

Teaching machines to anticipate catastrophes

Marcus Lapeyrolerie^{a,1} and Carl Boettiger^{a,1}

Tipping points may be responsible for some of the most dramatic transformations we observe, from the outbreak of a disease to the collapse of an ecosystem. By their very nature, they are also some of the most difficult phenomena to predict. Could new innovations using machine learning (ML) manage to cut this Gordian knot? In PNAS, Bury et al. (1) present a compelling illustration of that potential. By training a deep neural network on a massive corpus of simulations, they demonstrate how supervised learning successfully distinguishes which cases are rapidly approaching a tipping point and which are not.

Tipping points have been observed in a wide variety of complex systems, from brain waves to global climatic trends (2). The recognition that such a wide array of natural phenomena all share certain signature characteristics known as critical slowing down (CSD) (3) opened the door to the exciting possibility of discovering “universal indicators”—that is, early warning signs of an impending transition. Yet the past decade has demonstrated that such a task is easier said than done (4, 5).

A major weakness of current approaches based on CSD is that the pattern is simply too general, occurring in contexts where no catastrophic transition is imminent. Catastrophic transitions were first classified by mathematical topologist René Thom (6), whose work gave rise to catastrophe theory. CSD occurs not only in the tipping point transition known to mathematicians as the “fold bifurcation” but also in bifurcations which are not catastrophic, such as the transcritical and Hopf bifurcations (7). All three bifurcations have been observed in natural systems, where they can result in important changes to a system’s dynamics (5). For instance, a transcritical bifurcation occurs when either mutations or changes to social environment allow a virus’s basic reproductive number, R_0 , to exceed the critical threshold, one new infection for every infection (8). Only the fold bifurcation results in the free-fall transition to a distant alternative stable state and cannot be reversed by returning to the conditions observed before crossing the threshold: a phenomenon known as hysteresis.

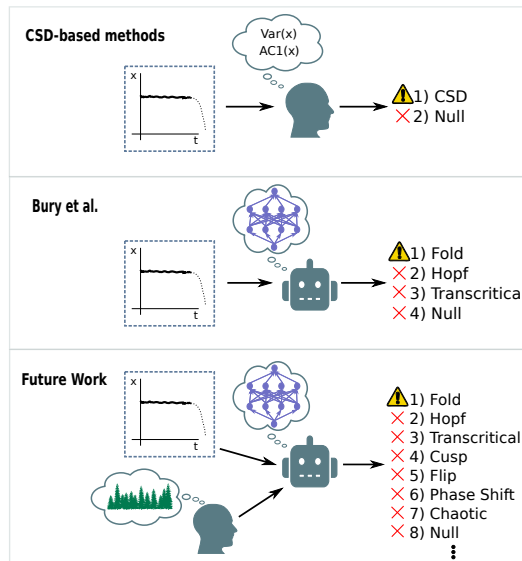


Fig. 1. CSD-based methods (2) which rely on simple statistical indicators may provide advanced warning of a critical transition but cannot distinguish between noncatastrophic and catastrophic bifurcations. By training a deep neural network on hundreds of thousands of simulations, the algorithm of Bury et al. (1) can reliably identify the catastrophic fold bifurcation from several alternatives. Future work may extend this capability to a larger array of possible transitions by not only training on a wider array of mathematical patterns but also allowing the neural network to include more of the biological context of each time series.

Not only do CSD-based methods produce false alarms for noncatastrophic bifurcations, but they also frequently fail to detect catastrophic bifurcations (9, 10). These methods require loads of rich, high-quality data to provide an adequate signal of the coming shift: Insufficient length, frequency, or precision in the data can easily mask warning signs. Both problems—the lack of specificity and the lack of sensitivity—happen because CSD does not look at the bigger picture. Bury et al. (1) find the bigger picture by examining

^aDepartment of Environmental Science, Policy, and Management, University of California, Berkeley, CA 94720

Author contributions: M.L. and C.B. analyzed data and wrote the paper.

The authors declare no competing interest.

Published under the [PNAS license](#).

See companion article, “Deep learning for early warning signals of tipping points,” [10.1073/pnas.2106140118](https://doi.org/10.1073/pnas.2106140118).

¹To whom correspondence may be addressed. Email: cboettig@berkeley.edu.

Published September 28, 2021.

the “normal form” equations of both catastrophic and noncatastrophic bifurcations.

CSD indicators effectively view the world through a pinhole, seeing what happens only to the “first-order” terms which are shared by both the catastrophic (fold bifurcation) and these noncatastrophic (Hopf, transcritical) bifurcations. The normal forms of each class of bifurcation reveal the subtler differences between them: specifically, “higher-order” characteristics which may reveal when a catastrophe is really approaching. Unfortunately, unlike CSD, whose signature is readily identified in common statistical indices like variance or autocorrelation, it is much harder to characterize these subtle higher-order terms in simple summary statistics. To overcome this considerable challenge, Bury et al. (1) turn to the most powerful classification method available: deep learning.

