

Interactive comment on “Generalization properties of neural networks trained on Lorenzsystems” by S. Scher and G. Messori

Anonymous Referee #3

Received and published: 31 July 2019

Summary:

This paper investigates the very topical question to what extent neural network models can be generalized in the context of dynamical systems. The authors use two well-known models which display chaotic behaviour: the Lorenz 63 and Lorenz 95 models. Two aspects are addressed. The first aspect is the representativity of a neural network which is trained on a severely limited set of training data from the Lorenz 63 model (in this case, the removal of one entire "wing" of the Lorenz butterfly, or just the tip of one wing). The second aspect relates to the representativity of a neural network model under a changing parameter or forcing, when this parameter is provided as input. These parameters are the sigma parameter in the Lorenz 63 model, and the forcing parameter F in the Lorenz 95 model.

Printer-friendly version

Discussion paper



The neural network is evaluated not only by its short-term predictive skill but also by its ability to reproduce the long-term characteristics of the original model. Three different metrics are proposed, the so-called density-selection and density-full approaches, and a third approach which rejects neural network models with fixed points or periodic solutions.

The authors show that indeed, selectively removing half the training data will lead to a NN model which performs very poorly outside the wing on which it was trained. They interpret this as follows: the NN model approximates a local, and not a global function. They try to pinpoint the source of this locality by looking at the spread of the neurons' activations.

Likewise, training the model for a certain range of sigma values does not yield a good model outside this range, even if sigma is provided as an input parameter. There is no clear benefit from including the forcing parameter F as an input parameter in the analogous experiment in the Lorenz 95 model.

General comments:

Given the recent enormous successes of machine learning, it's only natural that this approach is being adopted enthusiastically in many different branches of science. This makes it all the more important to highlight the limitations of machine learning methods in disciplines such as climate science. I applaud the authors' effort to provide a tangible example of where the NN approach breaks down and to investigate what the causes are.

The paper is easy to read and the results are presented in a clear and well-structured way. The reproducibility and transparency of the results are supported by the availability of the code on Zenodo. Overall the manuscript is of a high quality. As for the novelty of the results and the context, however, I have some remarks that may require a substantial revision of the manuscript.

[Printer-friendly version](#)[Discussion paper](#)

A first remark is that the first experiment seems quite artificial. The Lorenz 63 model is an extremely idealized model with a very peculiar bifurcation structure and distinct symmetries, making it less than ideal to represent a realistic general circulation model. A more realistic model which also displays regimes such as that of Charney and Straus (1980), would probably have been more suitable for this particular experiment. Of course, the authors mention these caveats, but in the end there are so many caveats that it seems that no conclusion can be drawn at all related to realistic climate models. Moreover, it is not because some insights related to machine learning map to more complex models, that this is a universal property. I'm not sure how the authors can address this issue without performing the analysis for a more realistic model. Nevertheless, I appreciate the value of this experiment, albeit in a more theoretical context of dynamical systems theory.

Secondly, the result of the forcing experiment of the Lorenz 95 model seems hardly surprising, as the forcing parameter is varied from the periodic regime to the turbulent regime. One cannot expect a neural network to predict qualitatively completely different behaviour.

Finally, it seems to me that similar studies must have been performed in the literature on neural networks, though not necessarily in the context of geophysics. I would encourage the authors to explore the literature on this. The recent groundbreaking success of deep learning was only possible thanks to the move from few to many hidden layers, and it appears that large deep learning networks have better generalization properties than smaller ones. It would therefore also be interesting to repeat the exercise for a deep neural network. See for example the work by Novak et al. (2018) or Wu et al. (2017) who investigate the source of these generalization properties, and references therein.

Specific comments and typographical errors:

p.1, L 9-10: In the abstract, the authors conclude that "These results outline challenges

[Printer-friendly version](#)[Discussion paper](#)

for a variety of machine-learning applications. [...]". The word "outline" (in the sense of summarize) goes a bit too far since the results shown are for two highly specific models and a very artificial set-up (an entire wing missing, training in periodic regime). I would just say that the results provide some examples.

Figure 3: Labels in the caption don't match with the relevant subfigures.

p. 11, L 16: lorenz -> Lorenz (2x)

p. 15, L 1: However, also the alternative methods suffer -> However, the alternative methods also suffer

p. 15, L 3: test test -> test

Check spelling consistency: throughout the manuscript, generalize / generalise are both used

References:

Charney J G, Straus DM (1980) Form-Drag Instability, Multiple Equilibria and Propagating Planetary Waves in Baroclinic, Orographically Forced, Planetary Wave Systems. *J Atmos Sci* 37: 1157-1176.

Roman Novak, Yasaman Bahri, Daniel A. Abolafia, Jeffrey Pennington, Jascha Sohl-Dickstein (2018) Sensitivity and Generalization in Neural Networks: an Empirical Study. arXiv:1802.08760.

Lei Wu, Zhanxing Zhu, Weinan E (2017) Towards Understanding Generalization of Deep Learning: Perspective of Loss Landscapes. arXiv:1706.10239.

Interactive comment on Nonlin. Processes Geophys. Discuss., <https://doi.org/10.5194/npg-2019-23>, 2019.

Printer-friendly version

Discussion paper

