


Learning chaotic dynamics via reservoir computing

Anton Pershin^{1, 2}

anton.pershin@physics.ox.ac.uk

¹Atmospheric, Oceanic and Planetary Physics, University of Oxford

²School of Mathematics, University of Leeds



Machine Learning Study Group, University of Leeds, Leeds, UK

10 March 2021

Origins


Machine-learning problem

Wish to learn the functional relationship between

→ Given input: $\mathbf{u}(n) \in \mathbb{R}^{N_u}$

→ Desired output: $\mathbf{y}_{target}(n) \in \mathbb{R}^{N_y}$

Dataset: $\{(\mathbf{u}(n), \mathbf{y}_{target}(n))\}_{n=1}^T$.




Time-series prediction

Suppose that n denotes the discrete time.

→ Feedforward Neural Networks: $\mathbf{x}(n) = \phi(\mathbf{u}(n), \mathbf{u}(n-1), \dots)$

→ Recurrent Neural Networks: $\mathbf{x}(n) = \phi(\mathbf{x}(n-1), \mathbf{u}(n))$




Time-series prediction

Suppose that n denotes the discrete time.

→ Feedforward Neural Networks: $\mathbf{x}(n) = \phi(\mathbf{u}(n), \mathbf{u}(n-1), \dots)$

→ Recurrent Neural Networks: $\mathbf{x}(n) = \phi(\mathbf{x}(n-1), \mathbf{u}(n))$




Recurrent neural networks (RNN)

Simple form of RNN:

$$\begin{aligned}\mathbf{x}(n) &= f(\mathbf{W}_{in}\mathbf{u}(n) + \mathbf{W}\mathbf{x}(n-1)), \\ \mathbf{y}(n) &= \mathbf{W}_{out}\mathbf{x}(n).\end{aligned}$$

where $f(\cdot)$ is usually $\tanh(\cdot)$, $\mathbf{W}_{in} \in \mathbb{R}^{N_x \times N_u}$, $\mathbf{W} \in \mathbb{R}^{N_x \times N_x}$, $\mathbf{W}_{out} \in \mathbb{R}^{N_y \times N_x}$.
RNNs can be viewed as universal approximators of dynamical systems¹



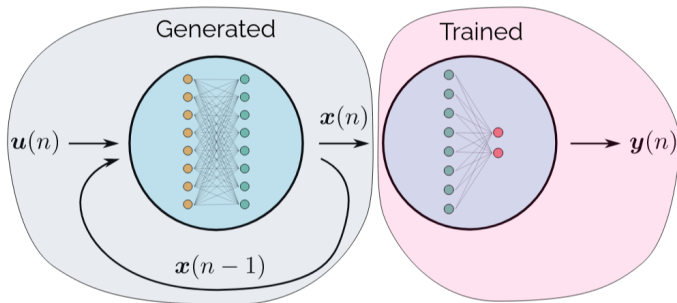
¹Funahashi and Nakamura, *Neural Networks* **6**, 801–806 (1993)

Reservoir computing

Reservoir computing denotes the approach where the recurrent part of RNN is generated/trained separately from the recurrence-free readout.

Reservoir-computing methods:

- Echo State Networks
- Liquid State Machines
- Backpropagation-Decorrelation training
- Temporal Recurrent Networks

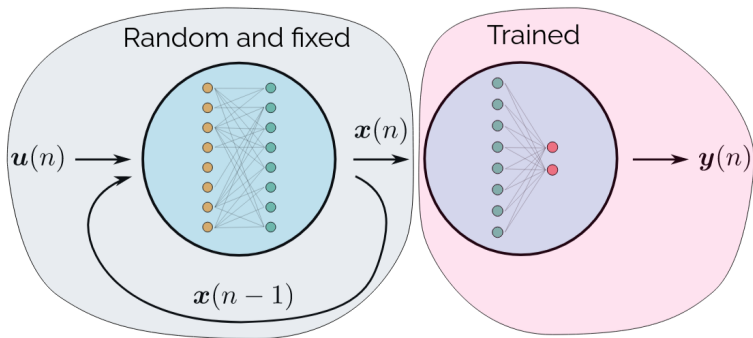


Echo State Networks

Echo State Network

$$\mathbf{x}(n) = f(\mathbf{W}_{in}\mathbf{u}(n) + \mathbf{W}\mathbf{x}(n-1)),$$
$$\mathbf{y}(n) = \mathbf{W}_{out}\mathbf{x}(n).$$

