (Alligood et al., 1996):

Definition 2 (Basin of attraction). Let $F: \mathbb{R}^M \to \mathbb{R}^M$ be a map. The *basin of attraction* of an attractor \mathcal{A} is the largest open set $B(\mathcal{A}) \subseteq \mathbb{R}^M$ (containing \mathcal{A}) such that for every point $x \in B(\mathcal{A})$, the iterates of x under the map F converge to \mathcal{A} in the forward limit:

$$B(\mathcal{A}) \ = \ \Big\{ \, oldsymbol{x} \in \mathbb{R}^M : d(F^k(oldsymbol{x}), \mathcal{A})
ightarrow 0 ext{ as } k
ightarrow \infty \Big\},$$

where F^k denotes the k-times iterated map F.

Stable and unstable manifolds are crucially important for the system dynamics not only because they dissect the state space into different regions of flow, but also because their intersections can give rise to complex types of dynamics like heteroclinic channels or chaos (Rabinovich et al., 2008; Perko, 2001).

Definition 3 (Homoclinic orbit). Let p be a saddle fixed point (or saddle cycle) of the map $F: \mathbb{R}^M \to \mathbb{R}^M$. A homoclinic orbit is a trajectory $\mathcal{O}_{hom} = \{x_n\}_{n \in \mathbb{N}}$ that connects p to itself, i.e., $x_n \to p$ as $n \to \pm \infty$. Hence,

$$\mathcal{O}_{hom} \subset W^s(\boldsymbol{p}) \cap W^u(\boldsymbol{p}).$$

If $x \neq p$ is a point where the stable and unstable manifolds of p intersect, then x is referred to as a homoclinic point or intersection (Patra and Banerjee, 2018; Perko, 2001). Such an intersection of the stable and unstable manifolds leads to a horseshoe structure associated with a fractal geometry and chaos (see Appx. Def. 5, (Wiggins, 1988)).

Definition 4 (**Heteroclinic orbit**). Let p and q be two *distinct* saddle fixed points (or saddle cycles) of the map F. A *heteroclinic orbit* $\mathcal{O}_{het} = \{x_n\}_{n \in \mathbb{N}}$ connects two different fixed points p and q, that is, $x_n \to p$ as $n \to +\infty$ and $x_n \to q$ as $n \to -\infty$. Heteroclinic intersections (or points) can be defined similarly to homoclinic intersections, and, like homoclinic points, inevitably lead to chaos (Wiggins, 1988; Patra and Banerjee, 2018).

3.3 LOCATING UN-/STABLE MANIFOLDS

For a PL system, locally the unstable (resp. stable) manifolds of a cyclic point p are simply given by the linear subspaces spanned by the eigenvectors with eigenvalues with positive (resp. negative) real part, and hence can easily be computed in closed form from the local Jacobians (A + WD(p)). However, as soon as we cross a border into a different linear subregion, the manifold may fold and start to follow different linear dynamics. For our algorithm, we can still exploit the fact that because the dynamics in the new subregion is again linear, the corresponding piece of the manifold will either be a (hyper-)plane segment or will have curvature along one or more directions according to Eq. 4, and will be of same dimensionality as the segment in the preceding linear subregion. To determine this segment, we only need a few support vectors to correctly position it in the new subregion. Assuming F_{θ} is invertible, we generate these by forward (unstable) or backward (stable) iterating the PLRNN

Algorithm 1 Manifold Construction

```
Input: N_{max}: Maximum number of iterations P: periodic point \sigma: Index indicating whether the stable or unstable manifold is to be computed E^{\sigma}: Stable/unstable eigenvectors \lambda^{\sigma}: Stable/unstable eigenvalues
```

Output: $\{S_k^{\sigma}, Q_k^{\sigma}\}_{k=0}^{k_{max}}$: Sampled points and expression for the manifold in each subregion

```
1: k=0
                                                                                                                                                      \begin{array}{l} \textbf{2:} \ \ Q_k^{\sigma} \leftarrow \mathsf{GetManifold}(P, E^{\sigma}, \lambda^{\sigma}) \\ \textbf{3:} \ \ S_k^{\sigma} \leftarrow \mathsf{SamplePoints}(Q_k^{\sigma}) \\ \textbf{4:} \ \ \textbf{for} \ n = 0 : N_{max} \ \textbf{do} \\ \textbf{5:} \ \ \ S_1^{\sigma} \leftarrow \mathsf{PropagateToNextRegion}(S_k^{\sigma}) \\ \end{array} 
                                                                                                                                        > Sample points on the manifold
                                                                                                                                 ▶ Propagate points to next subregion
                 \vec{F}^{\sigma} \leftarrow \text{STEP}(S_1^{\sigma})
                                                                                                                  ▷ backward or forward step for flow vectors
  6:
                 \begin{cases} F^{\sigma}\}_{m=0}^{M}, \{S_{1}^{\sigma}\}_{m=0}^{M} \leftarrow \text{Uniqueregions}() \\ \{E^{\sigma}\}_{m=0}^{M}, \{\lambda^{\sigma}\}_{m=0}^{M} \leftarrow \text{GetEigen}() \end{cases} 
  7:
                                                                                                                                  > Sort points into unique subregions
  8:
                                                                                                                         9:
                \quad {\bf for} \ m=0:M \ {\bf do}
10:
                        \begin{array}{l} Q_{k+1}^{\sigma} \leftarrow \operatorname{GetManifold}(\{F^{\sigma}\}_{m}, \{S_{1}^{\sigma}\}_{m}, \{E^{\sigma}\}_{m}, \{\lambda^{\sigma}\}_{m}) \\ S_{k+1}^{\sigma} \leftarrow \operatorname{SamplePoints}(Q_{k+1}^{\sigma}) \end{array} \triangleright
11:

    Sample stable manifold

12:
13:
14: end for
15: return (\{S_k^{\sigma}\}_{k=0}^{k_{max}}, \{Q_k^{\sigma}\}_{k=0}^{k_{max}})
                                                                                                                                                   ▶ Manifolds in every region
```

map F_{θ} . Invertibility also ensures a stable manifold continues along just one branch into a preceding subregion (instead of multiple disconnected sets). Inverting Eq. 2 yields

$$z_{t-1} = (A + WD_{t-1})^{-1}(z_t - h).$$
 (5)

A solution to this eq. needs to be self-consistent, however, i.e. the signs of the $z_{i,t-1}$ on the l.h.s. need to be consistent with the entries $d_{i,t-1}$ in D_{t-1} on the r.h.s. To address this, we introduce a simple heuristic: 1. Perform a backward step using the current linear subregion D_t . 2. Perform a forward step using the resulting candidate solution z_{t-1}^* . 3. If $z_t = F_{\theta}(z_{t-1}^*)$, we are done. 4. If not, we take the candidate's linear subregion $D(z_{t-1}^*)$ to attempt a new inversion. 5. If this fails, we start checking the neighboring subregions by iteratively flipping bits in D_{t-1} . See Algorithm 2 in Appx.H.2 for details. Generally, this algorithm works for any ReLU based RNN. Its efficiency comes from the fact that usually the candidate will lie in the correct linear subregion, and – if this is not the case – it usually crosses only into neighboring regions.

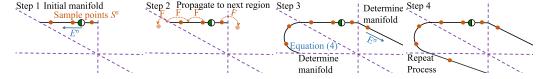


Figure 1: Illustration of the iterative procedure for computing stable manifolds with subregion boundaries indicated in purple-dashed. Step 1: The stable manifold (black) computation is initialized using the stable eigenvector (blue) of the saddle point (green), and sample points (orange) are placed along it. Step 2: These points are propagated until they enter another linear subregion, where the flow field is evaluated. Step 3: Using this, a new segment of the manifold is determined. Step 4: Repeating this process iteratively reconstructs the full global structure of the stable manifold. An example for a trained model can be found in Fig. 18.

We begin by using SCYFI (Eisenmann et al., 2023) to identify all fixed and cyclic points. In the linear subregion which harbors the saddle point of interest, the stable or unstable manifold is spanned by the n_{man} corresponding eigenvectors. To compute the global manifold, we sample seed points on this local manifold and propagate them backward or forward in time until a new linear subregion is

reached, then take one additional step to capture the local flow. If the number of stable eigenvectors in the new subregion is the same as in the original subregion and the eigenvalues are non-degenerate, we can usually represent the hyperplane or the curved manifold using the eigenvectors (Eq. 4). Otherwise we perform PCA (or kernel-PCA in the curved case), using the first n_{man} components to span the manifold. We repeat this process, resampling and propagating, until all subregions have been visited or a desired depth is achieved (see Algo. 1 and Fig. 1). In systems with highly disparate timescales, this sampling process progresses fast along fast eigendirections and thereby tends to underrepresent slow eigendirections, causing numerical issues and poor approximations. To address this, we adjust sampling density inversely to eigenvalue magnitude. This balances contributions from different directions, improves accuracy for the slower directions, and ensures a more uniform representation.

3.4 Enforcing map invertibility by regularization

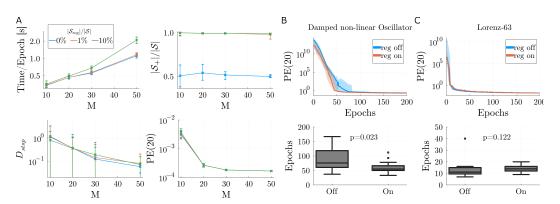


Figure 2: A) Runtime (top-left) and reconstruction quality (bottom) as a function of latent space dimensionality M for different proportions $|\mathcal{S}_{\text{reg}}|/|\mathcal{S}|$ of subregions for which invertibility was enforced by regularization ($\lambda=0.1\exp(M)$ in Eq. 6). Means across 100 different training runs on the Lorenz-63 system \pm SD are shown. Reconstruction quality was assessed through (dis-)agreement in attractor geometry (bottom-left; D_{stsp} , see Appx. J) and 20-step-ahead prediction error (bottom-right). Top-right: Relative proportion of subregions with a positive determinant of the Jacobian (det(J) > 0), \mathcal{S}_+ , as a function of latent space dimensionality M for different proportions $|\mathcal{S}_{\text{reg}}|/|\mathcal{S}|$. Medians \pm interquartile range are shown. B) Top: 20-step-ahead prediction error, PE(20), as a function of the number of training epochs when the invertibility regularization, Eq. 6, was turned off (blue) vs. on (orange), for a damped nonlinear oscillator. Bottom: Convergence to a predefined performance criterion (PE(20) \leq 0.5) was significantly faster with the regularization turned on vs. off. Median across 20 trained models, error bands = interquartile range. C) Same as B for Lorenz-63 system.

In designing our algorithm, we relied on invertibility of the RNN map F_{θ} in the sense that $\exists ! \mathbf{D}_{t-1}$ for which $F_{\theta} \left(F_{\theta}^{-1}(\mathbf{z}_t, \mathbf{D}_{t-1}) \right) = \mathbf{z}_t$. This is a reasonable assumption as the flow map $\phi(t, \mathbf{x}_0)$ for most underlying ODE systems of interest, which we attempt to approximate, is invertible due to Picard-Lindelöf (Perko, 2001). However, empirically, invertibility of F_{θ} is not always guaranteed (depending on the quality of the approximation). Thus, we enforce this condition by regularization. F_{θ} is invertible, if the determinants of the Jacobian matrices of neighboring subregions have the same sign (sign condition) (Fujisawa et al., 1972). This can be enforced (for any type of ReLU-based RNN) by adding the following regularization term to the loss:

$$\mathcal{L}_{\text{reg}} = \lambda \cdot \frac{1}{|\mathcal{S}_{\text{reg}}|} \sum_{i \in \mathcal{S}_{\text{reg}}} \max(0, -\det(\boldsymbol{J}_i)), \qquad (6)$$

computed across a small subset S_{reg} of linear subregions, where $J_i = A + WD_i$ is the Jacobian in subregion i, and λ a regularization parameter. As shown in Fig. 3.4A, strategic sampling of only 1% of linear subregions, which are traversed by actual trajectories (cf. (Brenner et al., 2024a)), hardly affects runtime (Fig. 3.4A, top) and reconstruction performance (Fig. 3.4A, bottom) while still ensuring almost full invertibility (Fig. 3.4A, right). Since invertible flows are an inherent property of many, if not most, dynamical systems of scientific interest, we would furthermore expect that this regularization does not hamper, or even improves, the reconstruction of dynamical systems from data,

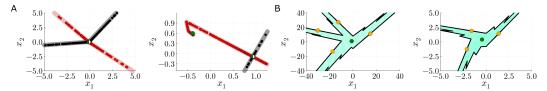


Figure 3: Model validation. A) Two examples of saddle points (half-green) with stable (gray solid lines) and unstable (red solid lines) manifolds determined by our algorithm, and points (black/ red dots, respectively) sampled by the analytical resp. backward/ forward map, showing that these all fall onto the analytically determined manifolds. B) Basins of attraction (light green, confirmed by sampling initial conditions and tracing their trajectories) of a stable fixed point (green dot) delineated by the stable manifold (black) of a 4-cycle (left) or 3-cycle (right).

in particular if the systems carry an intrinsic time reversibility (Huh et al., 2020). This is confirmed in Fig. 3.4B-C which shows that with the invertibility regularization in place, training of a PLRNN on a 10d damped nonlinear oscillator converges significantly faster to a good solution (p = 0.023, Mann-Whitney U-test) than without the regularization (Fig. 3.4B), while hardly affecting performance on other systems like the chaotic Lorenz-63 (Lorenz, 1963) system (Fig. 3.4C; p = 0.122).

4 DELINEATING BASINS OF ATTRACTION

Basic methods validation We first validate our algorithm on a simple toy example, a two-dimensional PL map for which we have analytical forms for the inverse and the fixed points (see Appx. D.1 for details). We chose parameters to produce a simple test case with only a single saddle point. Fig. 3A confirms that points produced by backward (resp. forward) iterating the map all lie on the stable (resp. unstable) manifold as determined by our algorithm. In this simple 2d example, the manifolds correspond to line segments. Changing the PL map's parameters slightly, we obtain a stable fixed point coexisting with a saddle period-4 (Fig. 3B, left) or period-3 (Fig. 3B, right) cyclic point (Gardini et al., 2009). Tracing back the stable manifold of the period-4 or period-3 saddle using our algorithm, we obtain the boundaries of the basins of attraction of the resp. fixed point (Fig. 3B). Fig. 3B further confirms this solution for the basin perfectly agrees with the 'classical' (not scalable) numerical approach of drawing initial conditions on a grid in state space, and observing their behavior in the limit $t \to \infty$.

Bistable Duffing system The Duffing (1918) system is a simple 2d nonlinear oscillator that can exhibit bistability between two spiral point attractors in certain parameter regimes (see Appx. I.1 for more details). We trained a shPLRNN (M=2,H=10; (Hess et al., 2023)) on this system using sparse teacher forcing (Mikhaeil et al., 2022), and used SCYFI (Eisenmann et al., 2023) to determine the fixed points. The basin boundary between the two spiral point attractors is the stable manifold of a saddle node in the center (Fig. 4A, and – as computed by our algorithm – agrees with the trajectory flows of the *true* system in blue.

Multistable choice paradigm Simple models of decision making in the brain assume multistability between several choice-specific attractor states, to which the system's state is driven as one or the other choice materializes (Wang, 2002). We trained an ALRNN (M=15, P=6) (Brenner et al., 2024a) to perform a simple 2-choice decision making task taken from (Gerstner et al., 2014), and, as before, use SCYFI to find fixed points and algorithm 1 to determine the stable manifold of a saddle separating the two basins of attraction corresponding to the two choices. The basin boundary consists of different linear pieces (hyperplanes) and is visualized in Fig. 4B in a 3d subspace (of the 15d system) together with the reconstructed system's fixed points (green) and some trajectories of the *true* system in blue.

