Multistability in input-driven recurrent neural networks

Andrea Ceni [1] Peter Ashwin [1] Lorenzo Livi [1,2] Claire Postlethwaite [3]

[1]University of Exeter, UK

[2]University of Manitoba, Canada

[3]University of Auckland, New Zealand

December 12th, 2019 – Exeter, UK

Outline

1. Recurrent neural networks
2. Open the Black Box
3. Beyond autonomous systems: the Echo idea
4. Local Point Attractors and the Echo Index
5. Current research

Feedforward/Recurrent Neural Networks

Recurrent neural networks (RNNs) are dynamical systems inspired by the brain

Feedforward/Recurrent Neural Networks

Recurrent neural networks (RNNs) are dynamical systems inspired by the brain

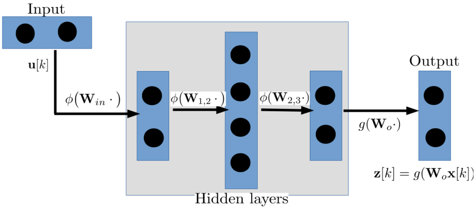

- Nonlinear (trainable) mapping from the input space to the output space:

$$\mathbf{z}[k] = \mathbf{F}(\mathbf{u}[k]) \qquad (1)$$

where $\mathbf{F}(\cdot) =$

$$g\left(\mathbf{W}_o\phi\left(\mathbf{W}_{2,3}\phi\left(\mathbf{W}_{1,2}\phi(\mathbf{W}_{in}\,\cdot)\right)\right)\right).$$

Feedforward/Recurrent Neural Networks

Recurrent neural networks (RNNs) are dynamical systems inspired by the brain

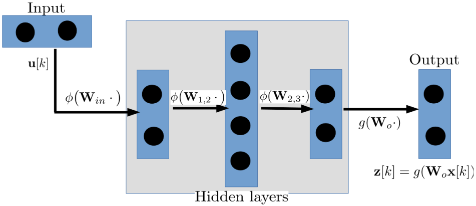

- Nonlinear (trainable) mapping from the input space to the output space:

$$\mathbf{z}[k] = \mathbf{F}(\mathbf{u}[k]) \qquad (1)$$

where $\mathbf{F}(\cdot) =$

$$g\left(\mathbf{W}_o\phi\left(\mathbf{W}_{2,3}\phi\left(\mathbf{W}_{1,2}\phi\left(\mathbf{W}_{in}\,\cdot\right)\right)\right)\right).$$

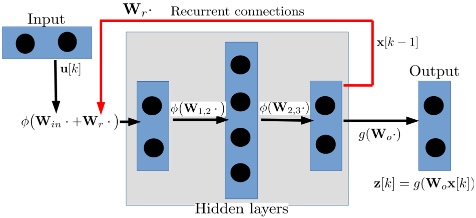

- Nonautonomous (trainable) dynamical system:

$$\mathbf{x}[k] = \mathbf{G}(\mathbf{x}[k-1], \mathbf{u}[k]) \qquad (2)$$

where $\mathbf{G}(\bullet, *) =$

$$\phi\left(\mathbf{W}_{2,3}\phi\left(\mathbf{W}_{1,2}\phi\left(\mathbf{W}_{in} * + \mathbf{W}_r \bullet\right)\right)\right).$$

This loop makes RNNs suitable to deal with temporal data.

Reservoir computing machines

Let's simplify: just 1 layer but very large (*reservoir*).

Reservoir computing machines

Let's simplify: just 1 layer but very large (*reservoir*).

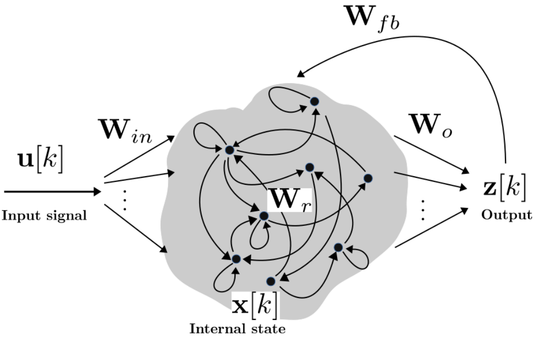

Figure 1: Representation of a RNN, with $\psi(x) = \mathbf{W}_o x$.

