

Science Will Win | Season 4
Episode 3 | Software
Finalized Script

JEREMIAH

It is said that the average adult makes over 30,000 decisions each day.

[Music]

From the moment we wake, we're making decisions: what to wear, what to eat, how we spend our time, what we say and how we say it.

Many of these decisions simply happen without you even realizing it. But what you don't see is that within your brain a series of intricate actions make these decisions possible.

[SFX of crowd chatter]

You might imagine your brain as a council with some 86 billion neurons as its members. Much like committees within a council, these neurons form different groups representing potential choices. When you make a decision, the neuron committees rapidly send messages with evidence supporting the choices they represent. It's like trillions of emails being sent to one another, each one in the matter of milliseconds.

[Email notification pings SFX]

These neuron "committees" then compete with each other and vote on the choices. When they gather enough evidence to support a choice—and an information threshold is reached—they vote on the options. The most active group—that is the one with the strongest evidence—wins. And just like that, a decision is made.

[Sound cue to indicate finality]

Our brain's neural networks power just about everything we do. But what if I told you that artificial neural networks—which loosely mimic the ones in your brain—also power much of the artificial intelligence used today?

[Theme music]

JEREMIAH

This is Science Will Win. I'm your host, Jeremiah Owyang. I'm an entrepreneur, investor, and tech industry analyst. I'm passionate about emerging technologies and the ways they can shape our world.

In today's episode, we're diving deep into the third element that forms AI: software. We'll get into the neural networks, deep learning and machine learning at the heart of the AI technology that scientists are using today in the process of drug development and discovery.

[Theme music fades out]

Charlotte Allerton:

I think what's unique about the work we do in healthcare and the pharmaceutical industry is it really enables us to bring together fundamental science with an ability to make a difference for human health.

[Music]

JEREMIAH OWYANG

This is Charlotte Allerton.

Charlotte Allerton:

I started working in the industry over 30 years ago, firstly actually as a synthetic chemist in our medicinal chemistry department. And then over time I got more involved in the design of our medicines.

JEREMIAH

She leads preclinical and translational sciences at Pfizer. There, she and her team design and advance potential new therapeutic medicines that can be evaluated in patients.

It's a nuanced and challenging job to bring a potential drug to approval. Charlotte and her team spend a lot of time testing and analyzing different molecules before a potential drug candidate is even put through a clinical trial.

Charlotte Allerton:

It depends on potency and then lots of characteristics which more describe how will the body treat that molecule? So how quickly will the human liver want to get rid of the molecule from the body? Will it cross the blood-brain barrier? And will it cross the gut wall after it's been swallowed, dissolved? Will it then, you know,

cross the gut wall and enter into the blood circulation to get to the right place? And there's lots of different parameters that can be measured to assess that and we now have machine learning tools to predict that.

[Music fades out]

JEREMIAH

The machine learning tools that Charlotte and her team at Pfizer are using today, have improved over time thanks to a wealth of data, and increasingly powerful computing technology. But, there's a long, winding road of history that got us to this point.

Let's roll back the clock to understand how computers learned to think like humans—well, sort of.

[Clock ticking SFX and Music]

JEREMIAH

We start in the late 1930s, with a 15-year-old runaway from Detroit named Walter Pitts. Walter had just landed in Chicago, in search of a mathematician he admired. That is, until one of his friends introduced him to a professor of psychiatry named Warren McCulloch. McCulloch was trying to better understand how the human brain worked. And he had big ideas about the future of human and machine interactions. In the 1962 documentary "The Living Machine," he speaks of a world where machines are integrated into society, operating like humans.

Warren McCulloch: Insofar as machines might survive man, they would only carry on in the sense—the same direction that man would've carried on if he could've. So to speak, the machines would be standing on our shoulders.

JEREMIAH

Walter Pitts had no permanent place to live at the time and he showed incredible intellectual promise. So, McCulloch took him in. Together, they'd spend evenings pondering the question of whether the nervous system could be seriously likened to a computing device. This collaboration—a meeting of minds generations apart—led to their breakthrough 1943 paper titled "A Logical Calculus of the Ideas Imminent in Nervous Activity." The paper described the brain like an electrical circuit.

