

# **This is the reason Demis Hassabis started DeepMind**

A year after it took biologists by surprise, AlphaFold has changed how researchers work and set DeepMind on a new course.

[Will Douglas Heaven](#)



In March 2016 Demis Hassabis, CEO and cofounder of DeepMind, was in

Seoul, South Korea, watching his company's AI make history. AlphaGo, a computer program trained to master the ancient board game Go, played a five-game match against Lee Sedol, a top Korean pro with the second-highest number of international championship wins to his name at the time. Many consider Go the world's most complex board game; it takes years to master.

Lee predicted he would beat DeepMind's AI in a "landslide," but AlphaGo won 4–1. Its victory shocked Go and AI experts alike—and changed the world's perception of what AI can do.

But while the DeepMind team was celebrating, Hassabis was already thinking about an even bigger challenge. He remembers standing backstage with David Silver, who led the development of AlphaGo: "I said to him, 'Now is the time.'"

Watching DeepMind's AI play Go, Hassabis realized that his company's technology was ready to take on one of the most important and complicated puzzles in biology, one that researchers had been trying to solve for 50 years: predicting the structure of proteins.

The three-dimensional structure of proteins determines how they behave and interact in the body. But a large number of important proteins have structures that biologists still don't know. Using AI to accurately predict them would offer an invaluable tool to help understand diseases, from cancer to covid. Proteins are a primary target for many drugs and a key ingredient in new therapeutics. Quickly unlocking their structures would fast-track the development of new therapies and vaccines.

In 2020 DeepMind, which is owned by Alphabet, revealed AlphaFold2, an AI that could predict the shape of proteins down to the nearest atom. "It's the most complex thing we've ever done," says Hassabis.

AlphaFold's success is part of a bigger story, too, signaling a change of


direction for the AI lab. The company's focus is shifting from games to science, where it hopes to have a bigger real-world impact. Taking on scientific problems is the culmination of what Hassabis set out to achieve, and it's what he wants to be known for. "This is the reason I started DeepMind," he says. "In fact, it's why I've worked my whole career in AI."

Hassabis has been thinking about proteins on and off for 25 years. He was introduced to the problem when he was an undergraduate at the University of Cambridge in the 1990s. "A friend of mine there was obsessed with this problem," he says. "He would bring it up at any opportunity—in the bar, playing pool—telling me if we could just crack protein folding, it would be transformational for biology. His passion always stuck with me."

That friend was Tim Stevens, who is now a Cambridge researcher working on protein structures. "Proteins are the molecular machines that make life on earth work," Stevens says.

Nearly everything your body does, it does with proteins: they digest food, contract muscles, fire neurons, detect light, power immune responses, and much more. Understanding what individual proteins do is therefore crucial for understanding how bodies work, what happens when they don't, and how to fix them.

A protein is made up of a ribbon of amino acids, which chemical forces fold up into a knot of complex twists and twirls. The resulting 3D shape determines what it does. For example, hemoglobin, a protein that ferries oxygen around the body and gives blood its red color, is shaped like a little pouch, which lets it pick up oxygen molecules in the lungs. The structure of SARS-CoV-2's spike protein lets the virus hook onto your cells.



A model generated by AlphaFold shows how amino acids fold to form a protein.

COURTESY OF DEEPMIND

The catch is that it's hard to figure out a protein's structure—and thus its function—from the ribbon of amino acids. An unfolded ribbon can take  $10^{300}$  possible forms, a number on the order of all the possible moves in a game of Go.

Predicting this structure in a lab, using techniques such as x-ray crystallography, is painstaking work. Entire PhDs have been spent working out the folds of a single protein. The long-running CASP (Critical Assessment of Structure Prediction) competition was set up in 1994 to speed things up by pitting computerized prediction methods against each other every two years. But no technique ever came close to matching the accuracy of lab work. By 2016, progress had been flatlining for a decade.

Within months of its AlphaGo success in 2016, DeepMind hired a handful of

biologists and set up a small interdisciplinary team to tackle protein folding. The first glimpse of what they were working on came in 2018, when DeepMind won CASP 13, outperforming other techniques by a significant margin. But beyond the world of biology, few paid much attention.