$$x[k] := \phi(\mathbf{W}_r x[k-1] + \mathbf{W}_i u[k] + \mathbf{W}_{fb} z[k-1]), \quad (3)$$

$$z[k] := \psi(x[k]). \quad (4)$$

RNNs as input-driven system

Recurrent neural networks (RNNs) are **nonautonomous** **nonlinear** **high-dimensional** discrete-time dynamical systems

$$x[k] := \alpha\phi(\underbrace{\mathbf{W}_i}_{N_r \times N_i} u[k] + \underbrace{\mathbf{W}_r}_{N_r \times N_r} x[k-1] + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} z[k-1]) + (1-\alpha)x[k-1], \qquad \mathbf{N_r} \text{ very large.}$$

$$z[k] := \psi(x[k]).$$

RNNs as input-driven system

Recurrent neural networks (RNNs) are **nonautonomous** **nonlinear** **high-dimensional** discrete-time dynamical systems

$$x[k] := \alpha\phi(\underbrace{\mathbf{W}_i}_{N_r \times N_i} u[k] + \underbrace{\mathbf{W}_r}_{N_r \times N_r} x[k-1] + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} z[k-1]) + (1-\alpha)x[k-1], \qquad \mathbf{N_r} \text{ very large.}$$

$$z[k] := \psi(x[k]).$$

$$\begin{aligned} \psi \in C^1(X, Y), &\ \text{(it depends on the specific task to solve,} \\ &\ \text{it reduces the effective dimensionality of the dynamics, } Y \subseteq X) \\ \phi \in C^1(\mathbb{R}, (-L, L)), &\ \text{monotonically increasing, } \phi' \text{ has unique max point in 0. (sigmoid)} \\ \alpha \in (0,1) &\ \text{tunes the time-scale of the RNN dynamics.} \end{aligned}$$

RNNs as input-driven system

Recurrent neural networks (RNNs) are **nonautonomous** **nonlinear** **high-dimensional** discrete-time dynamical systems

$$x[k] := \alpha\phi(\underbrace{\mathbf{W}_i}_{N_r \times N_i} u[k] + \underbrace{\mathbf{W}_r}_{N_r \times N_r} x[k-1] + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} z[k-1]) + (1-\alpha)x[k-1], \qquad \mathbf{N_r} \text{ very large.}$$

$$z[k] := \psi(x[k]).$$

$$\begin{aligned} \psi \in C^1(X, Y), &\ \text{(it depends on the specific task to solve,} \\ &\ \text{it reduces the effective dimensionality of the dynamics, } Y \subseteq X) \\ \phi \in C^1(\mathbb{R}, (-L, L)), &\ \text{monotonically increasing, } \phi' \text{ has unique max point in 0. (sigmoid)} \\ \alpha \in (0,1) &\ \text{tunes the time-scale of the RNN dynamics.} \end{aligned}$$

$$x[k] = G(u[k], x[k-1]), \qquad x[k] \in X \subset \mathbb{R}^{N_r} \text{ compact}, \quad u[k] \in U \subset \mathbb{R}^{N_i} \text{ compact} \qquad (5)$$

Training a RNN

$$x[k] = G(u[k], x[k-1]), \qquad G(u,x) := (1-\alpha)x + \alpha\phi(\underbrace{\mathbf{W}_r}_{N_r \times N_r} x + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} \psi(x) + \underbrace{\mathbf{W}_i}_{N_r \times N_i} u)$$

Training a RNN

$$x[k] = G(u[k], x[k-1]), \qquad G(u,x) := (1-\alpha)x + \alpha\phi(\underbrace{\mathbf{W}_r}_{N_r \times N_r} x + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} \psi(x) + \underbrace{\mathbf{W}_i}_{N_r \times N_i} u)$$

TRAINING = tune parameters $(\mathbf{W}_r, \mathbf{W}_{fb}, \mathbf{W}_i, \psi, \alpha)$ in order to minimise a Loss function.

- clever definition of a loss function of the parameters;
- clever choice of an optimisation algorithm.

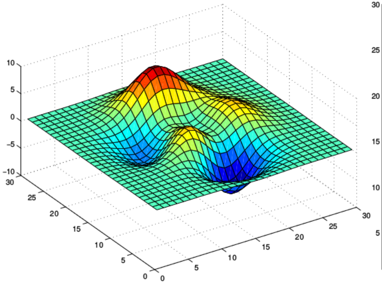

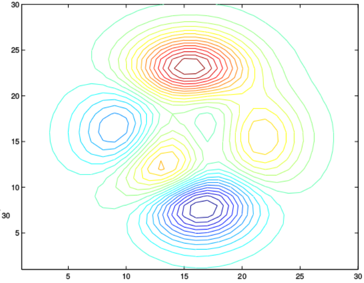

Opening the black box

Figure 2: Beer [3], Sussillo and Barak [6], Tiňo et al. [7]

A. Ceni, P. Ashwin, and L. Livi. Interpreting recurrent neural networks behaviour via excitable network attractors.
Cognitive Computation, pages 1–27, 2019.
doi: 10.1007/s12559-019-09634-2

Opening the black box

Figure 2: Beer [3], Sussillo and Barak [6], Tiňo et al. [7]

Is it possible to provide a mechanistic model of the functioning of a trained RNN?

