

Where once were black boxes, a new statistical tool illuminates

June 22 2022



How do you figure out how to alter a gene so that it makes a usefully different protein? The job might be imagined as interacting with a complex machine (at left) that sports a vast control panel filled with thousands of unlabeled switches, which all affect the device's output somehow. A new tool called LANTERN figures out which sets of switches—rungs on the gene's DNA ladder—have the largest effect on a given attribute of the protein. It also summarizes how the user can tweak that attribute to achieve a desired effect, essentially transmuting the many switches on our machine's panel into another machine (at right) with just a few simple dials. Credit: B. Hayes / NIST

Researchers at the National Institute of Standards and Technology (NIST) have developed a new statistical tool that they have used to predict protein function. Not only could it help with the difficult job of altering proteins in practically useful ways, but it also works by methods that are fully interpretable—an advantage over the conventional artificial intelligence (AI) that has aided with protein engineering in the past.

The new tool, called LANTERN, could prove useful in work ranging from producing biofuels to improving crops to developing new disease treatments. Proteins, as building blocks of biology, are a key element in all these tasks. But while it is comparatively easy to make changes to the strand of DNA that serves as the blueprint for a given [protein](#), it remains challenging to determine which specific base pairs—rungs on the DNA ladder—are the keys to producing a desired effect. Finding these keys has been the purview of AI built of deep neural networks (DNNs), which, though effective, are notoriously opaque to human understanding.

Described in a new paper published in the *Proceedings of the National Academy of Sciences*, LANTERN shows the ability to predict the genetic edits needed to create useful differences in three different proteins. One is the spike-shaped protein from the surface of the SARS-CoV-2 virus that causes COVID-19; understanding how changes in the DNA can alter this spike protein might help epidemiologists predict the future of the pandemic. The other two are well-known lab workhorses: the LacI protein from the *E. coli* bacterium and the green fluorescent protein (GFP) used as a marker in biology experiments. Selecting these three subjects allowed the NIST team to show not only that their tool works, but also that its results are interpretable—an important characteristic for industry, which needs predictive methods that help with understanding of the underlying system.

"We have an approach that is fully interpretable and that also has no loss in predictive power," said Peter Tonner, a statistician and computational

biologist at NIST and LANTERN's main developer. "There's a widespread assumption that if you want one of those things you can't have the other. We've shown that sometimes, you can have both."

The problem the NIST team is tackling might be imagined as interacting with a complex machine that sports a vast control panel filled with thousands of unlabeled switches: The device is a gene, a strand of DNA that encodes a protein; the switches are base pairs on the strand. The switches all affect the device's output somehow. If your job is to make the machine work differently in a specific way, which switches should you flip?

Because the answer might require changes to multiple base pairs, scientists have to flip some combination of them, measure the result, then choose a new combination and measure again. The number of permutations is daunting.

"The number of potential combinations can be greater than the number of atoms in the universe," Tonner said. "You could never measure all the possibilities. It's a ridiculously large number."

Because of the sheer quantity of [data](#) involved, DNNs have been tasked with sorting through a sampling of data and predicting which base pairs need to be flipped. At this, they have proved successful—as long as you don't ask for an explanation of how they get their answers. They are often described as "black boxes" because their inner workings are inscrutable.

"It is really difficult to understand how DNNs make their predictions," said NIST physicist David Ross, one of the paper's co-authors. "And that's a big problem if you want to use those predictions to engineer something new."

LANTERN, on the other hand, is explicitly designed to be understandable. Part of its explainability stems from its use of interpretable [parameters](#) to represent the data it analyzes. Rather than allowing the number of these parameters to grow extraordinarily large and often inscrutable, as is the case with DNNs, each parameter in LANTERN's calculations has a purpose that is meant to be intuitive, helping users understand what these parameters mean and how they influence LANTERN's predictions.

The LANTERN model represents protein mutations using [vectors](#), widely used mathematical tools often portrayed visually as arrows. Each arrow has two properties: Its direction implies the effect of the mutation, while its length represents how strong that effect is. When two proteins have vectors that point in the same direction, LANTERN indicates that the proteins have similar function.

These vectors' directions often map onto biological mechanisms. For example, LANTERN learned a direction associated with protein folding in all three of the datasets the team studied. (Folding plays a critical role in how a protein functions, so identifying this factor across datasets was an indication that the model functions as intended.) When making predictions, LANTERN just adds these vectors together—a method that users can trace when examining its predictions.

Other labs had already used DNNs to make predictions about what switch-flips would make useful changes to the three subject proteins, so the NIST team decided to pit LANTERN against the DNNs' results. The new approach was not merely good enough; according to the team, it achieves a new state of the art in predictive accuracy for this type of problem.

"LANTERN equaled or outperformed nearly all alternative approaches with respect to prediction accuracy," Tonner said. "It outperforms all

other approaches in predicting changes to LacI, and it has comparable predictive accuracy for GFP for all except one. For SARS-CoV-2, it has higher predictive accuracy than all alternatives other than one type of DNN, which matched LANTERN's accuracy but didn't beat it."

LANTERN figures out which sets of switches have the largest effect on a given attribute of the protein—its folding stability, for example—and summarizes how the user can tweak that attribute to achieve a desired effect. In a way, LANTERN transmutes the many switches on our machine's panel into a few simple dials.

"It reduces thousands of switches to maybe five little dials you can turn," Ross said. "It tells you the first dial will have a big effect, the second will have a different effect but smaller, the third even smaller, and so on. So as an engineer it tells me I can focus on the first and second dial to get the outcome I need. LANTERN lays all this out for me, and it's incredibly helpful."

Rajmonda Caceres, a scientist at MIT's Lincoln Laboratory who is familiar with the method behind LANTERN, said she values the tool's interpretability.

"There are not a lot of AI methods applied to biology applications where they explicitly design for interpretability," said Caceres, who is not affiliated with the NIST study. "When biologists see the results, they can see what mutation is contributing to the change in the protein. This level of interpretation allows for more interdisciplinary research, because biologists can understand how the algorithm is learning and they can generate further insights about the biological system under study."

Tonner said that while he is pleased with the results, LANTERN is not a panacea for AI's explainability problem. Exploring alternatives to DNNs more widely would benefit the entire effort to create explainable,

trustworthy AI, he said.

"In the context of predicting genetic effects on [protein function](#), LANTERN is the first example of something that rivals DNNs in predictive power while still being fully interpretable," Tonner said. "It provides a specific solution to a specific problem. We hope that it might apply to others, and that this work inspires the development of new interpretable approaches. We don't want predictive AI to remain a black box."

More information: Peter D. Tonner et al, Interpretable modeling of genotype–phenotype landscapes with state-of-the-art predictive power, *Proceedings of the National Academy of Sciences* (2022). [DOI: 10.1073/pnas.2114021119](#)

Provided by National Institute of Standards and Technology

Citation: Where once were black boxes, a new statistical tool illuminates (2022, June 22) retrieved 24 May 2024 from <https://phys.org/news/2022-06-black-statistical-tool-illuminates.html>

<p>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.</p>
--