Lorenz-63 attractor The Lorenz-63 model of atmospheric convection is probably the most famous example of a chaotic attractor. We reconstruct this system with a shPLRNN ($M=3,\,H=20$). Besides the chaotic attractor and two unstable spiral points, the system has a saddle node for which we compute the stable and unstable manifolds. Fig. 4C confirms that the manifolds computed by algorithm 1 for the shPLRNN (left), agree well with those determined by numerical integration of the

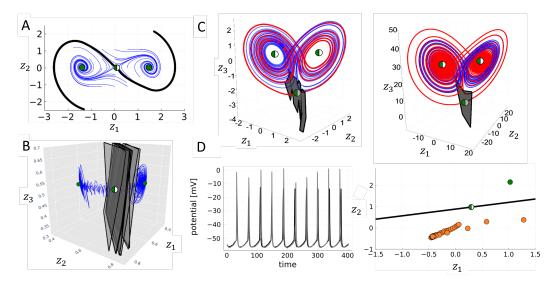


Figure 4: A) Reconstruction of the Duffing system by a shPLRNN (M=2, H=10). Trajectories drawn from the actual Duffing system in blue, identified fixed points in green, and in gray the stable manifold of the saddle in the center separating the two basins of attraction as determined by our algorithm. B) 3d subspace of the state space of an ALRNN (M=15, P=6) trained on a 2-choice decision making task, with true trajectories in blue. Two point attractors (green) were identified, with the stable manifold (black/gray) of a saddle (half-green) in the center separating the two basins. The stable manifold (basin boundary) consists of different planar pieces – note that for visualization these are projected down from a truly 15d space into a 3d subspace (accounting for some of the 'folded' appearance). C) Reconstruction for a shPLRNN (M=3, H=20) trained on the Lorenz63 system, with true trajectories in blue. The Lorenz63 system has two saddle-spirals in the center of the two lobes and a saddle at the bottom, which were correctly located by the shPLRNN. In black and red are the stable and unstable manifolds, respectively, of the saddle as identified by our algorithm (left), while on the right as computed by numerical continuation of the *original* Lorenz-63 system. The close agreement indicates the shPLRNN has correctly recovered the state space structure of the underlying system, although having been trained on trajectories from the actual attractor only. D) ALRNN (M = 25, P = 6) trained on electrophysiological recordings from a cortical neuron. Left: Time series of membrane voltage (true: gray, model-simulated: black); right: 2d projection of the ALRNN's state space with stable manifold of a saddle (black) separating the basins of attraction of a stable fixed point (green) and the 38-cycle (orange dots) corresponding to the spiking process. Note that the true stable manifold is a 24d curved object, which for visualization purposes is represented here by a locally linear approximation in the shown 2d subspace.

original Lorenz ODE system (right). This example also proves that the shPLRNN faithfully captured the geometrical structure of the state space, beyond just reconstruction of the chaotic attractor itself.

Empirical example: Single cell recordings In Figure 4D we trained an ALRNN (M=25, P=6) on membrane potential recordings from a cortical neuron (Hertäg et al., 2012). The trained ALRNN contains a 38-cycle which corresponds to the rhythmic spiking activity in the real cell. In addition it has a stable fixed point and a saddle whose stable manifold (determined by our algorithm) separates the stable cycle from the stable fixed point, as illustrated in Fig. 4D (right), where we visualized the manifold through a locally linear approximation. Although we can compute the full global 24-dimensional manifold, its inherent curvature in the 25 dimensional state space makes it impossible to visualize it directly. Many types of cortical cells exhibit this type of bistability between spiking activity and a stable equilibrium near the resting or a more depolarized potential (Izhikevich, 2007; Durstewitz and Gabriel, 2007), and this example demonstrates how our algorithm can be utilized to reveal the structure of the state space supporting this type of dynamics from real cells.

5 Homo-/Heteroclinic orbits and detection of chaos

The stable and unstable manifolds can be used to identify homoclinic and heteroclinic orbits, as defined in Sect. 3.2 (Def. 3, 4). Intersections between the stable and unstable manifolds of a saddle p lead to homoclinic points, or to heteroclinic points of two saddles $p \neq q$. The existence of such points inevitably gives rise to a complex fractal geometry (a so-scalled horseshoe structure, see Def. 5) and thus chaos (Wiggins, 1988). Finding such intersections is therefore highly illuminating for determining the dynamical behavior of a system. This is illustrated for a simple 2d PL map in Fig. 5A, where homoclinic intersections were identified by algorithm 1. For this 2d case, we can in fact analytically determine the presence of homoclinic points, as worked out in Appx. H.3, the results of which agree with algorithm 1. Fig. 5B illustrates the resulting chaotic attractor (which lies on the unstable manifold), Fig. 5C the bifurcation diagram as one varies the model's bias term as a control parameter, and Fig. 5D the system's two Lyapunov exponents across the chaotic range of h_1 (confirming the presence of 'robust chaos', as the Lyapunov exponents do not change across the chaotic regime (Banerjee et al., 1998)).

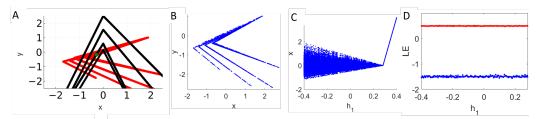


Figure 5: A) Stable (black) and unstable (red) manifolds of a saddle point (green dot) as identified by our algorithm. Note the homoclinic intersections between these two manifolds directly form the chaotic attractor. B) Structure of the chaotic attractor caused by these homoclinic intersections. C) Bifurcation diagram as a function of bias parameter h_1 . D) Lyapunov exponents across the h_1 -range for which the chaotic attractor exists.

6 CONCLUSIONS

Here we presented a novel semi-analytical algorithm for determining stable and unstable manifolds of fixed and cyclic points in ReLU-based RNNs. The algorithm analytically determines the local un-/stable manifolds within the subregion in which the cyclic point resides, and then by forward- or backward-iterating the RNN map F_{θ} collects a few support points for spanning their extensions into neighboring linear subregions. These manifolds are profoundly important for studying an RNN's dynamical repertoire, illuminating dynamical mechanisms in an underlying system reconstructed by the RNN, or the mechanisms by which an RNN solves a given task. Stable manifolds of saddle points segregate the state space into different basins of attraction, giving rise to the computationally important property of multistability (Durstewitz et al., 2000; Feudel et al., 2018). Intersections of stable and unstable manifolds, in turn, lead to homo- or heteroclinic points which produce a fractal geometry and chaos (Wiggins, 1988; Patra and Banerjee, 2018).

Limitations Limitations arise in the context of chaotic dynamics, where invariant manifolds may fold into fractal structures. While points from these manifolds may still be sampled, the analytic construction through curved/planar segments spanned by support vectors breaks down, as the intricate, self-similar geometry of fractals cannot be captured this way. Nevertheless, we still may be able to retrieve some important dynamical characteristics by determining homo- or heteroclinic intersections, as discussed in Section 5. Another limitation is that in the worst case scenario the algorithm may scale as 2^P with the number of linear subregions P. However, as shown in (Brenner et al., 2024a), the number of subregions utilized by trained PLRNNs quickly saturates, suggesting an at most polynomial scaling if one restricts attention to the domain explored by the data (see also Fig.19). In any case, we emphasize that this is the *first* algorithm for detecting un-/stable manifolds in ReLU-based RNNs.

Code for the algorithms developed here and results produced is available at [upon publication]. Large Language Models (LLMs) were used for implementation of standard routines and assistance with literature search.

REFERENCES

- K. T. Alligood, T. D. Sauer, and J. A. Yorke. *Chaos: An Introduction to Dynamical Systems*. Springer, New York, NY, 1996.
- S. Banerjee, J. A. Yorke, and C. Grebogi. Robust chaos. *Physical Review Letters*, 80, 1998.
 - N. Bertschinger and T. Natschläger. Real-time computation at the edge of chaos in recurrent neural networks. *Neural computation*, 16(7):1413–1436, 2004.
 - M. Brenner, F. Hess, J. M. Mikhaeil, L. F. Bereska, Z. Monfared, P. Kuo, and D. Durstewitz. Tractable Dendritic RNNs for Reconstructing Nonlinear Dynamical Systems. In *Proceedings of the 39th International Conference on Machine Learning*, pages 2292–2320. PMLR, June 2022. ISSN: 2640-3498.
 - M. Brenner, C. J. Hemmer, Z. Monfared, and D. Durstewitz. Almost-linear rnns yield highly interpretable symbolic codes in dynamical systems reconstruction. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 36829–36868. Curran Associates, Inc., 2024a. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/40cf27290cc2bd98a428b567ba25075c-Paper-Conference.pdf.
 - M. Brenner, F. Hess, G. Koppe, and D. Durstewitz. Integrating multimodal data for joint generative modeling of complex dynamics. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 4482–4516. PMLR, 21–27 Jul 2024b. URL https://proceedings.mlr.press/v235/brenner24a.html.
 - R. Cestnik and M. Abel. Inferring the dynamics of oscillatory systems using recurrent neural networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(6):063128, 06 2019. ISSN 1054-1500. doi: 10.1063/1.5096918. URL https://doi.org/10.1063/1.5096918.
 - S. Coombes, M. Şayli, R. Thul, R. Nicks, M. A. Porter, and Y. M. Lai. Oscillatory networks: Insights from piecewise-linear modeling. *SIAM Review*, 66(4):619–679, 2024. doi: 10.1137/22M1534365. URL https://doi.org/10.1137/22M1534365.
 - G. Duffing. Erzwungene Schwingungen bei veränderlicher Eigenfrequenz und ihre technische Bedeutung. Number 41-42. Vieweg, 1918.
 - D. Durstewitz. A state space approach for piecewiselinear recurrent neural networks for reconstructing nonlinear dynamics from neural measurements. *PLoS Computational Biology*, 13(6), 2017.
 - D. Durstewitz and T. Gabriel. Dynamical basis of irregular spiking in nmda-driven prefrontal cortex neurons. *Cerebral cortex*, 17(4):894–908, 2007.
 - D. Durstewitz, J. Seamans, and T. Sejnowski. Neurocomputational models of working memory. *Nature Neurosci*, 3:1184–1191, 2000.
 - D. Durstewitz, G. Koppe, and M. I. Thurm. Reconstructing computational system dynamics from neural data with recurrent neural networks. *Nature Reviews Neuroscience*, 24(11):693–710, 2023.
- L. Eisenmann, Z. Monfared, N. Göring, and D. Durstewitz. Bifurcations and loss jumps in RNN training. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 70511–70547. Curran Associates, Inc., 2023.
 - M. Ester, H. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.
 - U. Feudel, AN. Pisarchik, and K. Showalter. Multistability and tipping: From mathematics and physics to climate and brain-minireview and preface to the focus issue. *Chaos*, 28, 2018.

T. Fujisawa, E. Kuh, and T. Ohtsuki. A sparse matrix method for analysis of piecewise-linear resistive networks. *IEEE Transactions on Circuit Theory*, 19(6):571–584, 1972. doi: 10.1109/TCT.1972. 1083550.

- V. Ganatra and S. Banerjee. Sketching 1d stable manifolds of 2d maps without the inverse. *International Journal of Bifurcation and Chaos*, 32(08):2250111, 2022.
- L. Gardini, V. Avrutin, M. Schanz, et al. Connection between bifurcations on the Poincaré Equator and the dangerous bifurcations. 2009.
- W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- P. Glendinning. Robust chaos revisited. European Physical Journal Special Topics, 226:1721–1738,
 2017.
 - P. A. Glendinning and D. J. W. Simpson. A constructive approach to robust chaos using invariant manifolds and expanding cones. *Discrete and Continuous Dynamical Systems-Series A*, 41(7): 3367–3387, 2021.
 - M. Golub and D. Sussillo. FixedPointFinder: A Tensorflow toolbox for identifying and characterizing fixed points in recurrent neural networks. *Journal of Open Source Software*, 3(31):1003, November 2018. ISSN 2475-9066.
- A. Gu and T. Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.
 - J. Guckenheimer and P. Holmes. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, volume 42. Springer Science & Business Media, 2013.
 - B. Hasselblatt and A. Katok. Handbook of dynamical systems. Elsevier, 2002.
- L. Hertäg, J. Hass, T. Golovko, and D. Durstewitz. An approximation to the adaptive exponential integrate-and-fire neuron model allows fast and predictive fitting to physiological data. *Frontiers in computational neuroscience*, 6:62, 2012.
 - F. Hess, Z. Monfared, M. Brenner, and D. Durstewitz. Generalized teacher forcing for learning chaotic dynamics. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 13017–13049. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/hess23a.html.
 - M. W. Hirsch, S. Smale, and R. L. Devaney. *Differential equations, dynamical systems, and an introduction to chaos*. Academic press, 2013.
 - J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
 - I. Huh, E. Yang, S. J. Hwang, and J. Shin. Time-reversal symmetric ode network. *Advances in Neural Information Processing Systems*, 33:19016–19027, 2020.
 - E. M. Izhikevich. *Dynamical systems in neuroscience*. MIT press, 2007.
 - G. E. Katz and J. A. Reggia. Using Directional Fibers to Locate Fixed Points of Recurrent Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8):3636–3646, August 2018. ISSN 2162-2388. doi: 10.1109/TNNLS.2017.2733544. Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- G. Koppe, S. Guloksuz, U. Reininghaus, and D. Durstewitz. Recurrent neural networks in mobile sampling and intervention. *Schizophrenia bulletin*, 45(2):272–276, 2019a.
 - G. Koppe, H. Toutounji, P. Kirsch, S. Lis, and D. Durstewitz. Identifying nonlinear dynamical systems via generative recurrent neural networks with applications to fmri. *PLOS Computational Biology*, 15(8):1–35, 2019b.

- P. Kowalczyk. Robust chaos and border-collision bifurcations in non-invertible piecewise-linear maps. *Nonlinearity*, 18:485, 2005.
- B. Krauskopf, H. M. Osinga, and Jorge Galán-Vioque. *Numerical continuation methods for dynamical systems*, volume 2. Springer, 2007.
 - E. N. Lorenz. Deterministic nonperiodic flow. *Journal of atmospheric sciences*, 20(2):130–141, 1963.
- L. McInnes, J. Healy, and S. Astels. hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11), mar 2017. doi: 10.21105/joss.00205. URL https://doi.org/10.21105%2Fjoss.00205.
 - J. M. Mikhaeil, Z. Monfared, and D. Durstewitz. On the difficulty of learning chaotic dynamics with RNNs. In *Advances in Neural Information Processing Systems*, 2022.
 - Z. Monfared and D. Durstewitz. Existence of n-cycles and border-collision bifurcations in piecewise-linear continuous maps with applications to recurrent neural networks. *Nonlinear Dynamics*, 101 (2):1037–1052, 2020a.
 - Z. Monfared and D. Durstewitz. Transformation of ReLU-based recurrent neural networks from discrete-time to continuous-time. In *International Conference on Machine Learning*, pages 6999– 7009. PMLR, 2020b. ISSN: 2640-3498.
 - J. Nassar, S. Linderman, M. Bugallo, and I. M. Park. Tree-structured recurrent switching linear dynamical systems for multi-scale modeling. In *International Conference on Learning Representations*, 2019.
 - A. Orvieto, S. L. Smith, A. Gu, A. Fernando, C. Gulcehre, R. Pascanu, and S. De. Resurrecting recurrent neural networks for long sequences. In *International Conference on Machine Learning*, pages 26670–26698. PMLR, 2023.
 - H. M. Osinga. Computing global invariant manifolds: Techniques and applications. In *Proceedings* of the International Congress of Mathematicians, volume 4, pages 1101–1123, 2014.
 - S. H. Park, B. Kim, C. M. Kang, C. C. Chung, and J. W. Choi. Sequence-to-sequence prediction of vehicle trajectory via lstm encoder-decoder architecture. In 2018 IEEE Intelligent Vehicles Symposium (IV), pages 1672–1678, 2018. doi: 10.1109/IVS.2018.8500658.
 - D. Patel and E. Ott. Using machine learning to anticipate tipping points and extrapolate to post-tipping dynamics of non-stationary dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 33(2), 2023.
 - M. Patra and S. Banerjee. Robust chaos in 3-d piecewise linear maps. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28, 2018.
 - B. Peng, E. Alcaide, Q. Anthony, A. Albalak, S. Arcadinho, S. Biderman, H. Cao, X. Cheng, M. Chung, M. Grella, et al. Rwkv: Reinventing rnns for the transformer era. *arXiv* preprint *arXiv*:2305.13048, 2023.
- U. Pereira-Obilinovic, J. Aljadeff, and N. Brunel. Forgetting leads to chaos in attractor networks.
 Phys. Rev. X, 13:011009, Jan 2023. doi: 10.1103/PhysRevX.13.011009. URL https://link.aps.org/doi/10.1103/PhysRevX.13.011009.
 - L. Perko. Differential Equations and Dynamical Systems. Springer New York, NY, 2001.
- J. A. Platt, S. G. Penny, T. A. Smith, T.n Chen, and H. D.I. Abarbanel. Constraining chaos: Enforcing dynamical invariants in the training of recurrent neural networks. *arXiv preprint arXiv:2304.12865*, 2023.
- M. Rabinovich, R. Huerta, and G. Laurent. Transient dynamics for neural processing. *Science*, 321 (5885):48–50, 2008.
 - I. Roy, M. Patra, and S. Banerjee. Shilnikov-type dynamics in three-dimensional piecewise smooth maps. *Chaos Soliton. Fract.*, 133, 2020.