A. Ceni, P. Ashwin, and L. Livi. Interpreting recurrent neural networks behaviour via excitable network attractors.
Cognitive Computation, pages 1–27, 2019.
doi: 10.1007/s12559-019-09634-2

Excitable Network Attractors (ENAs)

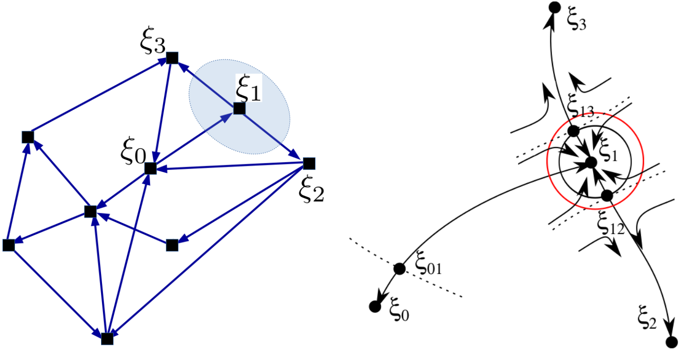

Figure 3: **Left:** directed graph representing an excitable network attractor. **Right:** excitable network attractor in phase space. See [1, 2, 8] for details.

Extracting ENA from a RNN trajectory

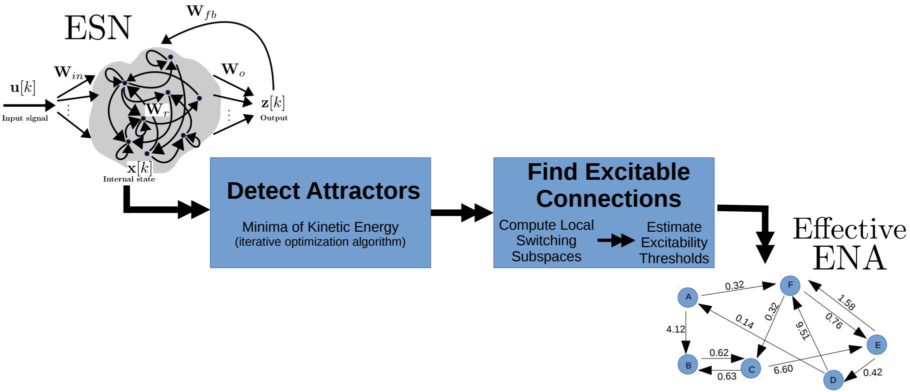

We showed that Excitable Network Attractors (ENAs) [1] can explain the functioning of RNNs in some tasks that require switching between attractors:

- useful to assess robustness to noise
- ENA models can be exploited to provide a mechanistic interpretation of errors occurring during the computation

Flip-Flop task: Learning memory states

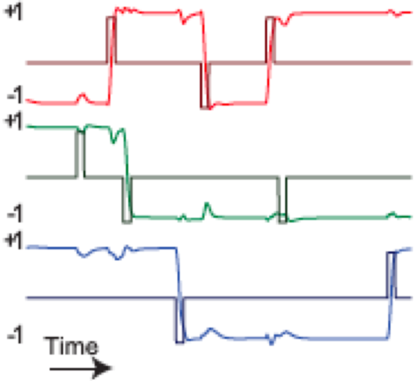

- 3D Input pulses —> 3D Output square wave
- The machine has to learn whether to switch a bit or not according to the current input provided and its current state

Mechanistic interpretation of errors

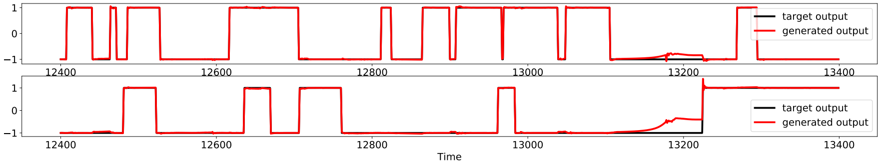

Mechanistic interpretation of errors

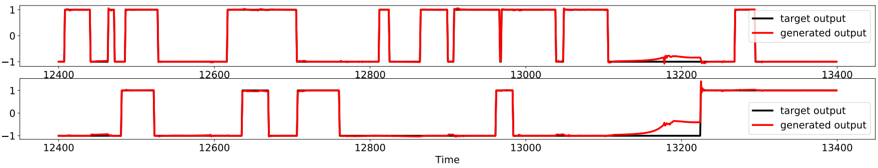

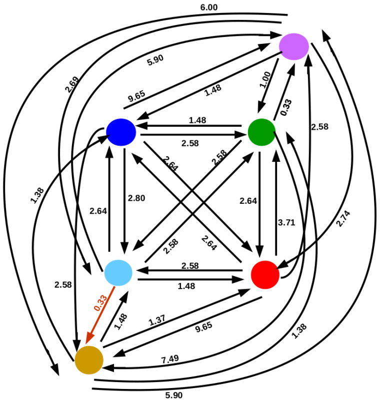

- Low overall mean squared error
- What went wrong there? ENA model can explain that behaviour!
- It learned 6 stable fixed points instead of the 4 attractors needed for the task
- Note also an excessively low excitability threshold (shown in red)

Going beyond autonomous dynamical systems ..

