Learning Simple Computations in Dynamical Systems by Example

Jason Z. Kim (jinsu1@seas.upenn.edu)
Department of Bioengineering, University of Pennsylvania
Philadelphia, PA, 19104

Danielle S. Bassett (dsb@seas.upenn.edu)
Departments of Bioengineering, Physics & Astronomy, Electrical & Systems Engineering, and Neurology, University of Pennsylvania
Philadelphia, PA, 19104

Abstract

The ability to store, represent, and manipulate information is a crucial element of processing systems. While computers are carefully engineered and manufactured to perform mathematical operations on data, neurobiological systems robustly provide similar functions with substantially more variability in physical brain connectivity across species and development. In addition, neural systems can recognize largely unstructured patterns of sensory inputs that cannot always be nicely represented as discrete static pieces of information. Here we present a dynamical system that can represent chaotic attractors, and learn a simple translation operation by example. Specifically, we train a sparse and randomly connected reservoir computer system to evolve along two translated chaotic Lorenz attractors with different initial conditions embedded in 3 dimensions. During training, we apply a fourth input that takes a unique constant value for each attractor. We demonstrate that by driving the trained reservoir with new values of this fourth input, the reservoir is able to extrapolate the translation of the Lorenz attractor. Together, our results provide a simple but powerful mechanism by which a general dynamical system can learn to manipulate internal representations of complex information.

Keywords: Dynamical Systems; Neural Network; Information; Reservoir; Recall; Memory;

Learning Chaotic Attractors Using Reservoirs

Reservoir computers are a class of dynamical recurrent neural networks that have been successfully used to classify and predict time-series data (Jaeger, 2010). The reservoir itself is a high-dimensional non-linear dynamical system with sparse and random coupling between states. Typically, the reservoir states are driven by a low-dimensional time-series signal, and measured to classify or reconstruct some desired output.

By training the outputs to approximate the inputs, the reservoir outputs can be recursively fed back as inputs, causing the reservoir to autonomously evolve according to the input signal. This closed-loop architecture has been shown to accurately reconstruct the dynamics of even chaotic attractors (Pathak, Hunt, Girvan, Lu, & Ott, 2018), and is a powerful demonstration of how a simple but large non-linear dynamical system can easily learn complex attractors. Here, we demonstrate that a simple external signal can be used to bias the reservoir dynamics when learning two separate attractors. After closing the loop, the constant external signal can be modulated outside of the training values to control the reservoir evolution along further translations.

Mathematical Framework

We consider an n-dimensional reservoir with states $\boldsymbol{r} \in \mathbb{R}^n$, coupled according to adjacency matrix $M \in \mathbb{R}^{n \times n}$. We also consider an input signal $\boldsymbol{s} \in \mathbb{R}^k$ coupled to the reservoir according to adjacency matrix $W_{\text{in}} \in \mathbb{R}^{n \times k}$, and the output signal $\hat{\boldsymbol{s}} \in \mathbb{R}^k$ as a linear combination of reservoir states according to $W_{\text{out}} \in \mathbb{R}^{k \times n}$. We consider the constant external control signal $c \in \mathbb{R}$ coupled to the reservoir by $\boldsymbol{v} \in \mathbb{R}^n$. Finally, we include a bias term $\boldsymbol{b} \in \mathbb{R}^n$. Then the open-loop reservoir evolves as

$$\frac{d\boldsymbol{r}(t)}{dt} = \gamma(-\boldsymbol{r}(t) + \tanh(M\boldsymbol{r}(t) + \sigma W_{\text{in}}\boldsymbol{s}(t)) + c(t)\boldsymbol{v} + \boldsymbol{b}).$$

The inputs are state trajectories of the 3-dimensional Lorenz dynamical system

$$\begin{aligned}\frac{dx}{dt} &= \sigma(y-x)\\ \frac{dy}{dt} &= x(\rho - z) - y\\ \frac{dz}{dt} &= xy - \beta z,\end{aligned}$$

with states $\boldsymbol{s} = (x,y,z)$, such that $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$ are the two attractors with the same parameters (σ, ρ, β) with different initial conditions, where $\boldsymbol{s_2}$ is translated.

The outputs are given as $\hat{\boldsymbol{s}} = W_{\text{out}}\boldsymbol{r}$, and we train the output matrix to minimize the quantity $\|\hat{\boldsymbol{s}} - \boldsymbol{s}\|_2^2$ at every time point. After training, we feed the predicted output back into the input to yield the closed-loop dynamics

$$\frac{d\boldsymbol{r}(t)}{dt} = \gamma(-\boldsymbol{r}(t) + \tanh(M^*\boldsymbol{s}(t) + c(t)\boldsymbol{v} + \boldsymbol{b})),$$

where $M^* = M^* \sigma W_{\text{in}} W_{\text{out}}$.