- T. K. Rusch and S. Mishra. Coupled oscillatory recurrent neural network (cornn): An accurate and (gradient) stable architecture for learning long time dependencies. *arXiv preprint arXiv:2010.00951*, 2020.
- T. K. Rusch, S. Mishra, N. B. Erichson, and M. W. Mahoney. Long expressive memory for sequence modeling. In *International Conference on Learning Representations*.
- D. Schmidt, G. Koppe, Z. Monfared, M. Beutelspacher, and D. Durstewitz. Identifying nonlinear dynamical systems with multiple time scales and long-range dependencies. *In proceedings of the Ninth International Conference on Learning Representations, ICLR*, 2021.
- H. T. Siegelmann and E. D Sontag. On the computational power of neural nets. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 440–449, 1992.
- D. J. W. Simpson. Border-collision bifurcations in r^n . SIAM Review, 58(2), 2016.
- D. J. W. Simpson. Detecting invariant expanding cones for generating word sets to identify chaos in piecewise-linear maps. *Journal of Difference Equations and Applications*, pages 1775–1789, 2022. URL DOI:10.1080/10236198.2022.2070009.
- D. J. W. Simpson. How to compute multi-dimensional stable and unstable manifolds of piecewise-linear maps. *arXiv* preprint arXiv:2310.09941, 2023.
- X. Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5): 955–968, 2002.
- X. Wang. Decision making in recurrent neuronal circuits. *Neuron*, 60(2):215–234, 2008.
- S. Wiggins. Global Bifurcation and Chaos. Springer-Verlag, New York, 1988.

APPENDIX:

A THE PLRNN AS 2D PIECEWISE-LINEAR MAP

The PLRNN is a piecewise linear map. As a simple example let us define a 2d map F on \mathbb{R}^2 with one boundary, for which we have simple ground truth systems for, by:

$$F(X) = \begin{cases} \mathbf{A}_{l} \cdot X + \mathbf{B}, & \text{if } x \leq 0, \\ \mathbf{A}_{r} \cdot X + \mathbf{B}, & \text{if } x \geq 0, \end{cases}$$
 (7)

where $X = \begin{pmatrix} x \\ y \end{pmatrix}$ is the position vector. The transformation matrices and vector are specified as follows:

$$m{A_l} = egin{pmatrix} au_l & c \ -\delta_l & d \end{pmatrix}, \quad m{A_r} = egin{pmatrix} au_r & c \ -\delta_r & d \end{pmatrix}, \quad m{B} = egin{pmatrix} h_1 \ h_2 \end{pmatrix}$$

This map can be reformulated as a PLRNN (Eq. 2) by defining the parameters A, W, h as:

$$\boldsymbol{A} = \begin{pmatrix} \tau_l & c \\ -\delta_l & d \end{pmatrix}, \quad \boldsymbol{W} = \begin{pmatrix} \tau_r - \tau_l & 0 \\ -\delta_r + \delta_l & 0 \end{pmatrix}, \quad \boldsymbol{h} = \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}$$

B GENERAL 2D PL MAPS

A 2-unit PLRNN is a 2d PL dynamical system whose phase space is generally split into 4 sub-regions by 4 borders; see Appx. F for more details. For our exposition it suffices to focus on just one border, however. Hence studying general 2d PL maps, we can investigate dynamics of 2d PLRNNs locally near one border. Also, as shown in Appx. F, in some cases, there are only two different sub-regions divided by one border for 2d PLRNNs. This means in such situations, a 2d PLRNN exactly has the form of a 2d PL map. Therefore, here we examine dynamics of a 2d PL map.

A 2d PL map T is a continuous map on \mathbb{R}^2 which is affine on each side of the line $\Sigma := \{(x,y)^\mathsf{T} \in \mathbb{R}^2 : x = 0\}$:

$$\begin{pmatrix} x^{(k+1)} \\ y^{(k+1)} \end{pmatrix} = T(x^{(k)}, y^{(k)})
= \begin{cases}
T_{\mathcal{L}}(x^{(k)}, y^{(k)}) = \underbrace{\begin{pmatrix} a_l & c \\ b_l & d \end{pmatrix}}_{\mathbf{A}_{\mathcal{L}}} \begin{pmatrix} x^{(k)} \\ y^{(k)} \end{pmatrix} + \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}; \quad x^{(k)} \leq 0
= \begin{cases}
T_{\mathcal{R}}(x^{(k)}, y^{(k)}) = \underbrace{\begin{pmatrix} a_r & c \\ b_r & d \end{pmatrix}}_{\mathbf{A}_{\mathcal{R}}} \begin{pmatrix} x^{(k)} \\ y^{(k)} \end{pmatrix} + \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}; \quad x^{(k)} \geq 0
\end{cases}$$
(8)

where $a_l, a_r, b_l, b_r, c, d, h_1, h_2 \in \mathbb{R}$. The phase space of the map Eq. 8 is divided into two different sub-regions $\mathcal{L} := \{(x,y)^\mathsf{T} \in \mathbb{R}^2 : x \leq 0\}$ and $\mathcal{R} := \{(x,y)^\mathsf{T} \in \mathbb{R}^2 : x \geq 0\}$ by the borderline Σ . The map Eq. 8 is continuous across the border, but its Jacobian matrix is discontinuous across Σ .

Under some conditions on the system parameters, there exists a well-defined and invertible coordinate change that transforms the general 2d PL map Eq. 8 into a 2d PL normal form map (Glendinning and Simpson, 2021; Simpson, 2022). For instance when T is generic in the sense that $T(\Sigma)$ intersects Σ at a unique point which is not a fixed point of T, or equivalently $c \neq 0$ and (1-d) h_1+c $h_2 \neq 0$; see (Glendinning and Simpson, 2021) for more details. However, it is not always possible to apply such a transformation. Therefore, to study 2d PLRNNs in a comprehensive way, we need to investigate the map Eq. 8 in a general case, without any restrictions on the parameters, where here we focus particularly on the case c=0 and/or (1-d) h_1+c $h_2=0$.

C STABLE MANIFOLDS WITH COMPLEX OR DEGENERATE EIGENVALUES

2d spiral For complex-conjugate eigenvalues in 2d, the associated eigenvectors span a two-dimensional invariant plane in which the motion forms a spiral (or rotation with expansion/ contraction) around the equilibrium. That is, within each linear subregion, the solution on an un-/stable manifold evolves in a straight or spiraling fashion. The spiral structure can be derived from the eigenvalues as follows. For simplicity we shift the coordinate system such that h=0. Suppose $A + WD_t$ has a *complex* eigenvalue $\lambda = r e^{i\theta}$ (where $|\lambda| = r < 1$) with corresponding eigenvector v = A + ib, $A, b \in \mathbb{R}^d$. A solution in the direction of this eigenvector evolves as

$$\boldsymbol{z}_t = (\boldsymbol{A} + \boldsymbol{W} \boldsymbol{D}_t)^t \boldsymbol{z}_0.$$

If z_0 lies in the span of A, b, then

$$\mathbf{z}_t = \operatorname{Re}(\lambda^t v) = r^t \Big[\cos(t \theta) \mathbf{A} - \sin(t \theta) \mathbf{b} \Big].$$
 (9)

General case

Consider a linear system of the form $z_{t+1} = Lz_t$.

Case (1): If the geometric multiplicity equals the algebraic multiplicity, then L is diagonalizable and admits a basis of linearly independent eigenvectors v_1, \dots, v_n . Thus, the orbits of the system can be expanded in terms of eigenvectors as

$$\mathbf{z}_t = c_1 \lambda_1^t \mathbf{v}_1 + c_2 \lambda_2^t \mathbf{v}_2 + \dots + c_n \lambda_n^t \mathbf{v}_n. \tag{10}$$

Case (2): If the geometric multiplicity of λ_i is less than the algebraic multiplicity, then the eigenvalue is said to be defective and \boldsymbol{L} admits a basis of generalized eigenvectors. In this case, \boldsymbol{L} has a Jordan block $\boldsymbol{J}_m(\lambda_i)$ of size m>1

$$J_m(\lambda_i) = \begin{pmatrix} \lambda_i & 1 & 0 & \dots & 0 \\ 0 & \lambda_i & 1 & \dots & 0 \\ 0 & 0 & \lambda_i & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & 1 \\ 0 & 0 & 0 & \dots & \lambda_i \end{pmatrix},$$

with the t-th power

$$J_{m}^{t}(\lambda_{i}) = \begin{pmatrix} \lambda_{i}^{t} & \binom{t}{1} \lambda_{i}^{t-1} & \binom{t}{2} \lambda_{i}^{t-2} & \cdots & \binom{t}{m-1} \lambda_{i}^{t-m+1} \\ 0 & \lambda_{i}^{t} & \binom{t}{1} \lambda_{i}^{t-1} & \cdots & \binom{t}{m-2} \lambda_{i}^{t-m+2} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_{i}^{t} & \binom{t}{1} \lambda_{i}^{t-1} \\ 0 & 0 & \cdots & 0 & \lambda_{i}^{t} \end{pmatrix}.$$

Moreover, a generalized eigenvector $w_m \neq 0$ of degree m, corresponding to the defective eigenvalue λ_i , satisfies

$$(\boldsymbol{L} - \lambda \boldsymbol{I})^m \boldsymbol{w}_m = 0, \quad \text{but } (\boldsymbol{L} - \lambda \boldsymbol{I})^{m-1} \boldsymbol{w}_m \neq 0,$$
 (11)

and L has m linearly independent generalized eigenvectors associated with λ_i . In fact, we can construct a chain of generalized eigenvectors $\{w_1, \dots, w_m\}$ such that

$$(L - \lambda I)w_m = w_{m-1}, (L - \lambda I)w_{m-1} = w_{m-2}, \dots, (L - \lambda I)w_2 = w_1,$$
 (12)

where w_1 is a regular eigenvector. Given a chain of length m, the orbit contribution corresponding to the Jordan block $J_m(\lambda_i)$ is given by

$$\mathbf{z}_{t}^{defective} = \lambda_{i}^{t} \left(c_{1} \mathbf{w}_{1} + c_{2} t \mathbf{w}_{2} + c_{3} \frac{t^{2}}{2!} \mathbf{w}_{3} + \dots + c_{m} \frac{t^{m-1}}{(m-1)!} \mathbf{w}_{m} \right).$$
 (13)

The full orbit of the system is a linear combination of contributions from the non-defective eigenvalues λ_i , $1 \le j \ne i \le n$, and the defective eigenvalue λ_i as

$$z_{t} = \sum_{\substack{j=1\\j\neq i}}^{n} c_{j} \lambda_{j}^{t} v_{j} + \lambda_{i}^{t} \sum_{r=1}^{m} c_{r} \frac{t^{r-1}}{(r-1)!} w_{r}.$$
(14)

This formula describes exponential evolution for all non-defective eigenvalues and polynomially modified exponentials for the defective eigenvalue λ_i .

For an affine system of the form $z_{t+1} = Lz_t + h$, Eq. 14 consists of a homogeneous part (determined by the eigenstructure of L) and a particular part due to the constant bias term h. In the presence of a defective eigenvalue λ_i with Jordan block of size m, the full orbit is given by

$$z_{t} = \sum_{\substack{j=1\\j\neq i}}^{n} c_{j} \lambda_{j}^{t} v_{j} + \lambda_{i}^{t} \sum_{r=1}^{m} c_{r} \frac{t^{r-1}}{(r-1)!} w_{r} + \sum_{k=0}^{t-1} L^{k} h.$$
 (15)

The final term $\sum_{k=0}^{t-1} \mathbf{L}^k \mathbf{h}$ accounts for the cumulative effect of the bias, modifying the orbit away from purely exponential or polynomial-exponential behavior.

D CHAOS

 Definition 5 (**Horseshoe structure**). A horseshoe structure is generated when there is a homoclinic/heteroclinic intersection between stable and unstable manifolds of a saddle fixed point and therefore an infinite number of similar intersections. The occurrence of a homoclinic/heteroclinic intersection implies the existence of a chaotic orbit (Wiggins, 1988).

Definition 6 (Robust chaos). A chaotic attractor is called robust if, for its parameter values, there are no periodic windows and coexisting attractors in some neighborhood of the parameter space such that the chaotic attractor is unique in that open subset (Banerjee et al., 1998; Patra and Banerjee, 2018). Accordingly, small perturbations in either the parameter or phase space will not destroy a robust chaotic attractor.

Many practical applications of neural networks depend on such a robust chaotic mode for reliable operation (Patra and Banerjee, 2018). Banerjee et al. (Banerjee et al., 1998) were the first to introduce the idea of robust chaos in the context of 2d PL normal form maps. They discussed the occurrence of robust chaos in 1d and 2d PWS systems, and derived the existence and stability conditions of robust chaos in 2d PL normal form maps. In (Kowalczyk, 2005), robust chaos and border collision bifurcations (BCBs) were studied for non-invertible 2d PL normal form maps. Later, in (Glendinning, 2017) the concept of robust chaos was revisited to provide a new set of tools for analyzing it. Afterwards, (Glendinning and Simpson, 2021) illustrated a constructive approach to examine robust chaos, in the original parameter regime of (Banerjee et al., 1998), based on invariant manifolds and expanding cones. Very recently, Simpson (Simpson, 2022) detected invariant expanding cones and presented a general method to identify robust chaotic attractors in 2d border-collision normal form maps. Here, we investigate the existence of chaos for the general PL system Eq. 8 where T is an invertible map. For this purpose, we first discuss its invertibility.

D.1 INVERTIBILITY OF THE MAP EQ. 8

The map Eq. 8 is invertible iff $D_{\mathcal{L}}D_{\mathcal{R}} > 0$ and

$$\boldsymbol{x}^{(k)} = T^{-1}(\boldsymbol{x}^{(k+1)})$$

$$= \begin{cases} T_{\mathcal{L}}^{-1} = \boldsymbol{A}_{\mathcal{L}}^{-1}(\boldsymbol{x}^{(k+1)} - \boldsymbol{h}); & \frac{\varphi^{\mathsf{T}}(\boldsymbol{x}^{(k+1)} - \boldsymbol{h})}{D_{\mathcal{L}}} \leq 0 \\ T_{\mathcal{R}}^{-1} = \boldsymbol{A}_{\mathcal{R}}^{-1}(\boldsymbol{x}^{(k+1)} - \boldsymbol{h}); & \frac{\varphi^{\mathsf{T}}(\boldsymbol{x}^{(k+1)} - \boldsymbol{h})}{D_{\mathcal{R}}} \geq 0 \end{cases} , \tag{16}$$

where $D_{\mathcal{L}/\mathcal{R}}$ is the determinant of $\mathbf{A}_{\mathcal{L}/\mathcal{R}}$, and $\varphi^{\mathsf{T}} = \mathbf{e}_1^{\mathsf{T}} adj(\mathbf{A}_{\mathcal{L}}) = \mathbf{e}_1^{\mathsf{T}} adj(\mathbf{A}_{\mathcal{R}}) = (d - c)$, (Simpson, 2016). Therefore

$$\begin{pmatrix} x^{(k)} \\ y^{(k)} \end{pmatrix} = T^{-1}(x^{(k+1)}, y^{(k+1)})$$

$$= \begin{cases} \frac{1}{D_{\mathcal{L}}} \begin{bmatrix} \begin{pmatrix} d & -c \\ -b_{l} & a_{l} \end{pmatrix} \begin{pmatrix} x^{(k+1)} \\ y^{(k+1)} \end{pmatrix} + \begin{pmatrix} c h_{2} - d h_{1} \\ b_{l} h_{1} - a_{l} h_{2} \end{pmatrix} \end{bmatrix}; & \Phi \leq 0 \\ \frac{1}{D_{\mathcal{R}}} \begin{bmatrix} \begin{pmatrix} d & -c \\ -b_{r} & a_{r} \end{pmatrix} \begin{pmatrix} x^{(k+1)} \\ y^{(k+1)} \end{pmatrix} + \begin{pmatrix} c h_{2} - d h_{1} \\ b_{r} h_{1} - a_{r} h_{2} \end{pmatrix} \end{bmatrix}; & \Phi \geq 0 \end{cases}$$
(17)

where

$$\Phi := \frac{\varphi^{\mathsf{T}}(\boldsymbol{x}^{(k+1)} - \boldsymbol{h})}{D_{\mathcal{L}/\mathcal{R}}}
= \frac{d}{D_{\mathcal{L}/\mathcal{R}}} (x^{(k+1)} - h_1) - \frac{c}{D_{\mathcal{L}/\mathcal{R}}} (y^{(k+1)} - h_2).$$
(18)

Denoting the border of T by $\Sigma = \{(x,y) \in \mathbb{R}^2 : x = 0\}$, the switching border of T^{-1} is the line

$$\ell^{\Sigma} =: T(\Sigma) = \{ (x, y) \in \mathbb{R}^2 : cy = dx - dh_1 + ch_2 \}, \tag{19}$$

along which T^{-1} is continuous.