where \mathbf{W}_{in} and \mathbf{W} are random sparse matrices and \mathbf{W}_{out} is to be trained.




Training

Minimisation of the residual sum of squares (RSS):

$$\min_{\mathbf{W}_{out}} \sum_{n=1}^T \|\mathbf{y}(n) - \mathbf{y}_{target}(n)\|_2^2 = \min_{\mathbf{W}_{out}} \sum_{n=1}^T \|\mathbf{W}_{out} \mathbf{x}(n) - \mathbf{u}(n+1)\|_2^2,$$

where $\mathbf{x}(n) = f(\mathbf{W}_{in} \mathbf{u}(n) + \mathbf{W} \mathbf{x}(n-1))$.



Training

Solution via the normal equation:

$$\mathbf{W}_{out}^T = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y},$$

where matrices \mathbf{X} and \mathbf{Y} are given by

$$\mathbf{X} = \begin{bmatrix} \text{---} & \mathbf{x}(1) & \text{---} \\ \text{---} & \mathbf{x}(2) & \text{---} \\ & \dots & \\ \text{---} & \mathbf{x}(T) & \text{---} \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} \text{---} & \mathbf{u}(1) & \text{---} \\ \text{---} & \mathbf{u}(2) & \text{---} \\ & \dots & \\ \text{---} & \mathbf{u}(T) & \text{---} \end{bmatrix}$$

Here internal states $\mathbf{x}(n)$ are obtained via solving the explicit recurrent equation:

$$\mathbf{x}(n) = f(\mathbf{W}_{in} \mathbf{u}(n) + \mathbf{W} \mathbf{x}(n-1)),$$


provided an initial condition $\mathbf{x}(0) = \mathbf{0}$.

Prediction

Generative mode:

$$\begin{aligned}\mathbf{x}(n) &= f(\mathbf{W}_{in}\mathbf{y}(n-1) + \mathbf{W}\mathbf{x}(n-1)), \\ \mathbf{y}(n) &= \mathbf{W}_{out}\mathbf{x}(n).\end{aligned}$$

We still need to specify the initial condition $\mathbf{x}(0)$.




Prediction

Generative mode:

$$\begin{aligned}\mathbf{x}(n) &= f(\mathbf{W}_{in}\mathbf{y}(n-1) + \mathbf{W}\mathbf{x}(n-1)), \\ \mathbf{y}(n) &= \mathbf{W}_{out}\mathbf{x}(n).\end{aligned}$$

We still need to specify the initial condition $\mathbf{x}(0)$.

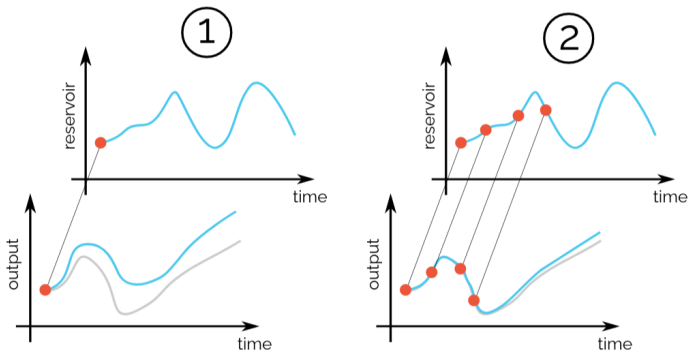


Prediction

Generative mode:

$$\begin{aligned}\mathbf{x}(n) &= f(\mathbf{W}_{in}\mathbf{y}(n-1) + \mathbf{W}\mathbf{x}(n-1)), \\ \mathbf{y}(n) &= \mathbf{W}_{out}\mathbf{x}(n).\end{aligned}$$

We still need to specify the initial condition $\mathbf{x}(0)$.