[Electrical zaps and switches SFX]

With its neurons a collection of tiny switches, each one turning on and off like a light bulb, working together to create thoughts and decisions. This was the first time a

comparison like that had been made. Emboldened by their work, McCulloch declared: “For the first time in the history of science, we know how we know.”

From their work, the idea of the neural network was born. In the decades that followed, neuroscientists, psychologists, mathematicians, computer scientists, and logicians set their sights on new theories and applications of neural nets.

Nearly a decade later, Frank Rosenblatt, a psychologist at Cornell University, turned the McCulloch-Pitts model into a physical product. He created the first artificial neural network, called the Perceptron. The system used an input-output relationship to mimic the brain’s decision-making process. Compared to the original concept of the neural net, the Perceptron was trainable. So, with each successive input, the machine would learn from the data and improve its ability to classify information. It was an oversimplified representation of the brain, but it laid the foundation for today’s artificial neural networks.

[Music ends]

In the decades to come, the hype around neural networks grew rapidly. By 1956, artificial intelligence had formally become a scientific field, intent on the potential of this technology. The media became crazed with the idea of “thinking machines” that could one day think just like us. Here’s Oliver G. Selfridge—mathematician and pioneer of early AI—on a 1961 episode of “Tomorrow,” which aired on CBS.

Oliver G. Selfridge: I’m convinced that machines can and will think. I don’t think for a very long time we’re going to have a difficult problem distinguishing a man from a robot, and I don’t think my daughter will ever marry a computer. But I think the computers will be doing the things that men do when they say they’re thinking.

JEREMIAH

But as it would turn out, these simple brain-mimicking models were not all they were cracked up to be.

[Music]

The human brain is complex, and real neurons have far more intricacies than the early artificial neural networks could account for.

In 1969, scientists at MIT published a book laying out all of the issues with neural nets: they weren’t fast enough, nor sophisticated enough to live up to the hype. In short, they

said the Perceptron could not solve important problems. This made waves in the scientific community and effectively shut down funding for neural networks for the next 10-12 years. With that, the “AI Winter” had begun.

[Music ends]

It wasn't until 1982 that this long winter began to thaw. Around the world renewed efforts on neural network innovation were cropping up, and conferences dedicated to neural nets emerged. With this renewed enthusiasm, advancements in memory, trainability, and multilayer capabilities for neural nets increased. Simultaneously, the computer game industry was booming, and computer scientists became interested in how they could marry hardware like GPUs and CPUs, with the architecture of neural networks.

By the 1990s, neural networks were back in, but this time, they exceeded the lofty expectations.

[Music]

And today, neural networks power the different components of artificial intelligence like machine learning—which helps AI learn patterns from past experience—and deep learning—which processes data through computation.

Biomedical professionals have now employed AI software in their work for more than a decade. It comes into play from the first stages of drug development, when scientists look at large molecules like proteins, and envision how they could optimize them to become the perfect clinical candidate.

Daniel:

There's the very first stage when you're thinking about which of the many proteins in the human body do you want to, for instance, inhibit to try to ideally cure a certain disease.

JEREMIAH

That's Daniel Ziemek, you've heard from him throughout the season. He's the Vice President of Integrative Biology and Systems Immunology at Pfizer. He oversees teams of computational biologists who work on creating new medicines. Many of these medicines contain monoclonal antibodies, which are laboratory-made proteins used to help treat diseases. These are different than the antibodies naturally produced by our immune systems.

Daniel:

At some point, the time comes where you want to either develop a chemical molecule or a so-called antibody—which is a little bit of a bigger molecule that has certain properties—that you want to then use to inhibit the function of the protein you have selected.

JEREMIAH

These antibodies target certain molecular and cellular processes that are dysregulated by disease.

The design of these antibodies—as well as other types of therapeutics—can be further enhanced if the structure of the molecule, or potential medicine and its target, are known.

Daniel:

There have been experimental methods to find protein structures now for a while, but they're expensive, they take a lot of time.

JEREMIAH

There's a whole field of science dedicated to predicting the structure of proteins. The scientists doing this work rely on clear metrics and data to learn from, and become more accurate in their predictions. Well, AI software also uses data to train and improve. Scientists realized this, and decided to try using AI to help them make these predictions.