Remark D.1. For $c \neq 0$, the switching border of T^{-1} is the line

$$\ell^{\Sigma} = \{ (x, y) \in \mathbb{R}^2 : y = \frac{d}{c} x - \frac{d h_1}{c} + h_2 \}, \tag{20}$$

and $T(\Sigma)$ intersects Σ at a unique point. While for c=0 and $d\neq 0$ we have

$$\ell^{\Sigma} = \{ (x, y) \in \mathbb{R}^2 : x = h_1 \}, \tag{21}$$

which means $T(\Sigma)$ and Σ never intersect at any point or are coinciding with each other.

When $D_{\mathcal{L}}D_{\mathcal{R}} \leq 0$, then either there is no inverse for T or there are two inverses. If $D_{\mathcal{L}}D_{\mathcal{R}} > 0$, then there is a unique inverse T^{-1} for T. In this case, for $c \, d \neq 0$, the nature of T^{-1} depends on the sign of $D_{\mathcal{L}}$, $D_{\mathcal{R}}$, c, d and $y^{(k+1)} - \frac{d(x^{(k+1)} - h_1)}{c} - h_2$; see Table 1. As shown in Table 2 (left), for c = 0 the nature of the inverse map can be determined based on the sign of $D_{\mathcal{L}}$, $D_{\mathcal{R}}$, d and $x^{(k+1)} - h_1$. Note that, when c = 0, for $a_{l/r} = 1$ or d = 1 the matrix $A_{\mathcal{L}/\mathcal{R}}$ has an eigenvalue 1. Similarly, for d = 0, it depends on the sign of $D_{\mathcal{L}}$, $D_{\mathcal{R}}$, c and $y^{(k+1)} - h_2$; Table 2 (right).

Next we examine the theoretical conditions for the existence of chaotic orbits. In this regard, the fold structure of the stable and unstable manifolds will be discussed in detail. In particular it will be shown that, unlike 2d PL normal form maps, for general 2d PL maps the unstable manifold does not fold along the *x*-axis.

D.2 NATURE OF THE STABLE AND UNSTABLE MANIFOLDS

Let T be invertible and $\mathcal{O}_{\mathcal{L}/\mathcal{R}}$ a saddle fixed point with stable and unstable eigenvalues λ_s and λ_u . The stable and unstable subspaces E^s and E^u of $\mathcal{O}_{\mathcal{L}/\mathcal{R}}$ are lines crossing $\mathcal{O}_{\mathcal{L}/\mathcal{R}}$, generated by the stable and unstable eigenvectors \mathbf{v}_s and \mathbf{v}_u associated with λ_s and λ_u . Since T is a PL map, the local stable and unstable manifolds of $\mathcal{O}_{\mathcal{L}/\mathcal{R}}$ coincide with the stable and unstable subspaces E^s and E^u , respectively. The global stable and unstable manifolds $W^s(\mathcal{O}_{\mathcal{L}/\mathcal{R}})$ and $W^u(\mathcal{O}_{\mathcal{L}/\mathcal{R}})$ have a complex PL structure due to the PL nature of T. In fact, as mentioned before, any point on the y-axis maps to the line ℓ^Σ as defined in Eq. 19. Since the linear map changes across the y-axis, the

$D_{\mathcal{L}}$	$D_{\mathcal{R}}$	c	d	$y^{(k+1)} - \frac{d(x^{(k+1)} - h_1)}{c} - h_2$	Map
+ve	+ve	+ve	+ve	-ve	$T_{\mathcal{R}}^{-1}$
+ve	+ve	+ve	-ve	-ve	$T_{\mathcal{R}}^{-1}$
+ve	+ve	-ve	-ve	+ve	$T_{\mathcal{R}}^{-1}$
+ve	+ve	-ve	+ve	+ve	$T_{\mathcal{R}}^{-1}$
-ve	-ve	+ve	+ve	+ve	$T_{\mathcal{R}}^{-1}$
-ve	-ve	+ve	-ve	+ve	$T_{\mathcal{R}}^{-1}$
-ve	-ve	-ve	-ve	-ve	$T_{\mathcal{R}}^{-1}$
-ve	-ve	-ve	+ve	-ve	$T_{\mathcal{R}}^{-1}$
+ve	+ve	+ve	+ve	+ve	$T_{\mathcal{L}}^{-1}$
+ve	+ve	+ve	-ve	+ve	$T_{\mathcal{L}}^{-1}$
+ve	+ve	-ve	-ve	-ve	$T_{\mathcal{L}}^{-1}$
+ve	+ve	-ve	+ve	-ve	$T_{\mathcal{L}}^{-1}$
-ve	-ve	+ve	+ve	-ve	$T_{\mathcal{L}}^{-1}$
-ve	-ve	+ve	-ve	-ve	$T_{\mathcal{L}}^{-1}$
-ve	-ve	-ve	-ve	+ve	$T_{\mathcal{L}}^{-1}$
-ve	-ve	-ve	+ve	+ve	$T_{\mathcal{L}}^{-1}$
				·	

Table 1: Signs of $D_{\mathcal{L}}$, $D_{\mathcal{R}}$, c, d and $y^{(k+1)} - \frac{d(x^{(k+1)} - h_1)}{c} - h_2$ decide which map to apply.

$\overline{D_{\mathcal{L}}}$	$D_{\mathcal{R}}$	d	$x^{(k+1)} - h_1$	Map	$\overline{D_{\mathcal{L}}}$;	$D_{\mathcal{R}}$	c	$y^{(k+1)} - h_2$	Map
+ve	+ve	+ve	+ve	$T_{\mathcal{R}}^{-1}$	+v	e	+ve	+ve	-ve	$T_{\mathcal{R}}^{-1}$
+ve	+ve	-ve	-ve	$T_{\mathcal{R}}^{-1}$	+v	е	+ve	-ve	+ve	$T_{\mathcal{R}}^{-1}$
-ve	-ve	+ve	-ve	$T_{\mathcal{R}}^{-1}$	$-\mathbf{v}$	е	-ve	+ve	+ve	$T_{\mathcal{R}}^{-1}$
-ve	-ve	-ve	+ve	$T_{\mathcal{R}}^{-1}$	$-\mathbf{v}$	е	-ve	-ve	-ve	$T_{\mathcal{R}}^{-1}$
+ve	+ve	+ve	-ve	$T_{\mathcal{L}}^{-1}$	+v	е	+ve	+ve	+ve	$T_{\mathcal{L}}^{-1}$
+ve	+ve	-ve	+ve	$T_{\mathcal{L}}^{-1}$	+v	е	+ve	-ve	-ve	$T_{\mathcal{L}}^{-1}$
-ve	-ve	-ve	-ve	$T_{\mathcal{L}}^{-1}$	$-\mathbf{v}$	e _	-ve	-ve	+ve	$T_{\mathcal{L}}^{-1}$
-ve	-ve	+ve	+ve	$T_{\mathcal{L}}^{-1}$	-v	e _	-ve	+ve	-ve	$T_{\mathcal{L}}^{-1}$

Table 2: Left (c=0): signs of $D_{\mathcal{L}}$, $D_{\mathcal{R}}$, d and $x^{(k+1)}-h_1$, and right (d=0): signs of $D_{\mathcal{L}}$, $D_{\mathcal{R}}$, c and $y^{(k+1)}-h_2$ to determine the nature of T^{-1} .

unstable manifold must have different slopes on the two sides of ℓ^{Σ} . That is, the unstable manifold folds along the line ℓ^{Σ} (Fig.s 6-7) and all images of the fold points will also be fold points. Likewise, since under the action of T^{-1} the line ℓ^{Σ} maps to the y-axis, the stable manifold folds along the y-axis; see Fig. 7. Moreover, all the preimages of fold points are fold points. Because of the folds of $W^s(\mathcal{O}_{\mathcal{L}/\mathcal{R}})$ and $W^u(\mathcal{O}_{\mathcal{L}/\mathcal{R}})$, strictly speaking, they are not manifolds. However, they are called manifolds just as a matter of convention.

Consider system Eq. 8 with parameters

$$a_r = 1.5, b_r = -0.75, a_l = -1.77, b_l = -0.9, c = 0.6,$$

 $d = 0.15, h_2 = -0.4, h_1 = -0.7.$ (22)

For these values, $D_{\mathcal{L}}D_{\mathcal{R}}=0.18529>0$ which implies the map Eq. 8 is invertible and $\mathcal{O}_{\mathcal{L}}^*=(-0.28848,-0.16514)^{\mathsf{T}}$ is a saddle fixed point in the left sub-region. The matrix $A_{\mathcal{L}}$ has two real eigenvalues $\lambda_1=-1.4277$ and $\lambda_2=-0.1922$. Moreover, ℓ_u^* hits the border x=0 at $P_0=(0,-0.000582)^{\mathsf{T}}$, and one obtains $P_1=(-0.70035,-0.40009)^{\mathsf{T}}\in\mathcal{L}$, and $P_2=(0.29957,0.1703)^{\mathsf{T}}\in\mathcal{R}$. Since $\mathsf{L}(x_1,y_1)\cdot\mathsf{L}(x_2,y_2)=0.044708>0$, we need to determine the first and second fold points of the stable manifold. Calculating $\tilde{P}_0=(0,0.59343)^{\mathsf{T}}$

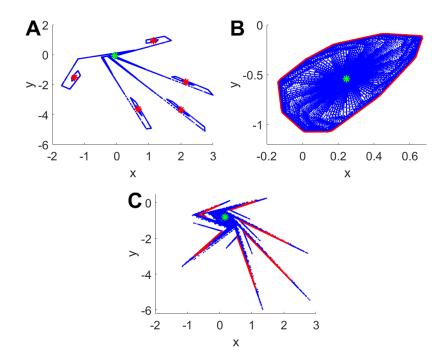


Figure 6: A) Unstable manifold (blue) corresponding to the saddle fixed point $\mathcal{O}_{\mathcal{L}}^* = (-0.0601, -0.0509)^{\mathsf{T}}$ (green star) and a stable 5-cycle (red stars) lying on it. Parameter settings: $a_r = 1.5, \ b_r = -1.58, \ a_l = -1.67, \ b_l = -0.9, \ c = 0.6, \ d = 0.1, \ h_2 = -0.1, \ h_1 = -0.13.$ B) Unstable manifold corresponding to an unstable fixed point (green star) and a stable quasi-periodic orbit (red) on it. Parameter settings: $a_l = 0.31, \ b_l = -0.6, \ a_r = -1.1, \ b_r = -1.54, \ , c = 0.9, \ d = 0.3, \ h_1 = 1, \ h_2 = 0.$ C) Unstable manifold of the unstable fixed point $\mathcal{O}_{\mathcal{R}}^* = (0.16978, -0.80815)^{\mathsf{T}}$ (green star) and a 2-band chaotic attractor (red) on it. Parameter settings: $a_r = 1.5, \ b_r = -1.69, \ a_l = -1.77, \ b_l = -0.9, \ c = 0.6, \ d = 0.15, \ h_2 = -0.4, \ h_1 = 0.4.$

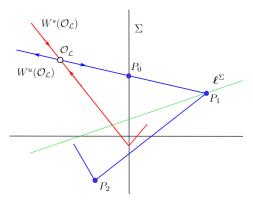


Figure 7: Folding structure of the stable and unstable manifolds of the map Eq. 8 for $c \neq 0$.

and $\frac{\varphi^{\mathsf{T}}(\tilde{P}_0-h)}{D_{\mathcal{L}}}=-1.7889<0$, we get $\tilde{P}_{-1}=(-1.7889,-4.1106)^{\mathsf{T}}$. Hence, we observe $\tilde{\mathsf{L}}(x_1,y_1)\cdot\tilde{\mathsf{L}}(x_2,y_2)=-1.0269<0$, which implies the unstable manifold intersects the stable manifold at $\tilde{P}_{hom}=(-0.28848,-0.16514)^{\mathsf{T}}$. Since $\tilde{x}_{home}\tilde{x}_{-1}=0.51606>0$ this is a homoclinic intersection and thus there must be a chaotic orbit, see Fig. 8. Moreover, fixing all parameter values while varying h_1 from negative to positive, a border collision bifurcation happens at $h_1\approx 0.282$ (Fig. 8B). For the parameter setting Eq. 22 we have $(1-d)h_1+ch_2=0.85h_1-0.24$. Therefore, at the bifurcation value $h_1\approx 0.282$ the term $(1-d)h_1+ch_2$ becomes zero. Hence, this is a bifurcation value that cannot be obtained by considering the 2d PL normal form map.

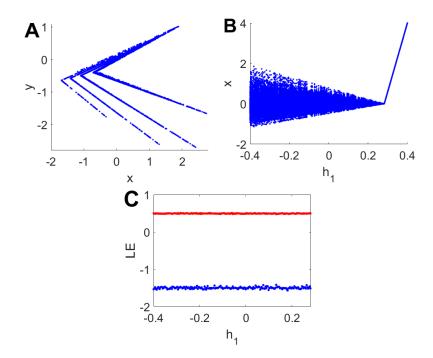


Figure 8: A) Phase portrait of a chaotic attractor due to a homoclinic interaction, B) bifurcation diagram, C) Lyapunov exponents (LEs). Parameter settings: $a_r = 1.5, \ b_r = -0.75, \ a_l = -1.77, \ b_l = -0.9, \ c = 0.6, \ d = 0.15, \ h_2 = -0.4, \ h_1 = -0.7.$

In Fig. 8C, the Lyapunov exponents are plotted while parameter h_1 is varied. The largest (maximum) Lyapunov exponent (red) is positive throughout the displayed range, whereas the second one (blue) remains negative. This implies the orbit is chaotic throughout a larger range of parameter values, and hence the system exhibits *robust chaos* not sensitive to smaller changes in parameter values. With parameter values close to those in Eq. 22 such that c=0, one still obtains chaotic attractors (Fig. 9). Recall that for c=0 the general 2d PL map Eq. 8 cannot be transformed into a 2d PL normal form map. This means none of the bifurcation points in Fig. 9 can be examined in the 2d PL normal form map.

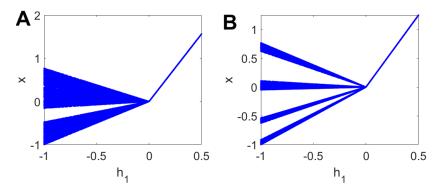


Figure 9: BCB diagrams for parameter values: A) $a_r = 0.68$, $b_r = -0.75$, $a_l = -1.77$, $b_l = -0.9$, c = 0, d = 0.15, $h_2 = -0.4$, $h_1 = -0.7$; and B) $a_r = 0.6$, $b_r = -0.75$, $a_l = -1.77$, $b_l = -0.9$, c = 0, d = 0.15, $h_2 = -0.4$, $h_1 = -0.7$.