Improvements

Echo State Property

This condition is guaranteed when $\alpha = \rho(\mathbf{W}) < 1$, so once \mathbf{W} is randomly generated, it needs to be re-scaled:

$$\mathbf{W} := \frac{1}{\alpha} \mathbf{W}.$$

Tikhonov regularisation

Helps improve stability and overfitting.

$$\min_{\mathbf{W}_{out}} \sum_{n=1}^T \|\mathbf{W}_{out} \mathbf{x}(n) - \mathbf{u}(n+1)\|_2^2 + \beta \|\mathbf{W}_{out}\|^2$$
$$\implies \mathbf{W}_{out}^T = (\mathbf{X}^T \mathbf{X} + \beta \mathbf{I})^{-1} \mathbf{X}^T \mathbf{Y}.$$

Noise immunisation

Helps improve stability and overfitting.

$$\mathbf{x}(n) = f(\mathbf{W}_{in} \mathbf{u}(n) + \mathbf{W} \mathbf{x}(n-1)) + \xi Z,$$

where $Z \sim \text{uniform}(-1, 1)$ and ξ is the noise strength.

Leaky integrator neurons


Echo State Networks can be viewed as Euler discretisation of the leaky-integrator-type ODE:

$$\frac{d\mathbf{x}}{dt} = -\mathbf{x} + f(\mathbf{W}_{in}\mathbf{u} + \mathbf{W}\mathbf{x})$$
$$\implies \mathbf{x}(n) = (1 - \Delta t)\mathbf{x}(n-1) + \Delta t f(\mathbf{W}_{in}\mathbf{u}(n) + \mathbf{W}\mathbf{x}(n-1)),$$

where $\Delta t \in [0; 1]$.

Input scaling

Scaling of the input weights \mathbf{W}_{in} and shifting the input $\mathbf{u}(n)$ helps control the "amount of non-linearity".







Applications

Lorenz 63 model

Long-term prediction

Lyapunov exponents

Correlation dimension



Haluszczynski and R ath, *Chaos* **29**, 103143 (2019)

Lorenz 63 model

Long-term prediction

Lyapunov exponents

Correlation dimension

TABLE II. Three largest Lyapunov exponents $\Lambda_1 \geq \Lambda_2 \geq \Lambda_3$ for the Lorenz system [Eq. (5)], and for the reservoir set up in the configuration of Fig. 1(b) for R1 and R2. Since the reservoir system that we employ is a discrete time system, while the Lorenz system is a continuous system, for the purpose of comparison, Λ_1 , Λ_2 , and Λ_3 are taken to be per unit time; that is, their reservoir values (columns 2 and 3) are equal to the reservoir Lyapunov exponents calculated on a per iterate basis divided by Δt .

	Actual Lorenz system	R1 system	R2 system
Λ_1	0.91	0.90	0.01
Λ_2	0.00	0.00	-0.1
Λ_3	-14.6	-10.5	-9.9

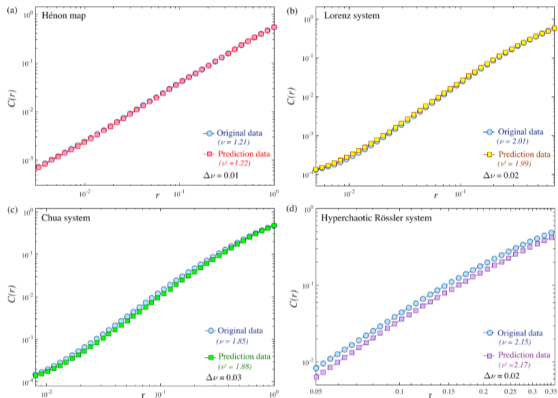
Pathak et al., *Chaos* **27**, 121102 (2017)