ENAs rely on autonomous dynamical systems theory: can handle only impulsive inputs. We want to go further and deal with RNNs driven by any type of inputs. This requires using concepts and notions from nonautonomous dynamical systems theory.

P. E. Kloeden and M. Rasmussen. *Nonautonomous Dynamical Systems.* Number 176. American Mathematical Society, 2011

The Echo State Property

Reservoir Computing paradigm: To function correctly some RNNs must possess the Echo State Property (ESP)

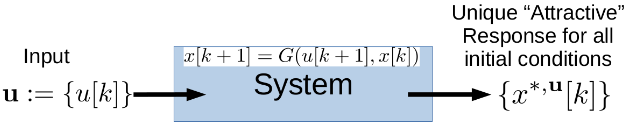

Figure 4: The system should transform the input into a unique response.

I. B. Yildiz, H. Jaeger, and S. J. Kiebel. Re-visiting the echo state property.
Neural Networks, 35:1–9, 2012.
doi: 10.1016/j.neunet.2012.07.005

The Echo State Property

Reservoir Computing paradigm: To function correctly some RNNs must possess the Echo State Property (ESP)

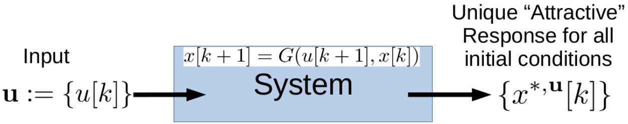

Figure 4: The system should transform the input into a unique response.

This implies that, given an input, the RNN will produce a unique response: it will "forget" any internal states and end up following a unique (possibly complex) attracting trajectory.

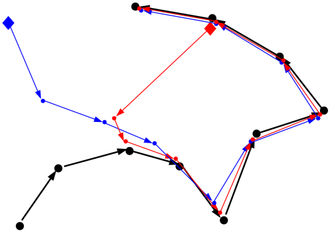

I. B. Yildiz, H. Jaeger, and S. J. Kiebel. Re-visiting the echo state property. *Neural Networks*, 35:1–9, 2012.
doi: 10.1016/j.neunet.2012.07.005

A generalisation of the ESP: multistable systems with inputs

Main Idea: to define the *Echo Index* $n = \mathcal{I}(\mathbf{u})$ of the system driven by input $\mathbf{u}$. This number quantifies the nonautonomous multistability under the action of the driving input

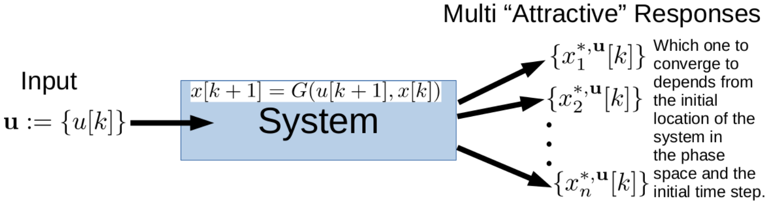

Example of Echo Index $n = 2$

$$u[k] \equiv A \;\Rightarrow\; x[k] = G(A, x[k-1])\,; \qquad u[k] \equiv B \;\Rightarrow\; x[k] = G(B, x[k-1])$$

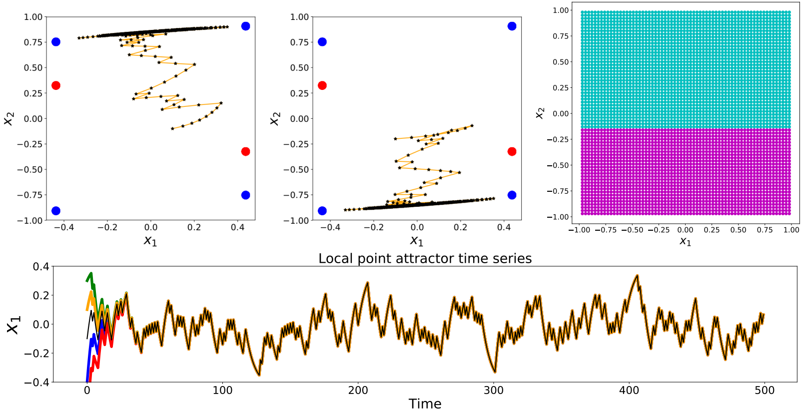

Figure 5: **Left:** Under the same realisation of input (switching between input A and B with $p_1 = p_2 = 0.5$), initial conditions converge to the upper local point attractor (left picture) or the bottom local point attractor (center picture). **Right:** Basins of attraction of the two local point attractors at time step $k = 0$. **Bottom:** Convergence of different initial conditions to one of the two local point attractors.