[Music fades out]

Daniel:

All these AI methods, machine learning methods, most of them work by having a so-called feature set that is describing the problem you're trying to solve. In the terms of protein structure it's these so-called amino acid that make up that protein. So you feed that in—many proteins in the form of amino acid sequence—and then you also give them true structures, structures that have been experimentally determined previously to basically look at, how does the amino acid sequence relate to the 3D structure that we have experimentally determined. The system learns—in this deep neural network—how to translate the amino acid sequence to the ultimate protein structure.

And this is really starting to come into pharmaceutical companies like Pfizer to directly impact how we think about protein structure prediction, and then what we know about when we develop these compounds, we say, these chemical molecules that can become future medicines.

JEREMIAH

In the drug design and development process, scientists work hard to finesse the properties of molecules that comprise potential therapeutics. This is necessary to make sure that medicines are safe enough to enter a clinical trial. Some of these properties include structure, size, atomic charge, and viscosity.

You might think of viscosity as the thickness of a fluid or how much it resists flow. For example, water has a very low viscosity, and honey has a high viscosity. Therapeutic antibodies can be measured by viscosity too.

[Music]

Enoch:

So viscosity is a big deal because in order to deliver the therapeutic dose of the drug you need to have molecules that aren't overly viscous because you can't deliver it through a narrow gauge needle.

JEREMIAH

Enoch Huang leads a group at Pfizer called Machine Learning and Computational Sciences. You heard him in our last episode.

Enoch:

It requires high quantities of highly purified protein. And so, that's why we don't have much training data to begin with.

[Music ends]

JEREMIAH

For context, an assay is a test or evaluation. It costs a lot of money to measure viscosity. So, scientists typically run assays for viscosity toward the end of a drug discovery process. But viscosity proved to be a challenging property to get right, which called for running even more assays.

Enoch:

It required multiple rounds of production and screening, in order to find a molecule that had sufficiently low viscosity while maintaining the other properties of interest such as affinity.

JEREMIAH

This refers to the strength of attraction between two substances, like the strength of an antibody binding to a target protein.

Enoch:

The principles of design for affinity are much well understood. So something you can design for and screen for relatively easily. Viscosity is a slow assay to run, and we could not reliably design for it, which means that you had to use trial and error, which means that often required multiple rounds before you actually find something that had sufficiently low viscosity. And so, the challenges that we had was in the context of a target that we unfortunately encountered a molecule that had great affinity, but unacceptably high viscosity.

[Music]

JEREMIAH

Enoch and his teams were familiar with using machine learning models to predict properties of small molecules with decent success. So, when they got stuck on the viscosity problem, they thought...hmm, why not apply it here too?

Enoch:

The inspirational moment was realizing that you can essentially leverage the same model architecture for entirely different purposes, which is large molecule optimization for viscosity.

The idea is, oh, what if we have a computational predictor where we're adapting an algorithm that was designed for computer vision towards solving this problem of viscosity, and it had to work in a context where we didn't have a lot of training data, maybe 125 examples. So, that pales in comparison to what we have in the small molecule where we have 500,000 structures and 2 million data points.

JEREMIAH

To do this, they used a type of neural network called the convolutional neural net. It was created by French-American computer scientist Yann LeCun to recognize patterns in images through the application of learned filters and parameters, but hadn't been tested in the particular context Enoch's team would use it for.

[Music ends]

Enoch:

The idea is by using the convolutional neural net, that the machine learning algorithm is able to discern those very subtle spatial relationship between atomic charges that result in higher load of viscosity.

JEREMIAH

Here, the convolutional neural net combs through molecular data, and looks for patterns in what causes viscosity.

Enoch:

And we developed this convolutional neural network in order to predict molecules that you haven't made and ask the model, is this likely to be viscous or not? So now we have a model, and we're gonna apply it.

JEREMIAH

Enoch and his team used the model to filter for parts of the drug molecule that had a favorable affinity and structural integrity, but could still be altered in terms of viscosity. The goal was to get low viscosity molecules, while leaving the other properties unchanged.

Enoch:

So we selected eight molecules that were predicted to have low viscosity and some seven of them had very good agreement model with only one misprediction and one of these seven then went on to be a clinical candidate. And it turns out that is actually quite predictive because in the example I gave, seven out of the eight cases were able to predict correctly that they were more favorable in terms of viscosity.

