


One of the Biggest Problems in Biology Has Finally Been Solved

Google DeepMind CEO Demis Hassabis explains how its AlphaFold AI program predicted the 3-D structure of every known protein

By Tanya Lewis on October 31, 2022



Demis Hassabis, software developer and co-founder of DeepMind Technologies. Credit: Tobias Hase/dpa/Alamy Live News



There's an age-old adage in biology: structure determines function. In order to understand the function of the myriad proteins that perform vital jobs in a healthy body—or malfunction in a diseased one—scientists have to first determine these proteins' molecular structure. But this is no easy feat: protein molecules consist of long, twisty chains of up to thousands of amino acids, chemical compounds that can interact with one another in many ways to take on an enormous number of possible three-dimensional shapes. Figuring out a single protein's structure, or solving the “protein-folding problem,” can take years of finicky experiments.

But earlier this year an artificial intelligence program called AlphaFold, developed by the Google-owned company DeepMind, predicted the 3-D structures of almost every known protein—about 200 million in all. DeepMind CEO Demis Hassabis and senior staff research scientist John Jumper were jointly awarded this year’s \$3-million Breakthrough Prize in Life Sciences for the achievement, which opens the door for applications that range from expanding our understanding of basic molecular biology to accelerating drug development.

DeepMind developed AlphaFold soon after its AlphaGo AI made headlines in 2016 by beating world Go champion Lee Sedol at the game. But the goal was always to develop AI that could tackle important problems in science, Hassabis says. DeepMind has made the structures of proteins from nearly every species for which amino acid sequences exist freely available in a public database.

Scientific American spoke with Hassabis about developing AlphaFold, some of its most exciting potential applications and the ethical considerations of highly sophisticated AI.

[An edited transcript of the interview follows.]

Why did you decide to create AlphaFold, and how did you get to the point where it can now fold practically every known protein?

We pretty much started the project roughly the day after we came back from the AlphaGo match in Seoul, where we beat Lee Sedol, the world [Go] champion. I was talking to Dave Silver, the project lead on AlphaGo, and we were discussing “What’s the next big project that DeepMind should do?” I was feeling like it was time to tackle something really hard in science because we had just solved more or less the pinnacle of games AI. I wanted to finally apply the AI to real-world domains. That’s always been the mission of DeepMind: to develop general-purpose algorithms that could be applied really generally across many, many problems. We started off with games because it was really efficient to develop things and test things out in games for various reasons. But ultimately, that was never the end goal. The end goal was [to develop] things like AlphaFold.

It’s been a mammoth project—about five or six years’ worth of work before CASP14 [Critical Assessment of Structure Prediction, a protein-folding competition]. We had an earlier version at the CASP13 competition, and that was AlphaFold 1. That was state of the art, you know, a good deal better than anyone had done before and I think one of the first times that machine learning had been used as the core component of a system to try and crack this problem. That gave us the confidence to push it even further. We had to reengineer things for AlphaFold 2 and put a whole bunch of new ideas in there and also bring onto the team some more specialists—biologists and chemists and biophysicists who worked in protein folding—and combine them with our engineering and machine-learning team.

I've been working on and thinking about general AI for my whole career, even back at university. I tend to note down scientific problems I think one day could be amenable to the types of algorithms we build, and protein folding was right up there for me always, since the 1990s. I've had many, many biologist friends who used to go on about this to me all the time.

Were you surprised that AlphaFold was so successful?

Yeah, it was surprising, actually. I think it's definitely been the hardest thing we've done, and I would also say the most complex system we've ever built. The Nature paper that describes all the methods, with the supplementary information and technical details, is 60 pages long. There are 32 different component algorithms, and each of them is needed. It's a pretty complicated architecture, and it needed a lot of innovation. That's why it took so long. It was really important to have all these different inputs from different backgrounds and disciplines. And I think something we do uniquely well at DeepMind is mix that together—not just machine learning and engineering.

But there was a difficult period after AlphaFold 1. What we did first was we tried to push AlphaFold 1 to the maximum. And we realized about six months after CASP13 that it was not going to reach the atomic accuracy we wanted to actually solve the problem and be useful to experimentalists and biologists. So I made the decision that we needed to go back to the drawing board and take the knowledge we had acquired, including where it worked and where it didn't work, and then see if we could really go back to almost a brainstorming stage with that experience and that knowledge and come up with a whole bunch of new ideas and new architectures. We did that, and ultimately that worked. But for about six months to a year after that reset, things got worse, not better. The AlphaFold 2 system, the early one, was much worse than AlphaFold 1. It can be very scary during the period where you seem to be going backward in terms of accuracy. Fortunately, that's where our experience in games and all the other AI systems we built before then came into play. I'd seen us go through that valley of death and then get out the other side.

Can you explain, on a very simple level, how AlphaFold works?

It's a pretty complicated thing. And we don't know a lot of things for sure. It's clear that AlphaFold 2 is learning something implicit about the structure of chemistry and physics. It sort of knows what things might be plausible. It's learned that through seeing real protein structures, the ones that we know of. But also, one of the innovations we had was to do something called self-distillation, which is: get an early version of AlphaFold 2 to predict lots of structures—but also to predict the confidence level in those predictions.