Lorenz 63 model

Long-term prediction

Lyapunov exponents

Correlation dimension



Chen et al., *Phys. Rev. E* **102**, 033314 (2020)


Kuramoto–Sivashinsky equation

Short-term prediction

Failure

Improved prediction

Lyapunov exponents



Pathak et al., *Chaos* **27**, 121102 (2017)

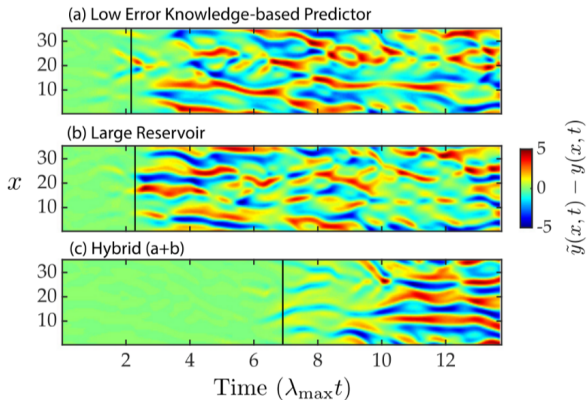
Kuramoto–Sivashinsky equation

Short-term prediction

Failure

Improved prediction

Lyapunov exponents



Pathak et al., *Chaos* **28**, 041101 (2018)



Kuramoto–Sivashinsky equation

Short-term prediction

Failure

Improved prediction

Lyapunov exponents




Pathak et al., *Chaos* **27**, 121102 (2017)

Rayleigh–Bénard convection

Long-term prediction

Snapshots

Statistics




Pandey et al., *Phys. Rev. Fluids* **5**, 113506 (2020)

Rayleigh–Bénard convection

Long-term prediction

Snapshots

Statistics




Pandey et al., *Phys. Rev. Fluids* **5**, 113506 (2020)

Comparative studies: Lorenz 96

RC vs. ANN vs. LSTM


RC vs. GRU vs. LSTM vs. Unit




Chattopadhyay et al., *Nonlin. Processes Geophys.* **27**, 373–389 (2020)

Comparative studies: Lorenz 96

RC vs. ANN vs. LSTM




(a) Reduced order observable ($d_o = 35$), $F = 8$




(c) Reduced order observable ($d_o = 35$), $F = 10$

RC vs. GRU vs. LSTM vs. Unit



(b) Full state ($d_o = 40$), $F = 8$



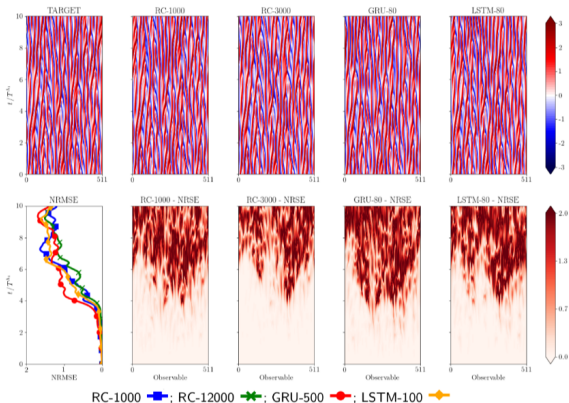
(d) Full state ($d_o = 40$), $F = 10$

Vlachas et al., *arXiv:1910.05266* (2019)

Comparative studies: K.-S. equation

RMS error


Performance



Vlachas et al., *arXiv:1910.05266* (2019)


Comparative studies: K.-S. equation

RMS error




(a) Valid prediction time in the test dataset

Performance



(b) Training time



(c) Average RAM memory requirement

Vlachas et al., *arXiv:1910.05266* (2019)

How can we use reservoir computing?

Building a dynamical system from observations only

Acceleration of numerical simulations

Super-parameterisation

Reduced-order modelling