D.2.1 HOMOCLINIC INTERSECTIONS FOR FURTHER ITERATIONS OF FOLD POINTS

In order to determine homoclinic intersections, it may be necessary to check further iterations of fold points for which we would need a recursive procedure. This will involve calculating the n-th power of a 2×2 matrix, and we therefore first prove the following Proposition:

Proposition D.2. Let $M = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$ have two distinct eigenvalues

$$\lambda_{1,2} = \frac{a+d}{2} \mp \frac{\sqrt{(a-d)^2 + 4bc}}{2} = \frac{\Gamma}{2} \mp \frac{\sqrt{\Gamma^2 - 4D}}{2},$$
 (23)

where Γ and \mathcal{D} are the trace and determinant of M. Then, for every $n \in \mathbb{N}$

$$\mathbf{M}^{n} = \begin{pmatrix} A_{n+1} - dA_{n} & cA_{n} \\ bA_{n} & dA_{n} - \mathcal{D}A_{n-1} \end{pmatrix}$$
 (24)

where

$$A_n = \frac{\lambda_1^n - \lambda_2^n}{\lambda_1 - \lambda_2}. (25)$$

Proof. In our case M is diagonalizable, so for $b \neq 0$

$$M = V \Lambda V^{-1}$$

$$= \frac{1}{\lambda_1 - \lambda_2} \begin{pmatrix} \frac{\lambda_1 - d}{b} & \frac{\lambda_2 - d}{b} \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \begin{pmatrix} b & d - \lambda_2 \\ -b & \lambda_1 - d \end{pmatrix}. \tag{26}$$

Therefore

$$M^{n} = V \Lambda^{n} V^{-1}$$

$$= \frac{1}{\lambda_{1} - \lambda_{2}} \begin{pmatrix} \frac{\lambda_{1} - d}{b} & \frac{\lambda_{2} - d}{b} \\ 1 & 1 \end{pmatrix} \begin{pmatrix} \lambda_{1}^{n} & 0 \\ 0 & \lambda_{2}^{n} \end{pmatrix} \begin{pmatrix} b & d - \lambda_{2} \\ -b & \lambda_{1} - d \end{pmatrix}$$

$$= \frac{1}{\lambda_{1} - \lambda_{2}} \times$$

$$\begin{pmatrix} \lambda_{2}^{n} (d - \lambda_{2}) - \lambda_{1}^{n} (d - \lambda_{1}) & -\frac{(d - \lambda_{1})(d - \lambda_{2})}{b} (\lambda_{1}^{n} - \lambda_{2}^{n}) \\ b(\lambda_{1}^{n} - \lambda_{2}^{n}) & \lambda_{1}^{n} (d - \lambda_{2}) - \lambda_{2}^{n} (d - \lambda_{1}) \end{pmatrix}$$

$$= \begin{pmatrix} A_{n+1} - d A_{n} & -\frac{d^{2} - \Gamma d + \mathcal{D}}{b} A_{n} \\ b A_{n} & d A_{n} - \mathcal{D} A_{n-1} \end{pmatrix}$$

$$= \begin{pmatrix} A_{n+1} - d A_{n} & c A_{n} \\ b A_{n} & d A_{n} - \mathcal{D} A_{n-1} \end{pmatrix}.$$
(27)

If b=0, then the eigenvalues of M are $\lambda_1=a$ and $\lambda_2=d$ $(a\neq d)$. Thus,

$$\boldsymbol{M} = \boldsymbol{V} \Lambda \boldsymbol{V}^{-1} = \begin{pmatrix} 1 & \frac{-c}{a-d} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \begin{pmatrix} 1 & \frac{c}{a-d} \\ 0 & 1 \end{pmatrix}, \tag{28}$$

and so (for $n \in \mathbb{N}$)

$$\boldsymbol{M}^{n} = \boldsymbol{V} \Lambda^{n} \boldsymbol{V}^{-1} = \begin{pmatrix} 1 & \frac{-c}{a-d} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a^{n} & 0 \\ 0 & d^{n} \end{pmatrix} \begin{pmatrix} 1 & \frac{c}{a-d} \\ 0 & 1 \end{pmatrix}$$
$$= \begin{pmatrix} a^{n} & \frac{c(a^{n}-d^{n})}{a-d} \\ 0 & d^{n} \end{pmatrix}, \tag{29}$$

which yields equation Eq. 24 for b = 0.

E MULTI-STABILITY

We next investigate multi-stability involving chaotic attractors. To do so, we need to obtain the regions of stability of various orbits and their overlap regions in the parameter space. In Sect. H.3, the existence of chaotic attractors was considered in a theoretical framework by finding necessary and sufficient conditions for the manifestation of homoclinic intersections. For k-periodic attractors ($k \ge 1$), the existence regions can be determined following a straightforward approach as briefly explained in Appx. G.

Remark E.1. The map Eq. 8 can exihit quasi-periodic orbits which are composed of an infinite number of points lying on an invariant closed curve, see Fig. 10. As illustrated in Fig. 10C, the largest Lyapunov exponent (red) is zero and the other one (blue) is negative which confirms the quasi-periodic orbit is stable. Multi-stability involving quasi-periodic attractors may therefore also be possible, but here we will focus on multi-stability of periodic and chaotic attractors.

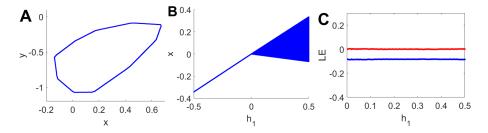


Figure 10: A) Phase portrait of a stable quasi-periodic orbit, B) bifurcation diagram, C) Lyapunov exponents (LEs). Parameter settings: $a_l = 0.31$, $b_l = -.6$, $a_r = -1.1$, $b_r = -1.54$, c = 0.9, d = 0.3, $h_1 = 1$, $h_2 = 0$.

Finding the overlapping stability regions of different periodic orbits, we can observe various MABs which result in different multi-stabilities. In Fig. 11A the MAB diagram shows multi-stability of 2-cycles and 3-cycles after the bifurcation occurred. Fig. 11B illustrates multi-stability of (i) 5-cycles and 2-cycles (before the bifurcation) and (ii) 3-cycles and fixed points (after the bifurcation). Note that in both cases the term $(1-d)h_1+ch_2$ vanishes at the bifurcation point, so that none of the bifurcation points in Fig. 11 can be obtained through the 2d PL normal form map. Finding the

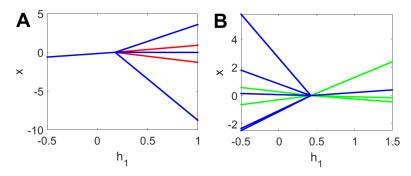


Figure 11: MAB diagrams for parameter settings: A) $c=0.71,\ d=0.2,\ b_l=-0.4,\ b_r=0.5,\ a_l=0.253,\ a_r=-2.83,\ h_2=-0.2;$ and B) $c=0.95,\ d=0.1,\ b_l=-1.21,\ b_r=-0.9,\ a_l=-2.98,\ a_r=-0.73,\ h_2=-0.4.$

existence regions for homoclinic intersections (see Sect. H.3), we can investigate multi-stability involving chaotic attractors, as illustrated in the following examples. Suppose that all system parameters except a_r have the same values as in Eq. 22. Fig. 12B represents multi-stability of a chaotic attractor and a stable 4-cycle for $a_r=0.99$. Changing h_1 from negative to positive values, a multiple attractor bifurcation occurs at $h_1\approx 0.282$. In Fig. 13B multi-stability of a 2-band chaotic attractor and a stable 3-cycle is apparent for $a_r=0.3$. In this case, again, the system undergoes an MAB at $h_1\approx 0.282$. Fig. 14 illustrates the coexistence of a chaotic and a 3-band chaotic attractor in two cases. As shown in Fig. 14A, there is an extended basin of attraction for the simple

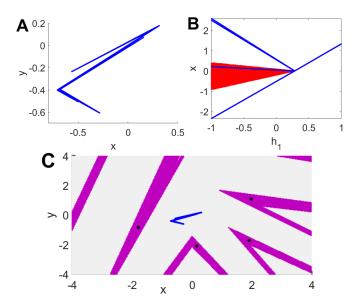


Figure 12: A) Phase portrait of a chaotic attractor, B) MAB diagram and multi-stability of a chaotic attractor (red) and a stable 4-cycle (blue). C) Basins of attraction of the coexisting attractors. Parameter settings: $a_r = 0.99, \ b_r = -0.75, \ a_l = -1.77, \ b_l = -0.9, \ c = 0.6, \ d = 0.15, \ h_2 = -0.4, \ h_1 = -0.7.$

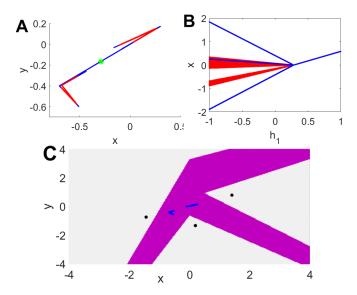


Figure 13: A) Unstable manifold (blue) of the saddle fixed point $\mathcal{O}_{\mathcal{L}}^* = (-0.28848, -0.16514)^{\mathsf{T}}$ (green star) and a 2-band chaotic attractor lying on it. B) MAB diagram and multi-stability of a 2-band chaotic attractor (red) and a stable 3-cycle (blue). C) Basins of attraction of the coexisting attractors. Parameter settings: $a_r = 0.3$, $b_r = -0.75$, $a_l = -1.77$, $b_l = -0.9$, c = 0.6, d = 0.15, $b_2 = -0.4$, $b_1 = -0.7$.

chaotic attractor whereas the 3-band attractor has a smaller basin. By contrast, the basins are more interwoven in Fig. 14B. Furthermore, in this case, one may consider the attractors in red as 2-band chaotic attractors based on the MAB diagram. Plotting the phase portrait, however, reveals they are 3-band chaotic attractors while projecting them onto either x or y axis yields 2-band attractors. Varying h_1 from negative to positive, the systems undergoes a MAB at which $(1-d)h_1 + ch_2 = 0$.

Fig. 15B demonstrates the multi-stability of a chaotic attractor and a stable 3-cycle before a MAB. The bifurcation results in another type of multi-stability involving a stable fixed point and a 4-cycle.

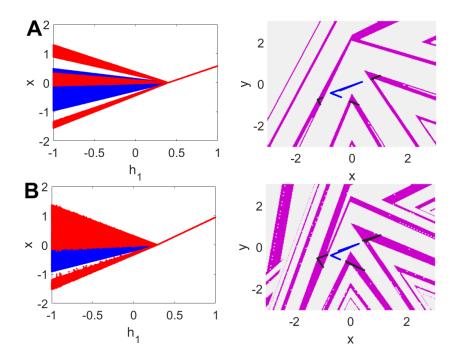


Figure 14: MAB diagram, multi-stability of a chaotic (blue) and a 3-band chaotic (red) attractor and their basins of attraction for parameter values: A) $a_r = 0.7$, $b_r = -0.75$, $a_l = -1.77$, $b_l = -0.9$, c = 0.6, d = 0.4, $h_2 = -0.4$, $h_1 = -0.7$; B) $a_r = 0.79$, $b_r = -0.75$, $a_l = -1.77$, $b_l = -0.9$, c = 0.6, d = 0.17, $h_2 = -0.4$, $h_1 = -0.7$.

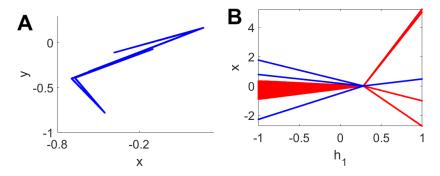


Figure 15: A) Phase portrait of a chaotic attractor, B) MAB diagram which shows the multi-stability of (i) a chaotic attractor (red) and a stable 3-cycle (blue) and (ii) a stable fixed point (blue) and a 4-cycle (red). Parameter settings: $a_r = 0.55$, $b_r = -1.5$, $a_l = -1.74$, $b_l = -0.9$, c = 0.6, d = 0.15, $h_2 = -0.4$, $h_1 = -0.7$.

F 2-UNIT PLRNNS

Consider the 2-unit PLRNN defined by

$$z_t = F(z_{t-1}) = (A + WD_{\Omega(t-1)})z_{t-1} + h$$

:= $W_{\Omega(t-1)}z_{t-1} + h$, (30)

where $z_t = (z_{1t}, z_{2t})^{\mathsf{T}} \in \mathbb{R}^2$ indicates the neural state vector at time $t = 1 \cdots T$, the vector $\boldsymbol{h} = (h_1, h_2)^{\mathsf{T}} \in \mathbb{R}^2$ is the bias term, the matrices

$$\mathbf{A} = \begin{pmatrix} a_{11} & 0 \\ 0 & a_{22} \end{pmatrix}, \quad \mathbf{W} = \begin{pmatrix} w_{11} & w_{12} \\ w_{21} & w_{22} \end{pmatrix}, \tag{31}$$

consist of all (linear) auto-regression weights and connection weights respectively, $D_{\Omega(t)} := \operatorname{diag}(\boldsymbol{d}_{\Omega(t)})$ is a diagonal matrix and $\boldsymbol{d}_{\Omega(t)} := (d_1, d_2)$ an indicator vector with $d_m(z_{m,t}) := d_m = 1$ whenever $z_{m,t} > 0$, m = 1, 2, and zeros otherwise (Durstewitz, 2017; Koppe et al., 2019b). There-

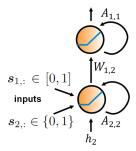


Figure 16: 2-unit PLRNN solution to addition problem (Schmidt et al., 2021).

fore, the phase space of system Eq. 30 is divided into 4 sub-regions by 4 hyper-surfaces as borders. Listing the 4 different configurations of $D_{\Omega(t)}$ as D_{Ω^k} , k = 1, 2, 3, 4, we define 4 matrices

$$W_{\Omega^k} := A + W D_{\Omega^k}. \tag{32}$$

In this case, in every indexed sub-region S_{Ω^k} , k=1,2,3,4, the system dynamics are governed by a different map as follows

$$\boldsymbol{z}_{t+1} = F(\boldsymbol{z}_t) = \boldsymbol{W}_{\Omega^k} \, \boldsymbol{z}_t + \boldsymbol{h}, \qquad \boldsymbol{z}_t \in S_{\Omega^k}. \tag{33}$$

Using the binary number system, all the sub-regions S_{Ω^k} 's can be defined as (Monfared and Durstewitz, 2020a;b)

$$S_{\Omega^{1}} = \hat{S}_{0} = \hat{S}_{(0\,0)_{2}^{*}} = \hat{S}_{0\,0} = \left\{ \boldsymbol{z}_{t} \in \mathbb{R}^{2}; \, z_{1t}, z_{2t} \leq 0 \right\},$$

$$S_{\Omega^{2}} = \hat{S}_{1} = \hat{S}_{(0\,1)_{2}^{*}} = \hat{S}_{1\,0} = \left\{ \boldsymbol{z}_{t} \in \mathbb{R}^{2}; z_{1t} > 0, z_{2t} \leq 0 \right\},$$

$$S_{\Omega^{3}} = \hat{S}_{2} = \hat{S}_{(1\,0)_{2}^{*}} = \hat{S}_{0\,1} = \left\{ \boldsymbol{z}_{t} \in \mathbb{R}^{2}; z_{2t} > 0, z_{1t} \leq 0 \right\},$$

$$S_{\Omega^{4}} = \hat{S}_{3} = \hat{S}_{(1\,1)_{2}^{*}} = \hat{S}_{1\,1} = \left\{ \boldsymbol{z}_{t} \in \mathbb{R}^{2}; z_{1t}, z_{2t} > 0 \right\},$$

$$(34)$$

where each subindex d of \hat{S} , $0 \le d \le 3$, is associated with a sequence $d_2 d_1$ of binary digits. The notation $(d_1 d_2)_2^*$ in building each corresponding sequence stands for the mirror image of the binary representation of d with M digits. By mirror image here we mean writing digits $d_1 d_2$ from right to left, i.e. $d_2 d_1$. Denoting switching boundaries $\Sigma_{ij} = \bar{S}_{\Omega^i} \cap \bar{S}_{\Omega^j}$ between every pair of successive sub-regions S_{Ω^i} and S_{Ω^j} with $i,j \in \{1,2,3,4\}$, we can rewrite map Eq. 30 as

$$\boldsymbol{z}_{t+1} = F(\boldsymbol{z}_{t}) = \begin{cases} F_{1}(\boldsymbol{z}_{t}) = \boldsymbol{W}_{\Omega^{1}} \boldsymbol{z}_{t} + \boldsymbol{h}; & \boldsymbol{z}_{t} \in \bar{S}_{\Omega^{1}} \\ F_{2}(\boldsymbol{z}_{t}) = \boldsymbol{W}_{\Omega^{2}} \boldsymbol{z}_{t} + \boldsymbol{h}; & \boldsymbol{z}_{t} \in \bar{S}_{\Omega^{2}} \\ F_{3}(\boldsymbol{z}_{t}) = \boldsymbol{W}_{\Omega^{3}} \boldsymbol{z}_{t} + \boldsymbol{h}; & \boldsymbol{z}_{t} \in \bar{S}_{\Omega^{3}} \\ F_{4}(\boldsymbol{z}_{t}) = \boldsymbol{W}_{\Omega^{4}} \boldsymbol{z}_{t} + \boldsymbol{h}; & \boldsymbol{z}_{t} \in \bar{S}_{\Omega^{4}} \end{cases}$$
(35)

If we consider the matrix W in Eq. 31 as

$$\boldsymbol{W} = \begin{pmatrix} w_{11} & 0 \\ w_{21} & 0 \end{pmatrix},\tag{36}$$

then, applying definition Eq. 32, we have

$$oldsymbol{W}_{\Omega^1} = oldsymbol{W}_{\Omega^3} = egin{pmatrix} a_{11} & 0 \ 0 & a_{22} \end{pmatrix} = oldsymbol{A},$$

$$\mathbf{W}_{\Omega^2} = \mathbf{W}_{\Omega^4} = \begin{pmatrix} a_{11} + w_{11} & 0 \\ w_{21} & a_{22} \end{pmatrix}. \tag{37}$$

In this case, there are only two different sub-regions divided by one border. That means considering $a_{11} = a_l$, $a_{11} + w_{11} = a_r$, $w_{21} = b_r$, $a_{22} = d$, the system Eq. 35 can be written as a 2-dimensional PL map with one switching boundary. On the other hand, studying generic 2d PL maps, we can investigate dynamics of 2-unit PLRNNs given by Eq. 35 locally near one border.