Numerical Simulations

The chaotic attractors, and both the open- and closed-loop reservoirs, were simulated using a 4th order Runge-Kutta approximation with a time step of $\Delta t = 0.001$. Both attractors $\boldsymbol{s}_1, \boldsymbol{s}_2$ were simulated for $T = 200$, and were used sequentially to drive the reservoir. For each attractor, the first $T_w = 50$ measurements were discarded to remove transient states, and the last $T_t = 150$ measurements were used for training W_{out}.

The reservoir was started at an initial state of $\boldsymbol{r} = 0$ with $n = 2000$ nodes. The coupling matrix elements were drawn from a uniform random distribution $W \in [-1,1]^{n\times n}$ with 0.05 density. Each row of W_{in} had only one element which was also drawn from a uniform random distribution $[-1,1]$. The bias vector $\boldsymbol{b}$ was drawn uniformly from $[-1,1]^n$. Attractor $\boldsymbol{s_2}$ was translated 10 units about the z-axis.

Results

Using the typical Lorenz parameters $(\sigma,\rho,\beta) = (10,28,8/3)$, we generate $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$, where $\boldsymbol{s_2}$ is shifted by 10 units along the positive z-axis. The reservoir used an input of $c = 5$ when training of $\boldsymbol{s_1}$, and $c = 10$ when training on $\boldsymbol{s_2}$. After training W_{out}, the closed-loop reservoir was evolved using input $c = 15$. The training data $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$ are shown as the bottom and middle blue invariant manifolds, and the output $\hat{\boldsymbol{s}} = W_{\text{out}}\boldsymbol{r}$ from the extrapolated forced input using $c = 15$ is shown as the green manifold (Fig. 1).

As can be seen, the reservoir is able to endogenously learn a translation operation on complex chaotic manifolds from examples, and execute this computation in the native reservoir representation with a simple modulation. We note that the reservoir initially begins by evolving on manifold $\boldsymbol{s_2}$ with $c = 10$. Then, by switching the control signal to $c = 15$, the reservoir evolves smoothly onto the translated green curve.

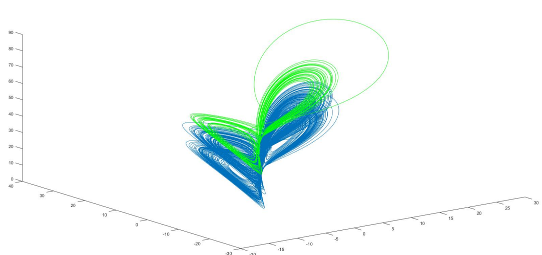

Figure 1: Plots of invariant manifolds from $\boldsymbol{s_1}$ and $\boldsymbol{s_2}$ as the non-translated and translated Lorenz attractors as the bottom two blue curves (trained using $c = 5$ and $c = 10$, respectively), and the transformed reservoir outputs $\hat{\boldsymbol{s}} = W_{\text{out}}\boldsymbol{r}$ using a new control signal $c = 15$. As expected, the reservoir output has a continued positive z-translation. The reservoir initially starts on the $\boldsymbol{s_2}$ manifold at $c = 10$, then transitions smoothly via a change in control signal to $c = 15$, as shown by the single large green loop.

Acknowledgments

J.Z.K. acknowledges support from National Institutes of Health T32-EB020087, PD: Felix W. Wehrli, and the National Science Foundation Graduate Research Fellowship No. DGE-1321851. D.S.B. acknowledges support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the ISI Foundation, the Paul Allen Foundation, the Army Research Laboratory (W911NF-10-2-0022), the Army Research Office (Bassett-W911NF-14-1-0679, Grafton-W911NF-16-1-0474, DCIST- W911NF-17-2-0181), the Office of Naval Research, the National Institute of Mental Health (2-R01-DC-009209-11, R01 MH112847, R01-MH107235, R21-M MH-106799), the National Institute of Child Health and Human Development (1R01HD086888-01), National Institute of Neurological Disorders and Stroke (R01 NS099348), and the National Science Foundation (BCS-1441502, BCS-1430087, NSF PHY-1554488 and BCS-1631550). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

References

Jaeger, H. (2010, dec). The echo state approach to analysing and training recurrent neural networks with an Erratum note. *GMD Report*, *1*(148), 1–47.

Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018, jan). Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach. *Physical Review Letters*, *120*(2), 024102. doi: 10.1103/PhysRevLett.120.024102