Cocycle formalism

Input space values (U, d_U) is a compact metric space.

Cocycle formalism

Input space values (U, d_U) is a compact metric space.

$$\Downarrow$$

Input signals space $\mathcal{U} := U^{\mathbb{Z}} = \{\mathbf{u} = \{u[k]\}_{k\in\mathbb{Z}} \ : \ u[k] \in U, \ \forall k \in \mathbb{Z}\}$ is a compact metric space with $d_{\mathcal{U}}(\mathbf{u}, \mathbf{v}) := \sum_{k\in\mathbb{Z}} \dfrac{d_U(u[k], v[k])}{2^{|k|}}$.

Cocycle formalism

Input space values (U, d_U) is a compact metric space.

$$\Downarrow$$

Input signals space $\mathcal{U} := U^{\mathbb{Z}} = \{\mathbf{u} = \{u[k]\}_{k\in\mathbb{Z}} \ : \ u[k] \in U, \ \forall k \in \mathbb{Z}\}$ is a compact metric space with $d_{\mathcal{U}}(\mathbf{u}, \mathbf{v}) := \sum_{k\in\mathbb{Z}} \dfrac{d_U(u[k], v[k])}{2^{|k|}}$.

Definition

The input-driven dynamics of (5) *is described through the cocycle mapping* $\Phi : \mathbb{Z}_0^+ \times \mathcal{U} \times \mathbf{X} \longrightarrow \mathbf{X}$ *as follows:*

$$\Phi(0, \mathbf{u}, x_0) := x_0, \qquad \forall x_0 \in \mathbf{X}, \forall \mathbf{u} \in \mathcal{U}, \tag{6}$$

$$\Phi(n, \mathbf{u}, x_0) := G(u[n], \Phi(n-1, \mathbf{u}, x_0)), \qquad \forall x_0 \in \mathbf{X}, \forall \mathbf{u} \in \mathcal{U}, \forall n \geq 1. \tag{7}$$

Local point attractor (nonautonomous fixed point)

Definition

Fix an input sequence $\mathbf{u} \in \mathcal{U}$. *We call a solution* $\{x[k]\}_{k\in\mathbb{Z}}$ *of the system* (5) *driven by input* $\mathbf{u}$ *a local point attractor if there exists a* $\mathbf{u}$*-positively invariant* nonautonomous set* $\{B[k]\}_{k\in\mathbb{Z}}$ *made by compact sets such that*

(i) $\lambda(B[k]) > 0$ *for all* $k \in \mathbb{Z}$*; and*

(ii) *For all sequences* $\{j_k\}_{k\in\mathbb{N}}$ *of integers we have*

$$\lim_{k\to\infty} h(\Phi(k, \sigma^{j_k}(\mathbf{u}), B[j_k]), x[j_k + k]) = 0.$$

If a $\mathbf{u}$*-positively invariant nonautonomous set* $\{B[k]\}_{k\in\mathbb{Z}}$ *made by compact sets satisfies (i)-(ii) we say it is uniformly attracted by the local point attractor* $\{x[k]\}_{k\in\mathbb{Z}}$*.*

Note that h denotes the Hausdorff semi-distance on $\mathbf{X}$, hence the limit above expresses a uniform (in space) convergence towards the local point attractor, because $h(\Phi(k, \sigma^{j_k}(\mathbf{v}), B[j_k]), x[j_k + k]) = \sup_{z\in B[j_k]} d_{\mathbf{X}}(\Phi(k, \sigma^{j_k}(\mathbf{v}), z), x[j_k + k])$.

*See the additional slide 22.

