


A method for designing neural networks optimally suited for certain tasks

March 30 2023, by Adam Zewe



Credit: Pixabay/CC0 Public Domain

Neural networks, a type of machine-learning model, are being used to help humans complete a wide variety of tasks, from predicting if someone's credit score is high enough to qualify for a loan to diagnosing whether a patient has a certain disease. But researchers still have only a limited understanding of how these models work. Whether a given model is optimal for certain task remains an open question.

MIT researchers have found some answers. They conducted an analysis of [neural networks](#) and proved that they can be designed so they are "optimal," meaning they minimize the probability of misclassifying borrowers or patients into the wrong category when the networks are given a lot of labeled [training data](#). To achieve optimality, these networks must be built with a specific architecture.

The researchers discovered that, in certain situations, the building blocks that enable a [neural network](#) to be optimal are not the ones developers use in practice. These optimal building blocks, derived through the new analysis, are unconventional and haven't been considered before, the researchers say.

In a paper published this week in the *Proceedings of the National Academy of Sciences*, they describe these optimal [building blocks](#), called activation functions, and show how they can be used to design neural networks that achieve better performance on any dataset. The results hold even as the neural networks grow very large. This work could help developers select the correct activation function, enabling them to build neural networks that classify data more accurately in a wide range of application areas, explains senior author Caroline Uhler, a professor in the Department of Electrical Engineering and Computer Science

(EECS).

"While these are new activation functions that have never been used before, they are simple functions that someone could actually implement for a particular problem. This work really shows the importance of having theoretical proofs. If you go after a principled understanding of these models, that can actually lead you to new activation functions that you would otherwise never have thought of," says Uhler, who is also co-director of the Eric and Wendy Schmidt Center at the Broad Institute of MIT and Harvard, and a researcher at MIT's Laboratory for Information and Decision Systems (LIDS) and its Institute for Data, Systems and Society (IDSS).

Joining Uhler on the paper are lead author Adityanarayanan Radhakrishnan, an EECS graduate student and an Eric and Wendy Schmidt Center Fellow, and Mikhail Belkin, a professor in the Halicioğlu Data Science Institute at the University of California at San Diego.

Activation investigation

A neural network is a type of machine-learning model that is loosely based on the human brain. Many layers of interconnected nodes, or neurons, process data. Researchers train a network to complete a task by showing it millions of examples from a dataset.

For instance, a network that has been trained to classify images into categories, say dogs and cats, is given an image that has been encoded as numbers. The network performs a series of complex multiplication operations, layer by layer, until the result is just one number. If that number is positive, the network classifies the image a dog, and if it is negative, a cat.

Activation functions help the network learn complex patterns in the input data. They do this by applying a transformation to the output of one layer before data are sent to the next layer. When researchers build a neural network, they select one activation function to use. They also choose the width of the network (how many neurons are in each layer) and the depth (how many layers are in the network.)

"It turns out that, if you take the standard activation functions that people use in practice, and keep increasing the depth of the network, it gives you really terrible performance. We show that if you design with different activation functions, as you get more data, your network will get better and better," says Radhakrishnan.

He and his collaborators studied a situation in which a neural network is infinitely deep and wide—which means the network is built by continually adding more layers and more nodes—and is trained to perform classification tasks. In classification, the network learns to place data inputs into separate categories.

'A clean picture'

After conducting a detailed analysis, the researchers determined that there are only three ways this kind of network can learn to classify inputs. One method classifies an input based on the majority of inputs in the training data; if there are more dogs than cats, it will decide every new input is a dog. Another method classifies by choosing the label (dog or cat) of the training data point that most resembles the new input.

The third method classifies a new input based on a weighted average of all the training data points that are similar to it. Their analysis shows that this is the only method of the three that leads to optimal performance. They identified a set of activation functions that always use this optimal classification method.

"That was one of the most surprising things—no matter what you choose for an activation function, it is just going to be one of these three classifiers. We have formulas that will tell you explicitly which of these three it is going to be. It is a very clean picture," he says.

They tested this theory on a several classification benchmarking tasks and found that it led to improved performance in many cases. [Neural network](#) builders could use their formulas to select an activation function that yields improved classification performance, Radhakrishnan says.

In the future, the researchers want to use what they've learned to analyze situations where they have a limited amount of data and for networks that are not infinitely wide or deep. They also want to apply this analysis to situations where data do not have labels.

"In [deep learning](#), we want to build theoretically grounded models so we can reliably deploy them in some mission-critical setting. This is a promising approach at getting toward something like that—building architectures in a theoretically grounded way that translates into better results in practice," he says.

More information: Adityanarayanan Radhakrishnan et al, Wide and deep neural networks achieve consistency for classification, *Proceedings of the National Academy of Sciences* (2023). [DOI: 10.1073/pnas.2208779120](https://doi.org/10.1073/pnas.2208779120)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Citation: A method for designing neural networks optimally suited for certain tasks (2023, March 30) retrieved 26 April 2024 from <https://techxplore.com/news/2023-03-method-neural-networks->

[optimally-tasks.html](#)

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.