One of the things we built in was this understanding of chemical bond angles and also evolutionary history using a process called multisequence alignment. These bring in some

constraints, which help to narrow the search space of possible protein structures. The search space is too huge to do by brute force. But obviously, real-world physics solves this somehow because proteins fold up in nanoseconds or milliseconds. Effectively, we're trying to reverse engineer that process by learning from the output examples. I think AlphaFold has captured something quite deep about the physics and the chemistry of molecules.

The fascinating thing about AI in general is that it's kind of a black box. But ultimately, it seems like it's learning actual rules about the natural world.

Yeah, it's almost learning about it in an intuitive sense. I think we'll have more and more researchers looking at protein areas that AlphaFold is not good at predicting and asking, "Are they actually disordered in biology when the protein doesn't have a clear shape, when it's not interacting with something? About 30 percent of proteins [from organisms with a nucleus] are thought to be disordered. A lot of those kinds of proteins are implicated in disease, such as neurodegeneration, because they might get tangled. And you can see how they might do so because they're just sort of floppy strings rather than forming structures.

The other really important thing we did with AlphaFold 2, which we don't do with machine-learning systems, was output a confidence measure on every single amino acid because we wanted a biologist to be able to really know which parts of the prediction they could rely on without needing to understand anything about the machine learning.

What are some of AlphaFold's most exciting applications?

We have a lot of really nice case studies from partners—early adopters—that have had a year to work with AlphaFold. They're doing an incredibly diverse set of things, from addressing antibiotic resistance to tackling plastic pollution by designing plastic-eating enzymes. I've been talking to [CRISPR pioneer] Jennifer Doudna about alfalfa crop sustainability—her team is trying to engineer crops to be a bit more sustainable in the face of climate change.

But there's also lots of really cool fundamental research being done with it. There was a whole special issue in *Science* on the nuclear pore complex. They solved the structure of one of the biggest proteins in the body. And I think three groups solved it at the same time from the cryo-EM [cryogenic electron microscopy] data—but they all needed AlphaFold predictions to augment the cryo-EM lower-resolution data in some places. So a combination of experimental structural data with AlphaFold turns out to be a real boon to structural biologists, which we weren't necessarily predicting.

And then in practical terms, almost every pharma company we've talked to is using AlphaFold. We'll probably never know what the full impacts are because obviously, they keep

that proprietary. But I like to think we've helped accelerate real cures to diseases and drug development maybe by a few years.

There's been a lot of hype around AI and everything it can do, especially for science and medicine. But AlphaFold seems to have a clear benefit.

I mean, it's for you to decide. But I would say I've had a lot of people tell me that it's the most concrete, useful case of AI doing something in science. I like the fact that we're delivering on the promise of AI. I mean, you could say "hype," but we try and let our work speak for itself.

I remember when we started in 2010, nobody was working on AI. And then now, 12 years later, it seems like everyone and their dog is talking about it. And in most cases, as I'm sure you have to sift through all the time, it's like they don't know what AI even is sometimes or they're misusing the term or it's not really impressive what's going on. But I think AlphaFold is a really good proof of concept or role model of what could happen. And I think we're going to see a lot more of that in the next decade—of AI really helping to genuinely accelerate some scientific breakthroughs—and we hope to be part of a lot more. We think it's just the beginning.

Stepping back a bit, AI has been in the news a lot lately, whether for producing intelligent language or creating digital art. Do you think AI has become more embedded in the public consciousness, and how should we think about its consequences?

Yeah, sure. We [at DeepMind] have our own internal versions of large language models and text-to-image systems, and we'll probably be releasing some of them at some point next year. It's really interesting seeing the explosion of developments. AlphaFold, obviously, is huge in the scientific community. But with language and image AIs, it's starting to break through into the mainstream, because obviously everyone knows about language and can appreciate images. You don't have to have any scientific expertise.

But I think we should always be thinking about the ethical issues, and that's one reason we haven't released ours yet. We're trying to be responsible about really checking what these models can do—how they can go off the rails, what happens if they're toxic, all of these things that are currently top of mind. It's our view that some of these systems are not ready to release to the general public, at least not unrestricted. But at some point that's going to happen. We have this phrase at DeepMind of "pioneering responsibly." And for me, that's about applying the scientific method to analyzing these systems and building these systems. I think a lot of times, especially in Silicon Valley, there's this sort of hacker mentality of like "We'll just hack it and put it out there and then see what happens." And I think that's exactly the wrong approach for technologies as impactful and potentially powerful as AI.

I've worked on AI my entire life because I think it's going to be the most beneficial thing ever to humanity, things like curing diseases, helping with climate, all of this stuff. But it's a dual-use technology—it depends on how, as a society, we decide to deploy it—and what we use it for.

ABOUT THE AUTHOR(S)



Tanya Lewis is a senior editor at *Scientific American* who covers health and medicine. [Follow Tanya Lewis on Twitter](#) *Credit: Nick Higgins*

Recent Articles by Tanya Lewis

The New COVID Booster Shot Could Save Your Life; Get One Now, FDA Expert Says

The Most Urgent Science, Health and Climate Issues in the 2022 Midterm Elections

New Halloween 'Scariant' Variants and Boosting Your Immunity: COVID, Quickly, Episode 41

Scientific American is part of Springer Nature, which owns or has commercial relations with thousands of scientific publications (many of them can be found at www.springernature.com/us). Scientific American maintains a strict policy of editorial independence in reporting developments in science to our readers.

© 2022 SCIENTIFIC AMERICAN, A DIVISION OF SPRINGER NATURE AMERICA, INC.

ALL RIGHTS RESERVED.