Definition of echo index

Definition

We say the system (5) *driven by input* $\mathbf{u} \in \mathcal{U}$ *has echo index n if there exist a number n of local point attractors* $\{x_1[k]\}_{k\in\mathbb{Z}}, \ldots, \{x_n[k]\}_{k\in\mathbb{Z}}$ *such that for all* $\eta > 0$ *there are nonautonomous sets* $\{B_1^\eta[k]\}_{k\in\mathbb{Z}}, \ldots, \{B_n^\eta[k]\}_{k\in\mathbb{Z}}$ *uniformly attracted by, respectively,* $\{x_1[k]\}_{k\in\mathbb{Z}}, \ldots, \{x_n[k]\}_{k\in\mathbb{Z}}$ *whose union at each time step get arbitrarily close to the full measure of the whole space, i.e.*

$$\lambda(X \setminus \bigcup_{i=1}^{n} B_i^\eta[k]) < \eta, \qquad \forall k \in \mathbb{Z}. \tag{8}$$

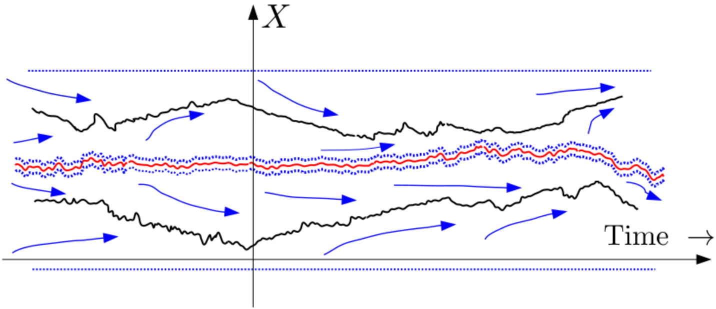

Existence of Local Point Attractors

The following theorem outlines a strategy for detecting local point attractors.

Theorem

Let be given an input sequence $\mathbf{v} \in \mathcal{U}$*. Suppose* $\{Q_n\}_{n\in\mathbb{Z}}$ *is a positively invariant nonautonomous set for input* $\mathbf{v}$ *made by nonempty compact sets. Suppose there exists a sequence of positive real numbers* $\{\mu_k\}_{k\in\mathbb{N}}$*, with* $\lim_{k\to\infty}\mu_k = 0$*, such that*

$$\forall n_0 \in \mathbb{Z},\ \forall x_0, y_0 \in Q_{n_0},\ \forall k > 0, \qquad d_X(\Phi(k, \sigma^{n_0}(\mathbf{v}), x_0), \Phi(k, \sigma^{n_0}(\mathbf{v}), y_0)) \leq \mu_k d_X(x_0, y_0). \tag{9}$$

Then, there exists a unique bi-infinite solution for input $\mathbf{v}$ *inside* $\{Q_n\}_{n\in\mathbb{Z}}$*. In addition, if* $\lambda(Q_s) > 0$*, for some* $s \in \mathbb{Z}$*, then such bi-infinite solution is a local point attractor.*

Existence of Local Point Attractors

The following theorem outlines a strategy for detecting local point attractors.

Theorem

Let be given an input sequence $\mathbf{v} \in \mathcal{U}$*. Suppose* $\{Q_n\}_{n\in\mathbb{Z}}$ *is a positively invariant nonautonomous set for input* $\mathbf{v}$ *made by nonempty compact sets. Suppose there exists a sequence of positive real numbers* $\{\mu_k\}_{k\in\mathbb{N}}$*, with* $\lim_{k\to\infty}\mu_k = 0$*, such that*

$$\forall n_0 \in \mathbb{Z},\ \forall x_0, y_0 \in Q_{n_0},\ \forall k > 0, \qquad d_X(\Phi(k, \sigma^{n_0}(\mathbf{v}), x_0), \Phi(k, \sigma^{n_0}(\mathbf{v}), y_0)) \leq \mu_k d_X(x_0, y_0). \tag{9}$$

Then, there exists a unique bi-infinite solution for input $\mathbf{v}$ *inside* $\{Q_n\}_{n\in\mathbb{Z}}$*. In addition, if* $\lambda(Q_s) > 0$*, for some* $s \in \mathbb{Z}$*, then such bi-infinite solution is a local point attractor.*

A simplified strategy is the following.

Theorem

Let μ *be a positive real such that* $\mu < 1$*. Suppose* Q_μ *be a* $\mathcal{U}$*-positively invariant nonempty compact set such that it is contained inside* $C(\mu) := \{x \in X \ : \ \sup_{u\in U} ||D_x G(u,x)|| \leq \mu\}$*, and suppose further that* Q_μ *is convex. Then, for all* $\mathbf{v} \in \mathcal{U}$ *the system (5) driven by the input sequence* $\mathbf{v}$ *admits a unique bi-infinite solution inside* Q_μ*. In addition, if* $\lambda(Q_\mu) > 0$ *then there is a unique local point attractor in* Q_μ*. Moreover, the convergence towards the local point attractor is ruled by* $\mu_k = \mu^k$*,* $k \in \mathbb{N}$*, i.e. strictly monotonic and linear.*