That changed when AlphaFold2 came out two years later. Its landslide victory in CASP 14 marked the first time an AI had predicted protein structure with an accuracy matching that of models produced in an experimental lab—often with margins of error just the width of an atom. Biologists were stunned by just how good it was.

Watching AlphaGo play in Seoul, Hassabis says, he'd been reminded of an online game called FoldIt, which a team led by David Baker, a leading protein researcher at the University of Washington, released in 2008. FoldIt asked players to explore protein structures, represented as 3D images on their screens, by folding them up in different ways. With many people playing, the researchers behind the game hoped, some data about the probable shapes of certain proteins might emerge. It worked, and FoldIt players even contributed to a handful of [new discoveries](#).

*"If we can mimic the pinnacle of intuition in Go, then why couldn't we map that across to proteins?"*

Hassabis played that game when he was a postdoc at MIT in his 20s. He was struck by the way basic human intuition could lead to real breakthroughs, whether making a move in Go or finding a new configuration in FoldIt.

"I was thinking about what we had actually done with AlphaGo," says Hassabis. "We'd mimicked the intuition of incredible Go masters. I thought, if we can mimic the pinnacle of intuition in Go, then why couldn't we map that across to proteins?"



The two problems weren't so different, in a way. Like Go, protein folding is a problem with such vast combinatorial complexity that brute-force computational methods are no match. Another thing Go and protein folding have in common is the availability of lots of data about how the problem could be solved. AlphaGo used an endless history of its own past games; AlphaFold used existing protein structures from the Protein Data Bank, an international database of solved structures that biologists have been adding to for decades.

AlphaFold2 uses attention networks, a standard deep-learning technique that lets an AI focus on specific parts of its input data. This tech underpins language models like GPT-3, where it directs the neural network to relevant words in a sentence. Similarly, AlphaFold2 is directed to relevant amino acids in a sequence, such as pairs that might sit together in a folded structure. "They wiped the floor with the CASP competition by bringing together all these things biologists have been pushing toward for decades and then just acing the AI," says Stevens.

Over the past year, AlphaFold2 has started having an impact. DeepMind has published a detailed description of how the system works and released the source code. It has also set up a public database with the European Bioinformatics Institute that it is filling with new protein structures as the AI predicts them. The database currently has around 800,000 entries, and DeepMind says it will add more than 100 million—nearly every protein known to science—in the next year.

A lot of researchers still don't fully grasp what DeepMind has done, says Charlotte Deane, chief scientist at [Exscientia](#), an AI drug discovery company based in the UK, and head of the protein informatics lab at the University of Oxford. Deane was one of the reviewers of the paper that DeepMind published on AlphaFold in the scientific journal *Nature* last year. "It's changed the questions you can ask," she says.

A handful of teams around the world have started using AlphaFold in work on

antibiotic resistance, cancer, covid, and more. Roland Dunbrack at the Fox Chase Cancer Center in Philadelphia is one early adopter. He leads a team that has been using computers to predict protein structures for years. Other teams at the lab then use these structures to guide their experiments.

AlphaFold has introduced an unprecedented level of accuracy to Dunbrack's work. "They are accurate enough to make biological judgments from, to interpret mutations in a cancer gene," he says of its predictions. "We always tried to do that with computer-generated models before, but we were often wrong."

When colleagues ask him to model proteins, Dunbrack says, he can now be more confident in what he gives them. Otherwise, he says, "I get really nervous, worried that they'll come back to me and say, 'We wasted all this money and your model was terrible—it didn't work.'"

AlphaFold can still make mistakes, but when it works well it can be hard to tell the difference between its predictions and a structure produced in the lab, says Dunbrack. He runs AlphaFold predictions on a computer platform called ColabFold, hosted by Harvard University and running on Google GPUs. "Every night I set one up before I go to sleep, and they take a few hours to run," he says.

"It's a super useful tool that everybody in my lab is using," says Kliment Verba, a structural biologist at the University of California, San Francisco. Verba mostly works on cancer, but in the early weeks of the covid-19 pandemic, he joined a loose consortium of researchers studying the SARS-CoV-2 virus. In particular, he wanted to figure out how its proteins hijacked host proteins.