G STABILITY REGIONS OF k-CYCLES FOR THE MAP EQ. 8

The following is a concise description of how to determine the stability regions of k-cycles for k = 1, 2, 3. For sake of clarity, we focus on basic orbits. The regions of stability of other k-periodic orbits can be obtained similarly.

The fixed points of the map Eq. 8 are given by

$$\mathcal{O}_{\mathcal{L}/\mathcal{R}} = \left(\frac{(1-d)h_1 + ch_2}{(1-d)(1-a_{l/r}) - b_{l/r}c}, \frac{b_{l/r}h_1 + (1-a_{l/r})h_2}{(1-d)(1-a_{l/r}) - b_{l/r}c}\right)^{\mathsf{T}}.$$
(38)

 $\mathcal{O}_{\mathcal{L}}$ and $\mathcal{O}_{\mathcal{R}}$ are admissible fixed points iff $x_{\mathcal{L}} < 0$ and $x_{\mathcal{R}} > 0$ respectively; otherwise they are virtual. Accordingly, the existence regions of the fixed points $\mathcal{O}_{\mathcal{L}}$ and $\mathcal{O}_{\mathcal{R}}$ are given by

$$E_{\mathcal{O}_{\mathcal{L}}} = \left\{ (h_1, h_2, a_l, b_l, c, d) \, \middle| \, \frac{(1 - d) h_1 + c h_2}{(1 - d)(1 - a_l) - b_l c} < 0 \right\},$$

$$E_{\mathcal{O}_{\mathcal{R}}} = \left\{ (h_1, h_2, a_r, b_r, c, d) \, \middle| \, \frac{(1 - d) h_1 + c h_2}{(1 - d)(1 - a_r) - b_r c} > 0 \right\}.$$
(39)

Let $D_{\mathcal{L}/\mathcal{R}}$ and $\mathcal{P}_{\mathcal{L}/\mathcal{R}}(\pm 1)$ be the determinants and characteristic polynomials of the Jacobian matrices. Then, considering the conditions $\mathcal{P}_{\mathcal{L}/\mathcal{R}}(\pm 1)>0$ and $D_{\mathcal{L}/\mathcal{R}}<1$, the stability region of fixed points can be derived as

$$S_{\mathcal{L}/\mathcal{R}} = \left\{ (h_1, h_2, a_{l/r}, b_{l/r}, c, d) \in E_{\mathcal{O}_{\mathcal{L}/\mathcal{R}}} \mid a_{l/r}d - b_{l/r}c < 1, \\ 1 \pm (a_{l/r} + d) + a_{l/r}d - b_{l/r}c > 0 \right\}.$$

$$(40)$$

Likewise, the stability region of the 2-cycle $\mathcal{O}_{\mathcal{RL}}$ is

$$S_{\mathcal{RL}} = \left\{ (h_1, h_2, a_{l/r}, b_{l/r}, c, d) \in E_{\mathcal{O}_{\mathcal{RL}}} \mid -1 < (a_r d - b_r c)(a_l d - b_l c) < 1, \\ - (a_r d - b_r c)(a_l d - b_l c) - 1 < c(b_l + b_r) + d^2 + a_l a_r < (a_r d - b_r c)(a_l d - b_l c) + 1 \right\},$$

$$(41)$$

in which

$$E_{\mathcal{O}_{\mathcal{RL}}} = \left\{ (h_1, h_2, a_l, b_l, c, d) \mid \frac{((1-d)h_1 + c h_2)(a_l + d + a_l d - b_l c + 1)}{(a_r d - b_r c)(a_l d - b_l c) - c(b_l + b_r) - d^2 - a_l a_r + 1} > 0, \right\}$$

 $\frac{\left((1-d)h_1 + c h_2\right)\left(a_r + d + a_r d - b_r c + 1\right)}{(a_r d - b_r c)(a_l d - b_l c) - c(b_l + b_r) - d^2 - a_l a_r + 1} < 0 \right\},\tag{42}$

represents the existence region of $\mathcal{O}_{\mathcal{RL}}$.

Analogously, the existence region of the 3-cycle $\mathcal{O}_{\mathcal{R}\mathcal{L}^2}$ is given by

$$S_{\mathcal{R}\mathcal{L}^{2}} = \left\{ (h_{1}, h_{2}, a_{l/r}, b_{l/r}, c, d) \in E_{\mathcal{O}_{\mathcal{R}\mathcal{L}^{2}}} \mid -1 < (a_{l} d - b_{l} c)^{2} (a_{r} d - b_{r} c) < 1, -(a_{l} d - b_{l} c)^{2} (a_{r} d - b_{r} c) - 1 < a_{l}^{2} a_{r} + d^{3} + c (a_{l} b_{l} + a_{l} b_{r} + a_{r} b_{l} + d(2 b_{l} + b_{r})) < (a_{l} d - b_{l} c)^{2} (a_{r} d - b_{r} c) + 1 \right\},$$

$$(43)$$

where

$$E_{\mathcal{O}_{\mathcal{R}\mathcal{L}^2}} = \left\{ (h_1, h_2, a_l, b_l, c, d) \, \middle| \, \frac{\left((1 - d)h_1 + c \, h_2 \right) G_1}{G} > 0, \right.$$

$$\left. \frac{\left((1 - d)h_1 + c \, h_2 \right) K_1}{G} < 0, \, \frac{\left((1 - d)h_1 + c \, h_2 \right) H_1}{G} < 0 \right\}, \tag{44}$$

and

$$G_{1} = a_{l}^{2} d^{2} + a_{l}^{2} d + a_{l}^{2} - 2 a_{l} b_{l} c d - a_{l} b_{l} c + a_{l} d^{2} + a_{l} d$$

$$+ a_{l} + b_{l}^{2} c^{2} - b_{l} c d + b_{l} c + d^{2} + d + 1,$$

$$G = -a_{l}^{2} a_{r} - d^{3} - c (a_{l} b_{l} + a_{l} b_{r} + a_{r} b_{l} + d (2 b_{l} + b_{r}))$$

$$+ (a_{r} d - b_{r} c) (a_{l} d - b_{l} c)^{2} + 1,$$

$$K_{1} = a_{r} + d + a_{l} a_{r} + b_{l} c + a_{r} d + a_{r} d^{2} + d^{2} + a_{l} a_{r} d - a_{l} b_{r} c$$

$$- b_{r} c d + a_{l} a_{r} d^{2} + b_{l} b_{r} c^{2} - a_{l} b_{r} c d - a_{r} b_{l} c d + 1,$$

$$H_{1} = a_{l} + d + a_{l} a_{r} + a_{l} d + b_{r} c + a_{l} d^{2} + d^{2} + a_{l} a_{r} d - a_{r} b_{l} c$$

$$- b_{l} c d + a_{l} a_{r} d^{2} + b_{l} b_{r} c^{2} - a_{l} b_{r} c d - a_{r} b_{l} c d + 1.$$

$$(45)$$

H ALGORITHMS

H.1 METHODOLOGICAL DETAILS

All experiments were run on a single CPU, specifically an Intel Xeon Gold 6132 with 512GB RAM and an Intel Xeon Gold 6248 with 832GB RAM.

For Fig. 3.4A, every 500 epochs 10 latent trajectories of length 1000 were produced to determine the pool of linear subregions \mathcal{S} traversed by the model. To exclude transients, the first 99 time steps were discarded from each trajectory. The regularized linear subregions \mathcal{S}_{reg} were chosen at random as a fraction (e.g. 1%, 10%, depending on the hyperparameter) from this pool \mathcal{S} for each epoch. To evaluate the invertibility ($\mathcal{S}_+/\mathcal{S}$), 20 trajectories of length 10000 were generated and a pool of subregions computed from. The first 99 time steps were again discarded to remove possible transients.

H.2 ADDITIONAL ALGORITHMS

In the main text, the algorithm was formulated for the PLRNN. The same procedure as described in Algo. 2 can analogously be applied to a 1-hidden layer PLRNN, called the shallow PLRNN

1458	Parameter	Lorenz63	Oscillator	Lorenz63	Duffing	Decision making	Lorenz63	Empirical
1459		Fig. 3.4A	Fig. 3.4B	Fig. 3.4C	Fig. 4A	Fig. 4B	Fig. 4C	Fig. 4D
	Model	ALRNN	ALRNN	ALRNN	shallowPLRNN	ALRNN	shallowPLRNN	ALRNN
1460	Latent dim	10/20/30/50	40	30	2	15	3	25
1461	Hidden dim	-	-	-	10	-	20	-
	#ReLUs	Latent dim - 3	15	8	-	6	-	6
1462	Sequence length	200	25	100	100	100	100	200
1463	Gaussian noise	0.0	0.0	0.0	0.0	0.01	0.05	0.02
	$\lambda_{ ext{invert}}$	0.0/0.1 · exp(Latent dim)	0.0/1e15	0.0/1e10	0.0	0.2	0.0	0.3
1464	Batch Size	16	16	16	32	16	16	16
1465	Epochs	10000	1000	1000	10000	20000	1000	2000
	Start LR	0.001	0.001	0.005	0.001	0.005	0.005	0.004
1466	TF interval	16	10	15	15	15	15	20

Table 3: Parameter configurations for the different experiments. TF = teacher forcing. LR = learning rate.

(shPLRNN) (Hess et al., 2023):

1467 1468

1469

1470 1471

14721473

1474

1476 1477

14781479

1480

1481 1482 1483

$$z_{t+1} = Az_t + W_1D_t(W_2z_t + h_2) + h_1$$
 (46)

where $A \in \mathbb{R}^{M \times M}$, $W_1 \in \mathbb{R}^{M \times H}$, $W_2 \in \mathbb{R}^{H \times M}$, $h_2 \in \mathbb{R}^H$, $h_1 \in \mathbb{R}^M$ and H > M. The inversion of this map yields

$$z_{t-1} = (A + W_1 D_{t-1} W_2)^{-1} (z_t - W_1 D_{t-1} h_2 - h_1)$$
(47)

Fallback algorithm For systems with complex dynamics, such as in Fig. 3B where trajectories jump between disjoint subregions, the primary algorithm may struggle to converge. In such cases, we

Algorithm 2 Backtracking Time Series in a ReLU based RNNs

```
1484
             1: z_T \leftarrow an initial State
1485
             2: \theta \leftarrow \text{Parameters}
1486
             3: Initialize list: S = [z_T]
1487
             4: for t = T : 1 do
1488
                      z_t = S[T - t]
             5:
1489
                      D_{t-1} \leftarrow \operatorname{diag}(z_t > 0)
                                                                                                   ▶ Initialize D as a diagonal matrix
                      z_{t-1}^* \leftarrow F^{-1}(\theta, D_{t-1}, z_t)
             7:
                                                                                                             ▶ Perform a backward step
                      \tilde{z}_t \leftarrow F(\theta, z_{t-1}^*)
             8:
                                                                                                                ⊳Perform a forward step
1491
             9:
                      if \tilde{z}_t = z_t then
1492
           10:
                           S \leftarrow S \cup \{z_{t-1}^*\}
                                                                                                              ⊳If forward step is correct
1493
                      else
           11:
1494
                           D_{t-1} \leftarrow \operatorname{diag}(z_{t-1}^* > 0)
                                                                                                       ⊳Update D with new candidate
           12:
1495
                           z_{t-1}^* \leftarrow F^{-1}(\theta, D_{t-1}, z_t)
           13:
                                                                                                                  ▶ Retry backward step
1496
                           \tilde{z}_t \leftarrow F(\theta, z_{t-1}^*)
                                                                                                                     ▶ Retry forward step
           14:
1497
                           if \tilde{z}_t = z_t then
           15:
1498
                                S \leftarrow S \cup \{z_{t-1}^*\}
           16:
                                                                                                             ▶ If forward step is correct
1499
           17:
1500
                                \tilde{z}_t \leftarrow \text{TryPreviousRegions}(\theta, D\_pool, z_t, z_{t-1}^*, \tilde{z}_t)
                                                                                                                  18:
1501
           19:
                                if \tilde{z}_t \neq z_t then
1502
           20:
                                     \tilde{z}_t \leftarrow \text{TryBitflips}(\theta, z_t, z_{t-1}^*, \tilde{z}_t)
                                                                                                  if \tilde{z}_t = z_t then
1503
           21:
                                          S \leftarrow S \cup \{z_{t-1}^*\}
           22:
1504
           23:
                                     else
1505
           24:
                                          return
1506
           25:
                                     end if
1507
                                end if
           26:
1508
                           end if
           27:
1509
                      end if
           28:
1510
           29: end for
1511
           30: return trajectory
```