Existence of Echo Index n for small-amplitude input perturbations

Definition

Let be given a (autonomous) map $F : \mathbf{X} \longrightarrow \mathbf{X}$ and a fixed point $x^* \in \mathbf{X}$ for the map F, i.e. $F(x^*) = x^*$. We call x^* a Uniformly Attracting Stable Point (UASP) if there is an $0 < M < 1$ and $\delta > 0$, such that

$$d_{\mathbf{X}}(F(z), x^*) < M d_{\mathbf{X}}(z, x^*), \qquad \forall z \in B_\delta(x^*). \tag{10}$$

Theorem

Suppose there is an input value $\bar{u} \in \mathbb{R}^i$, such that the corresponding autonomous map $F(x) := G(\bar{u}, x)$ induced by the constant input value $\bar{u}$ possesses a UASP x^*.
Then, it is possible to find an $r > 0$ such that, for any input signal $\mathbf{v} \in B_r(\bar{u})^{\mathbb{Z}}$ there exists a unique local point attractor confined in a neighbourhood of x^*.
Furthermore, the convergence towards the local point attractor is strictly monotonic and linear.

This result implies that:

- Slightly perturbing an autonomous system possessing n UASPs will lead to a nonautonomous system where n local point attractors coexist.
- Generally, ESP cannot hold for small-amplitude input signals.

Large-amplitude inputs induce RNNs with echo index 1

$$G(u, x) := (1-\alpha)x + \alpha\phi(\underbrace{\mathbf{W}_r}_{N_r \times N_r} x + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} \psi(x) + \underbrace{\mathbf{W}_i}_{N_r \times N_i} u)$$

Large-amplitude inputs induce RNNs with echo index 1

$$G(u,x) := (1-\alpha)x + \alpha\phi(\underbrace{\mathbf{W}_r}_{N_r\times N_r} x + \underbrace{\mathbf{W}_{fb}}_{N_r\times N_o} \psi(x) + \underbrace{\mathbf{W}_i}_{N_r\times N_i} u)$$

$\psi \in C^1(X)$

$\phi \in C^1(\mathbb{R}, (-L, L))$, monotonically increasing, ϕ' has unique max point in 0. (sigmoid)

Large-amplitude inputs induce RNNs with echo index 1

$$G(u,x) := (1-\alpha)x + \alpha\phi(\underbrace{\mathbf{W}_r}_{N_r \times N_r} x + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} \psi(x) + \underbrace{\mathbf{W}_i}_{N_r \times N_i} u)$$

$\psi \in C^1(X)$

$\phi \in C^1(\mathbb{R}, (-L, L))$, monotonically increasing, ϕ' has unique max point in 0. (sigmoid)

$x[k] = G(u[k], x[k-1])$, $x[k] \in X \subset \mathbb{R}^{N_r}$ compact , $u[k] \in U \subset \mathbb{R}^{N_i}$ compact

Large-amplitude inputs induce RNNs with echo index 1

$$G(u,x) := (1-\alpha)x + \alpha\phi(\underbrace{\mathbf{W}_r}_{N_r \times N_r} x + \underbrace{\mathbf{W}_{fb}}_{N_r \times N_o} \psi(x) + \underbrace{\mathbf{W}_i}_{N_r \times N_i} u)$$

$\psi \in C^1(X)$

$\phi \in C^1(\mathbb{R}, (-L, L))$, monotonically increasing, ϕ' has unique max point in 0. (sigmoid)

$$x[k] = G(u[k], x[k-1]), \qquad x[k] \in X \subset \mathbb{R}^{N_r} \text{ compact }, \quad u[k] \in U \subset \mathbb{R}^{N_i} \text{ compact}$$

Theorem

For all $\varepsilon > 0$, yielding the closed set $P_\varepsilon := \mathbb{R}^{N_i} \setminus \bigcup_{j=1}^{N_r} B_\varepsilon(H_j)$, where $H_j := \{u \in \mathbb{R}^{N_i} \ : \ (W_{in})_{(j)} \cdot u = 0\}$, there exists a $R_\varepsilon > 0$ such that for all compact sets K containing $B_{R_\varepsilon}(0)$ and for all input sequences $\mathbf{u} \in \left(K \setminus \left(P_\varepsilon \cup B_{R_\varepsilon}(0)\right)\right)^{\mathbb{Z}}$ there exists a unique local point attractor for the RNN (3) driven by $\mathbf{u}$.

Research directions

- To exploit the generalised ESP in:
    - Classification tasks
    - Counting problems
    - Diagnostic of errors
- To devise algorithms to compute local point attractors and related properties from data (reverse engineer RNNs driven by non trivial inputs)
    - Nonautonomous (i.e. time-dependent) excitable network attractors
- To look at the results of different training mechanisms for RNNs on the number, type, and properties of attractors in the underlying autonomous dynamics and how these affect the behaviour of the nonautonomous system

Thanks

Chaires de recherche du Canada
Canada Research Chairs
Canada

NVIDIA

EPSRC

References I

[1] P. Ashwin and C. Postlethwaite. Designing heteroclinic and excitable networks in phase space using two populations of coupled cells. *Journal of Nonlinear Science*, 26(2):345–364, 2016. doi: 10.1007/s00332-015-9277-2.