The expressiveness of deep neural networks has the potential to capture more-nuanced features and distinguish between the many different faces of critical transitions. Bury et al. (1) combine two well-recognized network architectures: convolutional neural networks (CNN), which have proven very successful in applications such as facial recognition, and long short-term memory (LSTM) networks, a recurrent architecture which has been critical in natural language processing and time series analysis (11). The authors then train their neural network against a massive corpus of data simulated from four classes of models: fold, Hopf, and transcritical bifurcations, as well as neutral processes. The trained neural network successfully identifies each of these four types in both empirical and model data with much higher sensitivity than existing CSD-based approaches.

Nevertheless, ML is only as good as the data with which it is trained. A machine which reliably distinguishes between cats and dogs will not successfully classify a catfish. The authors (1) have trained their deep learning algorithm in a world where only three kinds of bifurcations can exist (more precisely, in a world of two-dimensional, third-order polynomials). As they acknowledge, the real world is much richer—not only are there many other local bifurcation phenomena outside this set, but there are other dynamics, including chaotic dynamics (12), phase shifts (13), stochastic transitions (14), nonstationary or long transient dynamics (15), and other behavior only possible in higher-dimensional models. Adding these to the training data will dramatically increase the difficulty of the classification, likely requiring further innovations in the design and training of the neural network.

Even so, it is important to remember that these neural networks are being taught mathematics, not biology. While the ML algorithm

may be able to pick up the subtler signs of higher-order terms, it has had no access to the rich contextual information which a human researcher brings to their analysis. For example, is this a system where alternative stable states have been detected? Are there positive feedbacks, like cooperative social behavior, that could create the conditions required for a fold bifurcation? What are the time-scales of the processes involved, and of the changing environmental conditions in which they occur? If the ability to ignore system-specific information is part of the appeal of CSD, it has also been its greatest weakness (16). In contrast to statistical approaches, supervised learning is well positioned to make use of this biological context without the need for strong a priori assumptions about whether or how it is relevant. Future work should seek to incorporate this context into training data whenever possible, so that the machines may learn some biology along with their mathematics (Fig. 1).

The expressiveness of deep neural networks has the potential to capture more-nuanced features and distinguish between the many different faces of critical transitions. Bury et al. combine two well-recognized network architectures: convolutional neural networks (CNN), which have proven very successful in applications such as facial recognition, and long short-term memory (LSTM) networks, a recurrent architecture which has been critical in natural language processing and time series analysis.

By not only presenting their results but also providing the open source code necessary to reproduce both the training library and the deep learning classifier, Bury et al. (1) lay the groundwork for future research. The challenge of creating such a general-purpose detector of catastrophic transitions is immense—not a task for any single research team. But it need not be done in one step either. Domain scientists can help by contributing models and data of critical transitions that would extend the library far beyond its polynomial models, ideally including biological as well as mathematical annotations. The supervised ML problem becomes much harder as the library becomes more varied, but could become an open challenge to computer scientists and data scientists seeking to benchmark their algorithms against important scientific challenges.

- 1 T. M. Bury et al., Deep learning for early warning signals of tipping points. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2106140118 (2021).
- 2 M. Scheffer et al., Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009).
- 3 C. Wissel, A universal law of the characteristic return time near thresholds. *Oecologia* **65**, 101–107 (1984).
- 4 C. Boettiger, N. Ross, A. Hastings, Early warning signals: The charted and uncharted territories. *Theor. Ecol.* **6**, 255–264 (2013).
- 5 M. Scheffer, S. R. Carpenter, V. Dakos, E. van Nes, Generic indicators of ecological resilience: Inferring the chance of a critical transition. *Annu. Rev. Ecol. Evol. Syst.* **46**, 145–167 (2015).
- 6 R. Thom, *Structural Stability and Morphogenesis: An Outline of a General Theory of Models* (Addison-Wesley, Reading, MA, 1989).
- 7 Y. Kuznetsov, *Elements of Applied Bifurcation Theory* (Applied Mathematical Sciences, Springer Science & Business Media, 2013), vol. 112.
- 8 S. M. O'Regan, J. M. Drake, Theory of early warning signals of disease emergence and leading indicators of elimination. *Theor. Ecol.* **6**, 333–357 (2013).
- 9 C. Boettiger, A. Hastings, Quantifying limits to detection of early warning for critical transitions. *J. R. Soc. Interface* **9**, 2527–2539 (2012).
- 10 C. T. Perretti, S. B. Munch, Regime shift indicators fail under noise levels commonly observed in ecological systems. *Ecol. Appl.* **22**, 1772–1779 (2012).
- 11 Y. LeCun, Y. Bengio, G. Hinton, Deep learning. *Nature* **521**, 436–444 (2015).
- 12 A. Hastings, D. B. Wysham, Regime shifts in ecological systems can occur with no warning. *Ecol. Lett.* **13**, 464–472 (2010).
- 13 K. Tunstrom et al., Collective states, multistability and transitional behavior in schooling fish. *PLoS Comput. Biol.* **9**, e1002915 (2013).
- 14 C. Boettiger, From noise to knowledge: How randomness generates novel phenomena and reveals information. *Ecol. Lett.* **21**, 1255–1267 (2018).
- 15 A. Hastings et al., Transient phenomena in ecology. *Science* **361**, 6406 (2018).
- 16 C. Boettiger, A. Hastings, Tipping points: From patterns to predictions. *Nature* **493**, 157–158 (2013).