1513

1514

1515

1516

use a more robust fallback method (Alg. 3) that perturbs seed points along the analytically defined local manifold and iterates F_{θ} to generate dense support vectors. As manifolds can re-enter the same subregion in multiple folds, we apply HDBSCAN (McInnes et al., 2017; Ester et al., 1996) to cluster support vectors into distinct segments. Although computationally more demanding, this fallback reliably captures manifolds with discontinuous or folding structures when sequential tracing fails.

```
Algorithm 3 Finding stable/unstable manifolds: fallback algorithm
1538
             1: (P, E) \leftarrow SCYFI
                                                                                                      ⊳P: Fixed Point, E: Eigenvectors
1539
1540
             2: S^s \leftarrow \emptyset
                                                                                                                          ⊳Stable Manifolds
             S^u \leftarrow \emptyset
                                                                                                                      ⊳Unstable Manifolds
1541
             4: for i=1:N_1 do
1542
                      z_0 = P
                                                                                                       \trianglerightFor N_1 different initialisations
             5:
1543
                      for v^u \in E^u do
             6:
                           z_0 += v^u \cdot rand()
             7:
                                                                                                                  ⊳Perturbe into subspace
1545
                      end for
             8:
1546
                      T^u \leftarrow GetForwardTS(z_0)
             9:
1547
                      S^u \leftarrow S^u \cup \{T^u\}
            10:
1548
            11: end for
1549
            12: for i=1:N_2 do
1550
            13:
                      z_0 = P
                                                                                                      \triangleright For N_2 different initialisations
                      \text{ for } v^s \in E^s \ \ \text{ do }
            14:
1551
            15:
                           z_0 += v^s \cdot rand()
                                                                                                                   ▶ Perturb into subspace
1552
            16:
                      end for
1553
            17:
                      T^s \leftarrow GetBackwardTS(z_0)
1554
                      S^s \leftarrow S^s \cup \{T^s\}
            18:
1555
            19: end for
1556
           20: S^s \leftarrow \emptyset
                                                                                                        ▷ Piecewise linear manifold fits
1557
            21: \widetilde{S}^u \leftarrow \emptyset
1558
            22: for each D \in D_{\Omega} do
1559
                      S^s_\Omega \leftarrow S^s \cap D_\Omega
            23:
                                                                                                             1560
                      S_{\Omega}^{ii} \leftarrow S^u \cap D_{\Omega}
            24:
1561
           25:
                      (\tilde{C}^s_{\Omega}, C^u_{\Omega}) \leftarrow \mathrm{Fit}((S^s_{\Omega}, S^u_{\Omega}))
                                                                                                                   1562
                      \widetilde{S}^s \leftarrow \widetilde{S}^s \cup \{C^s_{\Omega}\}
           26:
1563
                      \widetilde{S}^u \leftarrow \widetilde{S}^u \cup \{C^u_\Omega\}
           27:
1564
           28: end for
1565
           29: return (\widetilde{S}^s, \widetilde{S}^u)
                                                                                                           ▶ Piecewise linear manifolds
```

```
1: function BACKWARDFORWARD(\Theta, D, z)
1567
                    z^* \leftarrow F - 1(\Theta, D, z)
           2:
1568
                    \tilde{z} \leftarrow F(\Theta, z^*)
           3:
1569
                    return z^*, \tilde{z}
           4:
1570
           5: end function
1571
           6: function TRYPREVIOUSREGIONS(\Theta, D pool, z, z*, \tilde{z})
1572
                    for D \in D\_pool do
                         z^*, \tilde{z} \leftarrow \text{BackwardForward}(\Theta, D, z)
1573
           8:
                         if \tilde{z} = z then
1574
           9:
          10:
                             return z^*
1575
          11:
                         end if
          12:
                    end for
          13: end function
          14: function TRYBITFLIPS(\Theta, z, z^*, \tilde{z})
                    for k = 1 : num\_relus do
          15:
1580
                         D\_versions \leftarrow generate\_bitflip\_k()
          16:
1581
          17:
                         for D \in D\_versions do
                             z^*, \tilde{z} \leftarrow \text{BackwardForward}(\Theta, D, z)
          18:
          19:
                             if \tilde{z} = z then
          20:
                                  return z^*
1585
          21:
                             end if
          22:
                         end for
          23:
                    end for
1587
          24: end function
```

H.3 Existence of homoclinic intersections

1590

1591 1592

1593

1594

1596

1597

1598

1602

1604

1607

1608

1609 1610

1611 1612

1613 1614 1615

1616

1617 1618

1619

A horseshoe structure is generated when there is a homoclinic intersection between stable and unstable manifolds of a saddle fixed point and therefore an infinite number of intersections. The occurrence of a homoclinic intersection implies the existence of a chaotic orbit (Wiggins, 1988). Here we will obtain a necessary and sufficient condition for the occurrence of homoclinic intersections in order to find a general condition for the existence of chaos. For this we consider the system Eq. 8 where T is invertible and the matrices $\mathbf{A}_{\mathcal{L}}$ and $\mathbf{A}_{\mathcal{R}}$ are non-singular and have no eigenvalue equal to 1. Note that as the parameter space of general 2d maps is eight-dimensional, the computation of manifolds could be more challenging than for 2d normal form maps. We approach this issue by establishing the conditions of homoclinic intersections using the equations of the stable and unstable eigenlines.

H.3.1 ANALYTICAL CONDITION FOR HOMOCLINIC INTERSECTION OF THE FIRST AND SECOND FOLD POINTS

Let $\mathcal{O}_{\mathcal{L}}^* = (x_{\mathcal{L}}^*, y_{\mathcal{L}}^*)^\mathsf{T}$ be an admissible saddle fixed point in the left sub-region \mathcal{L} . Since $\mathcal{O}_{\mathcal{L}}^*$ is a saddle, $A_{\mathcal{L}}$ has one stable and one unstable eigenvalue λ_s and λ_u , respectively. Let us denote the line generated by the associated stable eigenvector by $\boldsymbol{\ell}_s^*$ and the line produced by the corresponding unstable eigenvector by $\boldsymbol{\ell}_u^*$. Since the unstable eigenvector is $\boldsymbol{v}_u = (v_1, v_2)^\mathsf{T} = (\frac{\lambda_u - d}{b_l}, 1)^\mathsf{T}$, $\boldsymbol{\ell}_u^*$ hits the border x = 0 at $P_0 = (0, y_0)^\mathsf{T} \in \Sigma$ where

$$y_{0} = y_{\mathcal{L}}^{*} - x_{\mathcal{L}}^{*} \frac{v_{2}}{v_{1}}$$

$$= \frac{b_{l} h_{1} + (1 - a_{l}) h_{2}}{1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}}} - \frac{(1 - d) h_{1} + c h_{2}}{1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}}} \left(\frac{b_{l}}{\lambda_{u} - d}\right)$$

$$= \frac{\left(b_{l} h_{1} + (1 - a_{l}) h_{2}\right) (\lambda_{u} - d) - b_{l} \left(c h_{2} + (1 - d) h_{1}\right)}{(\lambda_{u} - d) (1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}})}.$$
(48)

The image of P_0 is the first fold point of the unstable manifold of $\mathcal{O}_{\mathcal{L}}^*$, and so all its images will also be fold points. Its coordinate is $P_1 = (c y_0 + h_1, d y_0 + h_2)^\mathsf{T}$. The image of the first fold point is

the second fold point $P_2 = (x_2, y_2)^T$ with coordinates

$$\begin{cases} x_2 = c(a_l + d)y_0 + (a_l + 1)h_1 + ch_2 \\ y_2 = (b_l c + d^2)y_0 + b_l h_1 + (d + 1)h_2 \end{cases}, \quad \text{if } c y_0 + h_1 < 0$$

$$\begin{cases} x_2 = c(a_r + d)y_0 + (a_r + 1)h_1 + ch_2 \\ y_2 = (b_r c + d^2)y_0 + b_r h_1 + (d + 1)h_2 \end{cases}, \quad \text{if } c y_0 + h_1 > 0$$

$$(49)$$

Now we check whether or not the points P_1 and P_2 are on opposite sides of the stable eigenline ℓ_s^* . When P_1 and P_2 are on opposite sides of ℓ_s^* , then the unstable manifold must have intersected the stable manifold. Thus, we have a homoclinic intersection which implies the occurrence of chaotic dynamics. Since the stable eigenvector is $\mathbf{v}_s = (\frac{\lambda_s - d}{b_l}, 1)^\mathsf{T}$ and the eigenline ℓ_s^* passes through $\mathcal{O}_{\mathcal{L}}^* = (x_{\mathcal{L}}^*, y_{\mathcal{L}}^*)^\mathsf{T}$, ℓ_s^* can be computed as

$$\ell_{s}^{*}: \frac{\lambda_{s} - d}{b_{l}} y - x + \frac{(1 - d) h_{1} + c h_{2}}{1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}}} - \left(\frac{\lambda_{s} - d}{b_{l}}\right) \frac{b_{l} h_{1} + (1 - a_{l}) h_{2}}{1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}}} =: L(x, y) = 0.$$
(50)

Now there are two possibilities:

Case I: $L(x_1, y_1) \cdot L(x_2, y_2) \leq 0$

If $L(x_1, y_1) \cdot L(x_2, y_2) < 0$, then P_1 and P_2 are on opposite sides of the stable manifold, while $L(x_1, y_1) \cdot L(x_2, y_2) = 0$ implies that at least one of the points P_1 and P_2 lies exactly on the stable manifold. In both cases there exists a homoclinic intersection, i.e., whenever we have

$$\left(\frac{\lambda_s - d}{b_l}(dy_0 + h_2) - (cy_0 + h_1) + \mathcal{M}\right) \left(\frac{\lambda_s - d}{b_l}\left((b_{l/r}c + d^2)y_0 + b_{l/r}h_1 + (d+1)h_2\right) - \left(c(a_{l/r} + d)y_0 + (a_{l/r} + 1)h_1 + ch_2\right) + \mathcal{M}\right) \le 0,$$
(51)

where $\mathcal{M}=rac{(1-d)\,h_1+c\,h_2}{1-\Gamma_{\mathcal{L}}+D_{\mathcal{L}}}-\left(rac{\lambda_s-d}{b_l}
ight)rac{b_l\,h_1+(1-a_l)\,h_2}{1-\Gamma_{\mathcal{L}}+D_{\mathcal{L}}}.$

The homoclinic intersection point $P_{hom} = (x_{hom}, y_{hom})^{\mathsf{T}}$ is the point of intersection between the line joining the two fold points and the stable eigenline ℓ_*^* , and hence given by

$$\begin{cases} x_{hom} = (x_2 - x_1)\beta + x_1 \\ y_{hom} = (y_2 - y_1)\beta + y_1 \end{cases},$$
 (52)

where

$$\beta = \frac{x_1 - \frac{\lambda_s - d}{b_l} - \mathcal{M}}{\frac{\lambda_s - d}{b_l} (y_2 - y_1) - (x_2 - x_1)}.$$
 (53)

Finally, we must ensure that the intersection happens before the stable eigenline hits the border. For this, we need to have $x_{home}\,x_{\mathcal{L}}^*>0$, which implies the intersection point P_{hom} and the fixed point $\mathcal{O}_{\mathcal{L}}^*$ are on the same side of Σ .

Case II: $L(x_1, y_1) \cdot L(x_2, y_2) > 0$

If $L(x_1, y_1) \cdot L(x_2, y_2) > 0$, we need to check whether the unstable manifold intersects with the part of the stable manifold which ensues after folding along the y-axis. For this we have to calculate more points of the global stable manifold. Since T is assumed to be invertible, the global stable manifold is formed by the union of all preimages (inverses) of any rank of the local stable set (a segment of the local stable eigenline). Assume, under the action of T^{-1} , the line ℓ_s^* maps to the y-axis and intersects

it at $\tilde{P}_0 = (0, \tilde{y}_0)^\mathsf{T} \in \Sigma$. Then \tilde{P}_0 is the first fold point of the stable manifold of $\mathcal{O}_{\mathcal{L}}^*$, and \tilde{y}_0 is given by

$$\tilde{y}_0 = \frac{\left(b_l h_1 + (1 - a_l)h_2\right)(\lambda_s - d) - b_l \left(c h_2 + (1 - d)h_1\right)}{(\lambda_s - d)(1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}})}.$$
(54)

The preimage of \tilde{P}_0 is the second fold point $\tilde{P}_{-1} = (\tilde{x}_{-1}, \tilde{y}_{-1})^\mathsf{T}$ with coordinates

$$\begin{cases} \tilde{x}_{-1} = \frac{1}{D_{\mathcal{L}}} \left(-c\tilde{y}_0 + ch_2 - dh_1 \right) \\ \tilde{y}_{-1} = \frac{1}{D_{\mathcal{L}}} \left(a_l \tilde{y}_0 + b_l h_1 - a_l h_2 \right) \end{cases}, \quad \text{if } \frac{\varphi^{\mathsf{T}} (\tilde{P}_0 - \boldsymbol{h})}{D_{\mathcal{L}}} \leq 0$$

$$\begin{cases} \tilde{x}_{-1} = \frac{1}{D_{\mathcal{R}}} \left(-c\tilde{y}_0 + ch_2 - dh_1 \right) \\ \tilde{y}_{-1} = \frac{1}{D_{\mathcal{R}}} \left(a_r \tilde{y}_0 + b_r h_1 - a_r h_2 \right) \end{cases}, \quad \text{if } \frac{\varphi^{\mathsf{T}} (\tilde{P}_0 - \mathbf{h})}{D_{\mathcal{R}}} \ge 0$$
 (55)

where

$$\frac{\varphi^{\mathsf{T}}(\tilde{P}_0 - \boldsymbol{h})}{D_{\mathcal{L}/\mathcal{R}}} = \frac{-d\,h_1}{D_{\mathcal{L}/\mathcal{R}}} - \frac{c}{D_{\mathcal{L}/\mathcal{R}}}(\tilde{y}_0 - h_2). \tag{56}$$

Now the line joining the two fold points \tilde{P}_0 and \tilde{P}_{-1} is given by $\tilde{L}(x,y)=0$ where

$$\tilde{\mathsf{L}}(x,y) := y - \tilde{y}_0 - \frac{\tilde{y}_{-1} - \tilde{y}_0}{\tilde{x}_{-1}} \, x. \tag{57}$$

If $\tilde{\mathsf{L}}(x_1,y_1) \cdot \tilde{\mathsf{L}}(x_2,y_2) < 0$, then P_1 and P_2 are on opposite sides of the stable manifold, and the unstable manifold intersects the stable manifold at $\tilde{P}_{hom} = (\tilde{x}_{hom}, \tilde{y}_{hom})^\mathsf{T}$ with

$$\begin{cases} \tilde{x}_{hom} = (x_2 - x_1)\tilde{\beta} + x_1 \\ \tilde{y}_{hom} = (y_2 - y_1)\tilde{\beta} + y_1 \end{cases},$$
 (58)

where

$$\tilde{\beta} = \frac{(\tilde{y}_0 - y_1)\tilde{x}_{-1} + (\tilde{y}_{-1} - \tilde{y}_0)x_1}{(y_2 - y_1)\tilde{x}_{-1} - (\tilde{y}_{-1} - \tilde{y}_0)(x_2 - x_1)}.$$
(59)

We need to have $\tilde{x}_{home} \, \tilde{x}_{-1} > 0$, which means the intersection point \tilde{P}_{hom} and the point \tilde{P}_{-1} are on the same side of Σ .

Remark H.1. An analogous procedure can be performed to analytically obtain homoclinic intersections for the fixed point $\mathcal{O}_{\mathcal{R}}^* \in \mathcal{R}$.

A recursive algorithm for determining homoclinic intersections

To investigate the existence of homoclinic intersections for the map Eq. 8, we use an algorithm similar to the one proposed in (Roy et al., 2020). As illustrated in Fig. 17, it is based on a recursive procedure as follows:

- 1. Let $\mathcal{O}_{\mathcal{L}}^* = (x_{\mathcal{L}}^*, y_{\mathcal{L}}^*)^\mathsf{T} \in \mathcal{L}$ be an admissible saddle fixed point, and λ_s , λ_u stable and unstable eigenvalues of $A_{\mathcal{L}}$. Assume that ℓ_s^* and ℓ_u^* are the stable and unstable eigenlines generated by the associated stable and unstable eigenvectors, respectively. Analogous to Sect. H.3.1, suppose that ℓ_u^* hits the border at $P_0 = (0, y_0)^\mathsf{T} \in \Sigma$ where y_0 is given by Eq. 48.
- 2. Consider the image of P_0 and assume this point is on the right side of the border, i.e. $P_1 \in \mathcal{R}$. Since this point has the coordinate $P_1 = (c y_0 + h_1, d y_0 + h_2)^\mathsf{T}$, it follows that

$$c y_0 + h_1$$

$$= c \frac{(b_l h_1 + (1 - a_l)h_2)(\lambda_u - d) - b_l(c h_2 + (1 - d)h_1)}{(\lambda_u - d)(1 - \Gamma_{\mathcal{L}} + D_{\mathcal{L}})} + h_1 > 0.$$
(60)

Moreover $P_1 \in \ell^{\Sigma}$ is the first fold point of the unstable manifold of $\mathcal{O}_{\mathcal{L}}^*$, and so all its images will also be fold points.

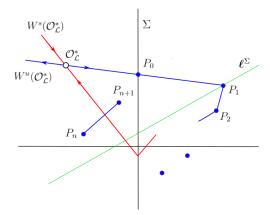


Figure 17: Schematic diagram to illustrate the procedure for finding homoclinic intersections.