[2] P. Ashwin and C. Postlethwaite. Sensitive finite state computations using a distributed network with a noisy network attractor. *IEEE Transactions on Neural Networks and Learning Systems*, 29(12):5847–5858, Dec. 2018. doi: 10.1109/TNNLS.2018.2813404.

[3] R. D. Beer. Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18(12):3009–3051, 2006. doi: 10.1162/neco.2006.18.12.3009.

[4] A. Ceni, P. Ashwin, and L. Livi. Interpreting recurrent neural networks behaviour via excitable network attractors. *Cognitive Computation*, pages 1–27, 2019. doi: 10.1007/s12559-019-09634-2.

[5] P. E. Kloeden and M. Rasmussen. *Nonautonomous Dynamical Systems*. Number 176. American Mathematical Society, 2011.

[6] D. Sussillo and O. Barak. Opening the black box: Low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Computation*, 25(3):626–649, 2013. doi: 10.1162/NECO_a_00409.

[7] P. Tiňo, B. G. Horne, and C. L. Giles. Attractive periodic sets in discrete-time recurrent networks (with emphasis on fixed-point stability and bifurcations in two-neuron networks). *Neural Computation*, 13(6):1379–1414, 2001. doi: 10.1162/08997660152002898.

References II

[8] O. Weinberger and P. Ashwin. From coupled networks of systems to networks of states in phase space. *Discrete & Continuous Dynamical Systems - B*, 23:2043, 2018. ISSN 1531-3492. doi: 10.3934/dcdsb.2018193.

[9] I. B. Yildiz, H. Jaeger, and S. J. Kiebel. Re-visiting the echo state property. *Neural Networks*, 35:1–9, 2012. doi: 10.1016/j.neunet.2012.07.005.

Positively invariant sets

Consider a nonautonomous system defined by a cocycle mapping $\Phi : \mathbb{Z}_0^+ \times \mathcal{U} \times \mathbf{X} \longrightarrow \mathbf{X}$ as in Definition 1, and let us fix an input sequence $\mathbf{u} \in \mathcal{U}$.

Definition

A family of nonempty compact sets $\mathbf{B} = \{B_n\}_{n\in\mathbb{Z}}$ is called a positively invariant nonautonomous set for input $\mathbf{u}$ (or simply $\mathbf{u}$-positively invariant) if

$$\Phi(s, \sigma^m(\mathbf{u}), B_m) \subseteq B_{s+m}.$$

for all $m \in \mathbb{Z}$ and $s \in \mathbb{Z}_0^+$.

Definition

A nonempty compact subset $B \subseteq X$ is called positively invariant for a family of input sequences $\mathcal{V} \subseteq \mathcal{U}$ (or $\mathcal{V}$-positively invariant) if for all $k \in \mathbb{Z}_0^+$ it holds that

$$\Phi(k, \mathbf{u}, B) \subseteq B, \qquad \forall \mathbf{u} \in \mathcal{V}.$$

Boundary solutions

All the initial conditions taken on the basin boundary converge to an unstable local point attractor representing a sort of nonautonomous saddle.

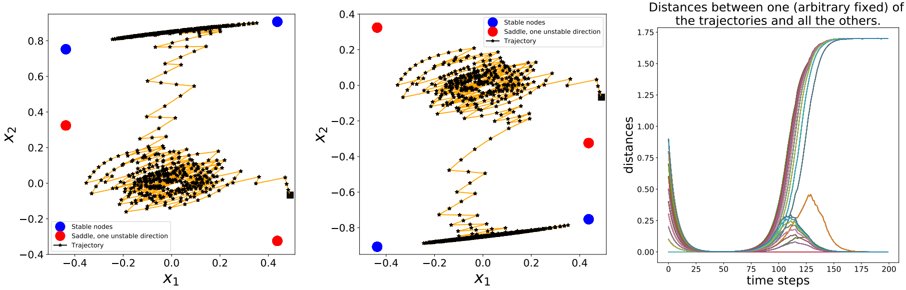

Figure 6: **Left:** initial condition is (0.49, −0.066761137519131756). The trajectory takes 354 steps before to cross the stable manifold of the left saddle. **Center:** initial condition is (0.49, −0.066761137519131757). **Right:** 100 initial conditions have been chosen around the separatrix. At each time step has been computed and plotted the distance between a (randomly chosen) reference trajectory and all the other trajectories.