Predicting hidden structure in dynamical systems

Daniel J. Gauthier[1] and Ingo Fischer[2]

[1]*Department of Physics, The Ohio State University, Columbus, OH, USA.*

[2]*Instituto de Física Interdisciplinar y Sistemas Complejos, IFISC (CSIC-UIB), Campus Universitat de les Illes Balears, Palma de Mallorca, Spain.*

Reservoir computers (RCs) are recurrent artificial neural networks that can solve a variety of machine learning tasks. They are attractive because they require comparably small training dataset sizes and allow for fast and straightforward training using linear regression(1). This training simplicity is the result of the reservoir architecture: the weights between the input layer and the core recurrent network (the 'reservoir') and the weights within the reservoir itself are kept fixed and are usually randomly assigned. Only the output layer is trained. Moreover, many different reservoir topologies, even very simple ones, have proven to perform well (2), facilitating hardware implementations of RCs (3,4). A particularly challenging task in time-series prediction is forecasting a chaotic signal, often used as a benchmark because of the sensitive dependence on errors and well-established prediction-horizon scale. Writing in Nature Machine Intelligence, Kim et al.(5) report a breakthrough in the field by demonstrating that a RC can predict response patterns generated by a dynamical system that are distinctly different from those it has seen during training. In other words, they teach the RC to modify its internal representation of the system.

Time-series prediction has been applied to various dynamical systems including a periodically forced pendulum, chaotic laser emission, and even the Earth's weather system. Mathematically, they are usually modelled using nonlinear ordinary or partial differential equations. In a typical setting for learning, sensors monitoring a physical system generate time-series data (possibly multidimensional), which are used to train a RC. Once trained for time-series prediction, the output of a RC can be fed back to its input (the so-called closed-loop configuration), resulting in an autonomous system that replicates the dynamics of the original system. Previous research demonstrated that the RC learns the underlying dynamical system and thus can predict subtle behaviours not seen explicitly during training (6). That is, it generates a model of the dynamical system. To date, research has mainly been confined to the case where the parameters of the system are kept fixed, although Klos et al.(7) have already realized a form of dynamical learning.

Once a RC is pre-trained and has learned some dynamics for particular control parameters via feedback of an error signal, it can predict unseen dynamics for different inputs or intermediate control parameters7; that is, their system has learned features of the dynamical system and its parameter dependence. But what happens when a control parameter is tuned beyond the range of its training values (the hand controlling the knob in Fig. 1)? Do prediction and replication of the dynamics fail?

Kim et al. address this question using an ingenious approach. They train the RC with an additional auxiliary signal that is related to an adjustable parameter (the parameter 'a' in Fig. 1), which is set to discrete values during training. When testing in the closed-loop configuration and while changing the parameter continuously and slowly in comparison to the characteristic timescale of the dynamics, they find that the RC can 'interpolate' between the behaviours on which it was trained and even extrapolate for parameters it has not seen. They also demonstrate changing the representation of the system by translating and rotating its attractor in phase space, where the attractor is the set of points corresponding to the time-course of the system variables after transients have died out. Interpolating time-series data using artificial neural networks has been demonstrated previously (8) but Kim et al. go well beyond interpolation by predicting the dynamics even across bifurcation points, as illustrated in Fig. 1

A bifurcation indicates a transition between qualitatively different dynamical behaviours of a nonlinear dynamical system, such as the predator–prey ecosystem illustrated in Fig. 1. Two other examples include the transition between laminar flow and periodic transport of a fluid contained between plates held at different temperatures, or the initiation of cell division (9). Kim et al. trained the RC using time-series data collected for different control parameter values just below a bifurcation, including the initial transient evolution. They then successfully used the trained closed-loop RC to predict the bifurcation's precise location and the response pattern after the bifurcation by slowly increasing the control parameter. This truly resembles Providentia's ability to foresee the unknown and demonstrates dramatically that the RC can learn an accurate model of the dynamical system.

The demonstrated model-free replication of a dynamical system from its time series, including stability properties, may have profound implications. It might be used to predict the moment of cell division by training a reservoir computer on previous observations of cell dynamics and real-time monitoring of the observed-cell behaviour. Another application includes predicting when a complex aircraft is about to begin unstable flight dynamics. However, more research is needed to clarify whether the surprising prediction abilities of RCs will also hold for stepwise changes of the control parameter and whether bifurcations in higher-dimensional dynamical systems can be predicted. Moreover, studies are needed to ascertain how the training data requirements scale in these cases, how robust the method is to noise, as well as whether the procedure of Kim et al. will work when only a subset of the dynamical variables are available during the testing phase. At the same time, these questions illustrate that the novel approach introduced by Kim et al. opens a fascinating field of fundamental and applied research in dynamical systems, machine learning and artificial intelligence.