Verba and his colleagues had produced part of the structure for the viral protein they were interested in, but they were missing a piece. Many proteins have multiple domains—densely folded sections, a few hundred amino acids long, that can each have a separate function. One domain might bind to



DNA, another might bind to another protein, and so on. "They're multiheaded beasts," says Dunbrack.

Structurally, domains are like knots in a rope, connected by loose, looping strands that flop around. In the protein he was studying, Verba's team had figured out the rough shape of the rope but not the detailed structure of all the knots. Without that detail, there was little they could say about how it worked.

They realized, though, that this protein was one of those DeepMind had already run through AlphaFold and shared online. AlphaFold's prediction wasn't perfect; the looping strands weren't quite right. But it had the shape of the protein's four domains. The researchers took AlphaFold's predictions for the domains and lined them up with the rough shape they had. It was remarkably close.

"I remember that moment when I saw it fit," says Verba. "It was amazing. We were now the only ones in the world with the full structure." They published their findings soon after.

Verba thinks AlphaFold's strength lies in finding structures for proteins that have not yet been fully studied. "Many of the proteins we care about have been studied for decades," he says. "People have spent careers chipping away at them, so we have a fairly good idea what they look like." But that still leaves a lot of uncharted territory.

Verba is interested in kinases, for example. Kinases are enzymes that play a crucial role in regulating the normal function of cells. If they stop working properly, they can cause cancer. Only around half of the 500 or so kinases in the human body are well understood; the remainder is known as the dark kinome.

Researchers like Verba and Dunbrack are interested in developing cancer drugs that target the kinome. But this is where AlphaFold's limitations kick

in.

Because working out the structure of a protein in the lab is costly, it is typically done only once the protein has been picked as a promising candidate—which might be months into the drug discovery process. The hope, Deane says, is that AlphaFold could reverse that sequence, making the pipeline move faster. “Now I can start with the structure—I can identify where it has pockets on the surface, places where I can bind drug molecules,” she says.

*“A lot of the time these small transformations are the crux of biological function.”*

Yet—as Deane acknowledges—you need more than a static structure to fully understand how a drug and a protein might interact. Proteins do not stay still; their structures can cycle through subtle reconfigurations. “A lot of the time these small transformations are the crux of biological function,” says Verba.

What’s more, a protein may be open to receiving a drug in one state but not others. And judging from what researchers are seeing so far, AlphaFold appears to predict the most common state of these structures, which may not be the state that is important for drug development.

Proteins can also change shape when drugs bind to them, which can affect how the drug works. In the worst-case scenario, a drug binding to a protein can have unpredictable knock-on effects on adjoining proteins, potentially even reversing what the drug was designed to do—for example, activating rather than inhibiting some function.

Ola Engkvist, head of molecular AI in discovery sciences at AstraZeneca, thinks that AI-generated structures will help identify drug targets eventually—but not yet. “To be transformational, AlphaFold needs to be followed by

better computational methods to understand protein dynamics and handle larger protein complexes," he says.

DeepMind plans to address many of these issues in the next version of the program. One line of work is to generate multiple variations of a protein's shape to try to capture its dynamics. The way a protein moves is governed by complex chemistry and physics, so a full, moving model may require feeding AlphaFold large amounts of extra information about this process. A downside of this approach could be that the information might act as a constraint, degrading the tool's predictive abilities.

Last summer, DeepMind released AlphaFold Multimer, which is designed to predict the structure of protein complexes—superstructures made of multiple proteins clumped together. But it is much less accurate than AlphaFold, and prone to more glaring errors.

Stupid mistakes are a feature of even the best AI. AlphaGo made a basic error in the one game it lost to Lee Sedol, says Hassabis. "You can think of it a bit like a bug," he says. "But the problem is that it's a bug in its knowledge—you can't just go in and debug it."

That's because you can't easily tinker with a neural network without fundamentally affecting how it works. "Hard-coded fixes damage an AI's ability to learn, because how does it know when to use them?" says Hassabis. "It goes against the point of learning."

Instead, DeepMind is gathering examples of AlphaFold's worst mistakes and training it to handle them properly. Hassabis wants researchers to break AlphaFold—to find what doesn't work—and share the results with his team so that they can make the next AlphaFold even better.

With AlphaFold, DeepMind is starting a new chapter. The company is investing in a team called AI for Science. It has produced a flurry of publications in the last few months, in fields from weather prediction to