JEREMIAH

The process Enoch described happens far before a potential drug would ever be tested in a clinical setting.

[Music]

Scientists go through numerous, painstaking rounds of trials to ensure that a drug candidate has every property it needs. With the help of AI software, they can better predict that it will.

Enoch:

So in effect, what we've done was done in one round of design and optimization, something that would've taken at least two or more rounds in traditional approaches because we couldn't design for it. So the best case was to hold onto what we had, make some changes, and then rescreen them. And that you may not have gotten in the second round, so there'd be a third round. But this is the first time that we were able to do it in a single round, which eliminated multiple cycles of design and testing, months of the discovery timeline, and of course, you

avoided a direct cost of experimentation, which I said was very expensive with each round. So we were able to deliver the clinical candidate months ahead of schedule because of the existence of this model.

[Music ends]

JEREMIAH

Behind AI software's ability to make predictions, is its remarkable capacity for juggling many drug properties at a time, easing the load on researchers and scientists. The discovery and pre-clinical development phase for one drug can take several years. So any amount of time saved—like a few months time—is valuable.

[Music]

With the help of machine and deep learning methods, scientists have started to be able to address many of the tricky problems in drug design and expedite their workflow tremendously. Despite this, there are still challenges that this technology presents.

Beth:

Some of the easiest-to-use tools are tools that have invested hundreds of thousands of dollars in really smart software engineers making tools really easy to use.

JEREMIAH

That's Beth Cimini.

Beth:

I'm the Associate Director for Bio image Analysis at the Broad Institute in Cambridge, Massachusetts.

JEREMIAH

A cell biologist by training, Beth is dedicated to teaching other biologists how to use software and making that software better. That means making it more accessible.

Beth:

I wish we lived in a universe where money didn't matter to science, and we could have everybody who wanted to be a scientist and they could afford to do every experiment they wanted to do, but unfortunately budgets are finite and so we always have to prioritize what we wanna do.

So having open source software, which is often though not the same thing as free software, making it so that people can actually afford to not just make the data in the first place, but get as much information as possible from the data, makes it again so that we get the most out of every experiment that we do.

[Music fades]

But open source software in the same way as sort of any infrastructure, it can be infrastructure as code. A single tool can support thousands or tens of thousands or hundreds of thousands of researchers all around the world and make it so that thousands of discoveries that couldn't have been done without it now do happen. And so the return on investment for something like that can potentially be huge.

JEREMIAH

With any new technology, there are always kinks to be worked out. Everchanging features of data like accuracy, relevance, volume, and diversity, all impact the processing reliability of AI.

[Music]

When it comes to the application of AI software in drug design and development, we're very much in the early stages. During the 1960s neural network craze, great minds of the time were envisioning a future where machines could do the work of scientists. We're now living in the time period they could only dream of—but it's not exactly what they had imagined. AI is getting smarter, but it is not replacing scientists.

[Music transitions]

Beth:

So far we're in this golden age where what mostly people have been able to use deep learning for really well is automate the stuff that was tedious and painful and nobody really wanted to do it and it was critical to get right, allowing you to get to actually the thinking about science.

JEREMIAH

AI technology has become a powerful partner, helping scientists do the work that will hopefully lead to a better future for us all.

[Music swells]

Next time on Science Will Win:

With the door to the future of AI in drug discovery wide open, we'll explore some of the industry's pioneering innovations with the people leading it—like Pfizer's Charlotte Allerton.

Charlotte Allerton:

I feel we sit at the cutting edge and we have great efficiency coming from those tools. I think the next wave, which we and many others are already working hard on is the explosion of biology data that has come about over the last 10, 15 years in terms of what you can measure in a human cell, the different omics technologies, the single cell capabilities where you can look in an individual cell at individual cellular components. It's a huge innovation opportunity right across the industry.

JEREMIAH

Science Will Win is created by Pfizer and hosted by me, Jeremiah Owyang. It's produced by Wonder Media Network. Please take a minute to rate, review, and follow Science Will Win wherever you get your podcasts. It helps new listeners to find the show.

Special thanks to all of our guests and the Pfizer research & development teams. And thank you for listening!

[Music fades]