3. Suppose that the orbit starting from P_0 will return to the border again. Let n be the the border return time defined as the minimum number of iterations needed for P_0 to cross the border and return to the left side again, i.e., such that all iterations $P_1, P_2, \cdots, P_{n-1}$ lie on the right hand side while $P_n \in \mathcal{L}$. Using a recursive method similar to the one proposed in (Roy et al., 2020), we can compute the n-th iteration, P_n , directly from P_0 . That is, there is no need to compute any other previous iterations $P_1, P_2, \cdots, P_{n-1}$. For this purpose, first we calculate P_n as

$$P_n = T^n(P_0) = \boldsymbol{A}_{\mathcal{R}}^n P_0 + (\boldsymbol{A}_{\mathcal{R}} - \boldsymbol{I})^{-1} (\boldsymbol{A}_{\mathcal{R}}^n - \boldsymbol{I}) \boldsymbol{h}.$$
 (61)

Suppose that $A_{\mathcal{R}}$ has two distinct eigenvalues, then, according to Proposition D.2, the matrix $A_{\mathcal{R}}^n$ has the form Eq. 24. Hence, for $P_0 = (0, y_0)^{\mathsf{T}}$ we have

$$P_{n} = \begin{pmatrix} A_{n+1} - dA_{n} & cA_{n} \\ b_{r}A_{n} & dA_{n} - \mathcal{D}_{\mathcal{R}} A_{n-1} \end{pmatrix} \begin{pmatrix} 0 \\ y_{0} \end{pmatrix} + \frac{-1}{\mathcal{P}_{\mathcal{R}}(1)} \times \\ \begin{pmatrix} 1 - d & c \\ b_{r} & 1 - a_{r} \end{pmatrix} \begin{pmatrix} A_{n+1} - dA_{n} - 1 & cA_{n} \\ b_{r}A_{n} & dA_{n} - \mathcal{D}_{\mathcal{R}} A_{n-1} - 1 \end{pmatrix} \begin{pmatrix} h_{1} \\ h_{2} \end{pmatrix} = \\ \begin{pmatrix} cA_{n} y_{0} - \frac{\left((1 - d) \left[A_{n+1} - dA_{n} - 1 \right] + cb_{r}A_{n} \right) h_{1} + c \left[A_{n} - \mathcal{D}_{\mathcal{R}} A_{n-1} - 1 \right] h_{2}}{\mathcal{P}_{\mathcal{R}}(1)} \\ \begin{pmatrix} (dA_{n} - \mathcal{D}_{\mathcal{R}} A_{n-1}) y_{0} - \frac{b_{r} \left[A_{n}(1 - \Gamma_{\mathcal{R}}) + A_{n+1} - 1 \right] h_{1} + b^{*} h_{2}}{\mathcal{P}_{\mathcal{R}}(1)} \end{pmatrix}$$

$$(62)$$

where A_n is given by Eq. 25; and $\Gamma_{\mathcal{R}}$, $\mathcal{D}_{\mathcal{R}}$, $\mathcal{P}_{\mathcal{R}}$ are the trace, determinant and characteristic polynomial of $A_{\mathcal{R}}$ respectively; and

$$b^* = (a_r - 1)(1 + \mathcal{D}_{\mathcal{R}} A_{n-1}) + (d - \mathcal{D}_{\mathcal{R}}) A_n.$$
 (63)

Now, by Eq. 62 we can compute all iterations and thus find the border return time n for which the unstable manifold passes the border again.

4. Finally, we calculate P_{n+1} and check whether or not the points P_{n+1} and P_n are on opposite sides of the stable eigenline ℓ_s^* , i.e. whether we obtain a homoclinic intersection. For this, let $P_n = (x_n, y_n)^\mathsf{T}$, $P_{n+1} = (x_{n+1}, y_{n+1})^\mathsf{T}$, and consider $\mathsf{L}(x, y)$ given by Eq. 50. If $\mathsf{L}(x_n, y_n) \cdot \mathsf{L}(x_{n+1}, y_{n+1}) < 0$, then P_n and P_{n+1} are on opposite sides of the stable manifold, while $\mathsf{L}(x_n, y_n) \cdot \mathsf{L}(x_{n+1}, y_{n+1}) = 0$ implies at least one of the points P_n or P_{n+1} lies exactly on the stable manifold. In both cases there exists a homoclinic intersection.

1798 1799

1801

1803

1804

1805

1806 1807

1815 1816 1817

1818 1819 1820

1821

1824

1825

1826 1827

1828

1829

1830

1831

1832

1834

1835

Algorithm 4 Investigating Homoclinic Intersections in 2D

```
1783
              1: procedure InvestigateHomoclinicIntersections
1784
                        \mathcal{O}_{\mathcal{L}}^* \leftarrow (x_{\mathcal{L}}^*, y_{\mathcal{L}}^*)^\mathsf{T}

    Admissible saddle fixed point

1785
                        (\lambda_s, \lambda_u) \leftarrow (\text{Stable Eigenvalue}, \text{Unstable Eigenvalue})
1786
                        (\ell_s^*, \ell_u^*) \leftarrow (\text{Stable Eigenline}, \text{Unstable Eigenline})
                        P_1 \leftarrow (cy_0 + h_1, dy_0 + h_2)^\mathsf{T} 
P_1 \in \ell^{\Sigma}
1787
              5:
                                                                                                  \triangleright Unstable eigenline hits the border at P_0
1788
                                                                                             \triangleright Image of P_0 on the right side of the border
              7:
                                                                                                 ▶ First fold point of the unstable manifold
1789
                        n \leftarrow \text{Border return time}
              8:
                                                                                 \triangleright Minimum iterations for P_0 to return to the left side
1790
                        P_n \leftarrow \boldsymbol{A}_{\mathcal{R}}^n P_0 + (\boldsymbol{A}_{\mathcal{R}} - \boldsymbol{I})^{-1} (\boldsymbol{A}_{\mathcal{R}}^n - \boldsymbol{I}) \boldsymbol{h}
                                                                                                                 \triangleright Compute P_n directly from P_0
              9:
1791
             10:
                        P_{n+1} \leftarrow \text{Next iteration}
                                                                                                           ▷ Check for homoclinic intersection
1792
                        if L(x_n, y_n) \cdot L(x_{n+1}, y_{n+1}) < 0 then
             11:
1793
             12:
                              Return "Homoclinic intersection exists"
             13:
1795
             14:
                              Return "No homoclinic intersection"
1796
             15:
                        end if
1797
             16: end procedure
```

Remark H.2. A similar algorithm can be devised for the fixed point $\mathcal{O}_{\mathcal{R}}^* \in \mathcal{R}$.

Remark H.3. Suppose that the first iteration of P_0 is on the left hand side $(P_1 \in \mathcal{L})$, but moves to the right side of the border after some iterations, i.e., there is some $k_0^* > 1$ such that $P_{k_0^*} \in \mathcal{R}$. Then a similar procedure could be applied for $P_{k_0^*}$ to find homoclinic intersections. In that case we can obtain the border return time needed for $P_{k_0^*}$ to return to the left hand side again. Using Proposition D.2, the iterations of $P_{k_0^*} = (x_0^*, y_0^*)^\mathsf{T}$ can be computed as

$$P_{k_{n}^{*}} = \begin{pmatrix} x_{n}^{*} \\ y_{n}^{*} \end{pmatrix} = \begin{pmatrix} A_{n+1} - dA_{n} & cA_{n} \\ b_{r}A_{n} & dA_{n} - \mathcal{D}_{\mathcal{R}}A_{n-1} \end{pmatrix} \begin{pmatrix} x_{0}^{*} \\ y_{0}^{*} \end{pmatrix} + \frac{-1}{\mathcal{P}_{\mathcal{R}}(1)} \times$$

$$\begin{pmatrix} 1 - d & c \\ b_{r} & 1 - a_{r} \end{pmatrix} \begin{pmatrix} A_{n+1} - dA_{n} - 1 & cA_{n} \\ b_{r}A_{n} & dA_{n} - \mathcal{D}_{\mathcal{R}}A_{n-1} - 1 \end{pmatrix} \begin{pmatrix} h_{1} \\ h_{2} \end{pmatrix}$$

$$= \begin{pmatrix} (A_{n+1} - dA_{n})x_{0}^{*} + cA_{n}y_{0}^{*} - \Theta \\ b_{r}A_{n}x_{0}^{*} + (dA_{n} - \mathcal{D}_{\mathcal{R}}A_{n-1})y_{0}^{*} - \frac{b_{r}\left[A_{n}(1-\Gamma_{\mathcal{R}}) + A_{n+1}-1\right]h_{1} + b^{*}h_{2}}{\mathcal{P}_{\mathcal{R}}(1)} \end{pmatrix}$$

$$(64)$$

where $\Theta = \frac{\left((1-d)\left[A_{n+1}-dA_{n}-1\right]+cb_{r}A_{n}\right)h_{1}+c\left[A_{n}-\mathcal{D}_{\mathcal{R}}A_{n-1}-1\right]h_{2}}{\mathcal{P}_{\mathcal{R}}(1)}$. Investigating the signs of $\mathsf{L}(x_{n}^{*},y_{n}^{*})\cdot\mathsf{L}(x_{n+1}^{*},y_{n+1}^{*})$ we can then check for the existence of homoclinic intersections.

I SYSTEMS

I.1 DUFFING SYSTEM

The *Duffing system* is a classical nonlinear dynamical system that models a damped and periodically driven oscillator with a nonlinear restoring force. Originally introduced by Georg Duffing (Duffing, 1918), the system is described by a second-order differential equation of the form

$$\ddot{x} + \delta \dot{x} + \alpha x + \beta x^3 = \gamma \cos(\omega t),$$

where x represents the displacement, δ is a damping coefficient, α and β determine the linear and nonlinear stiffness respectively, and $\gamma\cos(\omega t)$ is an external periodic forcing term. The Duffing system exhibits rich dynamical behavior, including periodic, quasi-periodic, and chaotic responses, making it a prototypical example in the study of nonlinear and chaotic dynamics.

I.2 LORENZ63

The *Lorenz63 system* (Lorenz, 1963) is a continuous-time dynamical system originally developed to model atmospheric convection. It describes the evolution of three state variables governed by the

nonlinear differential equations

$$\frac{dx_1}{dt} = \sigma(x_2 - x_1),
\frac{dx_2}{dt} = x_1(\rho - x_3) - x_2,
\frac{dx_3}{dt} = x_1x_2 - \beta x_3,$$

where x_1, x_2 , and x_3 denote, respectively, the convection rate, horizontal temperature difference, and vertical temperature difference. The parameters σ , ρ , and β correspond to physical constants related to the Prandtl number, Rayleigh number, and system geometry.

For specific values, e.g. $\sigma=10$, $\rho=28$, and $\beta=\frac{8}{3}$, the system exhibits chaotic dynamics. These settings give rise to the well-known "butterfly attractor," a prime example of deterministic chaos in low-dimensional systems.

J EVALUATION METRICS

J.1 GEOMETRICAL MEASURE D_{STSP}

Given probability distributions p(x) (estimated from ground truth trajectories) and q(x) (estimated from model-generated trajectories), $D_{\rm stsp}$ is defined as the Kullback-Leibler (KL) divergence

$$D_{\text{stsp}} := D_{\text{KL}}(p(\boldsymbol{x}) \parallel q(\boldsymbol{x})) = \int_{\mathbf{x} \in \mathbb{R}^N} p(\boldsymbol{x}) \log \frac{p(\boldsymbol{x})}{q(\boldsymbol{x})} d\boldsymbol{x}.$$
 (58)

For low-dimensional observation spaces, p(x) and q(x) can be estimated using binning (Koppe et al., 2019a; Brenner et al., 2022). The KL divergence is approximated as

$$D_{\text{stsp}} = D_{\text{KL}}(\hat{p}(\boldsymbol{x}) \parallel \hat{q}(\boldsymbol{x})) \approx \sum_{k=1}^{K} \hat{p}_{k}(\boldsymbol{x}) \log \frac{\hat{p}_{k}(\boldsymbol{x})}{\hat{q}_{k}(\boldsymbol{x})}.$$
 (59)

Here, $K = m^N$ is the total number of bins, with m bins per dimension. $\hat{p}_k(x)$ and $\hat{q}_k(x)$ are the normalized counts in bin k for ground truth and model-generated orbits, respectively.

In high-dimensional settings, Gaussian Mixture Models (GMMs), placed along the trajectories, are used (Brenner et al., 2022). This results in approximate probability distributions

$$\hat{p}(\boldsymbol{x}) = \frac{1}{T'} \sum_{t=1}^{T'} \mathcal{N}(\boldsymbol{x}; \{\boldsymbol{x}_t\}, \Sigma), \quad \hat{q}(\boldsymbol{x}) = \frac{1}{T'} \sum_{t=1}^{T'} \mathcal{N}(\boldsymbol{x}; \{\hat{\boldsymbol{x}}_t\}, \Sigma)$$

where the covariance matrix is given by $\Sigma = \sigma^2 \mathbf{1}_{N \times N}$ (σ is a hyperparameter), and $\{x_t\}, \{\hat{x}_t\}$ are samples from the true and generated orbits of length T'.

Using the Monte Carlo approximation from hershey2007approximating, the KL divergence in this case is estimated by

$$D_{\text{stsp}} = D_{\text{KL}}(\hat{p}(\boldsymbol{x}) \parallel \hat{q}(\boldsymbol{x})) \approx \frac{1}{n} \sum_{i=1}^{n} \log \frac{\hat{p}(\boldsymbol{x}^{(i)})}{\hat{q}(\boldsymbol{x}^{(i)})}$$
(60)

with n Monte Carlo samples $\{x^{(i)}\}$ randomly drawn from the GMM $\hat{p}(x)$ that represents the real data distribution.

J.2 PREDICTION ERROR PE

The *n*-step prediction error is defined as the mean squared error between ground truth data $\{x_t\}$ and *n*-step ahead predictions of the model $\{\hat{x}_t\}$, i.e.

$$PE(n) = \frac{1}{N(T-n)} \sum_{t=1}^{T-n} \|\boldsymbol{x}_{t+n} - \hat{\boldsymbol{x}}_{t+n}\|_{2}^{2}.$$
 (65)

K ADDITIONAL FIGURES

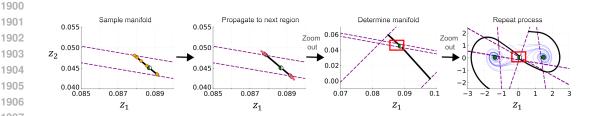


Figure 18: Illustration of the iterative procedure for computing stable manifolds with subregion boundaries of the shPLRNN ($M=2,\,H=10$) model in purple (dashed) Step 1: The stable manifold (black) is initialized using the stable eigenvector at the saddle point (half-green), and sample points (orange) are placed along it. Step 2: These points are propagated until they reach a new linear subregion, where the flow field is evaluated. Step 3: The updated manifold is given by the first principal component of points and the flow. Step 4: Repeating this process iteratively reconstructs the full global structure of the stable manifold (black), overlaid with the underlying GT flow field (blue)

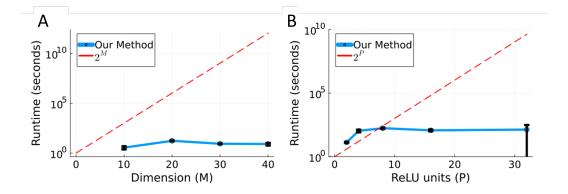


Figure 19: Algorithm runtime for determining the stable manifold of a saddle point averaged across 5 runs (error bars: standard deviation). A: For a *constant* number of linear subregions 2^P , Algo. 1's runtime hardly increases as a function of model size M for an ALRNN (P=2), confirming it is not significantly affected by the number of model parameters per se. B: For a *constant* model size M, Algo. 1's runtime increases much slower than 2^P for an ALRNN (M=40) when the manifold construction is restricted to the set of linear subregions explored by the data. All models were trained on the Duffing system for within-comparability, but we emphasize that the scaling may strongly depend on the system's actual dynamics and topological structure, such that general statements regarding scaling are therefore difficult. Runtime was determined on an Intel Core i5-1240P.