References

1. Lukoševičius, M. & Jaeger, H. *Comp. Sci. Rev.* **3**, 127–149 (2009).
2. Rodan, A. & Tino, P. *IEEE Trans. Neural Netw.* **22**, 131–144 (2011).
3. Appeltant, L. et al. *Nat. Commun.* **2**, 468 (2011).
4. Tanaka, G. et al. *Neural Netw.* **115**, 100–123 (2019).
5. Kim, J. Z., Lu, Z., Nozari, E., Pappas, G. J. & Bassett, D. S. *Nat. Mach. Intell. https://doi.org/10.1038/s42256-021-00321-2 (2021).*
6. Lu, Z., Hunt, B. R. & Ott, E. *Chaos* **28**, 061104 (2018).
7. Klos, C., Kossio, Y. F. K., Goedeke, S., Gilra, A. & Memmesheimer, R.-M. *Phys. Rev. Lett.* **125**, 088103 (2020).
8. Feldkamp, L. A., Puskorius, G. V. & Moore, P. C. *Info. Sci.* **98**, 217–235 (1997).
9. Ferrell, J. E. Jr. *Curr. Biol.* **22**, R458–R466 (2012).

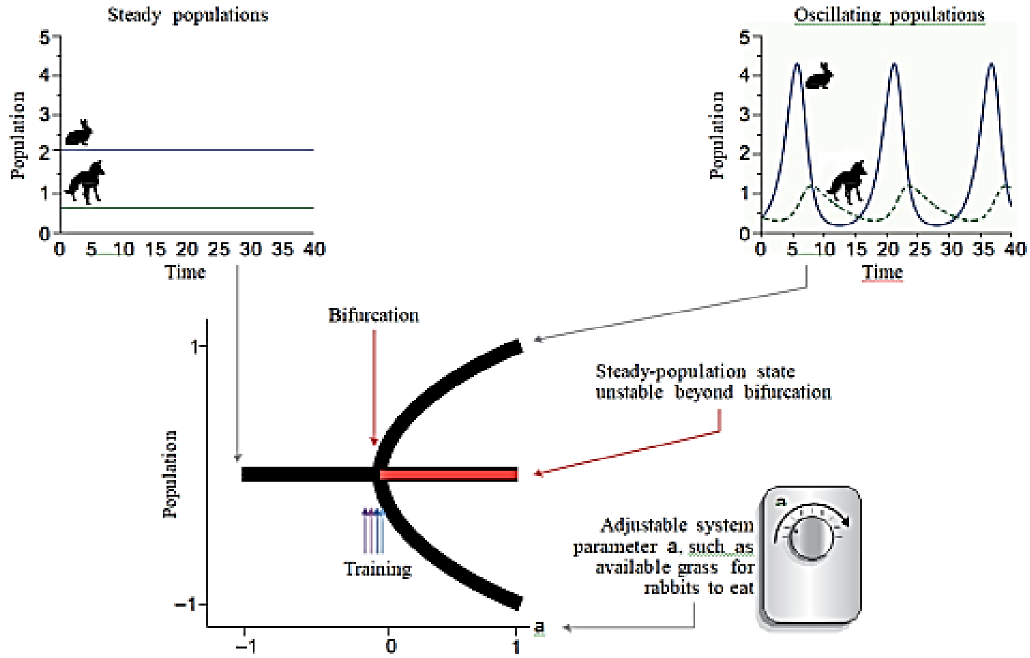

Fig. 1. A reservoir computer is trained on time-series data of a dynamical system — such as a fox– rabbit predator–prey system — below a bifurcation point. The curves indicate the population of the species as a function of an accessible system parameter a — an adjustable 'knob' such as the amount of food for the rabbits. The reservoir computer accurately predicts the behaviour of the ecosystem when the food supply is increased to the point where the population sizes begin to oscillate.

.

Competing interests

D.J.G. is a co-founder of ResCon Technologies, LLC, which focuses on commercializing reservoir computers for edge computing applications

Published in Nature Machine Intelligence, https://doi.org/10.1038/s42256-021-00329-8