



perform multiple computations at the same time, making them very effective for solving math problems. This innovation came from a computer scientist working with VLSI technology in the 1980s named H.T. Kung. Then a faculty member at Carnegie Mellon University and now at Harvard University, Kung gave a talk at the first VLSI conference.

“He figured out how to multiply whole rows of numbers, not just two at a time, on the VLSI chips,” Mead explains. “It’s called matrix multiplication, and it allowed for parallel processing. The idea was later rediscovered by NVIDIA and turned into GPUs.”

NVIDIA, the world’s leading developer of GPUs, also has its share of Caltech influences, including Bill Dally (PhD ’86), a former Caltech professor who is now the company’s chief scientist and senior vice president, and Anima Anandkumar, Bren Professor of Computing and Mathematical Sciences, who previously served as the company’s senior director of AI research from 2018 to 2023.

Like Hopfield, Anandkumar says physics inspires her work.

Even before Anandkumar joined Caltech in 2017, she says she “was fascinated by physics.” In 2011, she analyzed how the success of learning algorithms is tied to the phase transition in the Ising model, the same model upon which Hopfield built his network. “Hopfield gave us the starting tools for modern AI,” Anandkumar says.

Building Bridges Between Brains and Computers

Hopfield points to Mead as an early believer in his vision for neural networks. “Carver had me give a talk in the 1980s where people from Bell Labs would be,” Hopfield says, “and I remember thinking, I don’t know what to tell these people. Then I realized I could simply prove the theorem for the Hopfield network. The original proof is written on the back of hotel stationery that I still have.”

Vogt, the Caltech provost during this time, also believed in the viability of Hopfield’s efforts and ultimately green-lighted the formation of the CNS graduate option. “I don’t think CNS would have gotten going for another

year or two if it hadn’t been for Robbie Vogt,” Hopfield says. “He was a different kind of leader. He could do marvelous things.”


Hopfield saw CNS as a means for people with different backgrounds to converse and influence one another’s work, though he notes it was difficult to get both the Physics of Computation course and the CNS graduate option launched at Caltech. Other scientists, he says, were not convinced of the merits of the interdisciplinary effort.

“Before CNS, there was a clear gap between computer science and neurobiology,” he says. “The gap was something like having a set of people working on weather, and another set of people working on molecular physics and chemistry but having no one asking what the relationship was between weather and the molecular collisions, which were obviously at the bottom of it. The quality of the CNS incoming students was so high that the neurobiologists and engineers who had been skeptics rapidly became true believers, or at least willing participants.”

Today, the nearly 40 faculty members associated with Caltech’s CNS graduate option continue to study the human brain as a computational system both to develop new AI tools and better understand the brain’s own fundamental workings.

At a 30th anniversary celebration for the program held in 2017, many graduates of the program remembered the excitement of crossing boundaries between fields. Gabriel Kreiman (PhD ’02), a professor at Harvard Medical School, spoke at the event and credited the program’s rigor and collaborative nature for producing great science and scientists.

“The intellectual freedom to get together and go with all the other CNS people to the Athenaeum to have lunch and then spend three hours discussing the minutiae of one particular problem, or staying until the wee hours in one of the rooms where we have all of the computers and working together and fighting together about absolutely every problem in neuroscience and computational neuroscience...” said Kreiman at the event, “the magic, the spark of what happened here at CNS was completely unique.”



From top: Bill Dally (PhD ’86) and Anima Anandkumar.

Seeing the WHOLE BOARD

How Caltech is using AI to advance scientific discovery.

By Neil Savage

The lab of Caltech mathematician Sergei Gukov built an AI program that attempts to solve pure math problems as if it were playing a game like chess while improving along the way.



Some five years ago, mathematician Sergei Gukov began teaching himself how to build the neural networks that are the foundation of artificial intelligence, simply to see whether they might be useful in the realm of pure mathematics. He was, he admits now, skeptical whether the supremely complicated and complex questions being asked by pure math would be within the reach of AI's ability to process information.

Today, he says, he is no longer a skeptic, thanks to the time he has spent learning how to build those neural networks, in what he believed was an effort to show that they were irrelevant to the kind of work he was involved in. "When we think about people working on pure math, we usually think of someone sitting in the attic and proving theorems that are so esoteric that no other human can understand them, let alone machines," says Gukov, Caltech's John D. MacArthur Professor of Theoretical Physics and Mathematics.

Instead, he proved to himself that the opposite is true. To understand why, Gukov says it is important to recognize that solving hard math problems can be thought of as a sort of game. These problems involve assertions that mathematicians believe should be true, and their challenge is to prove that they *are* true. In other words, all these problems are essentially a search for the path from A to B. "We know the hypothesis, we know the goal, but connecting them is what's missing," he says.

What makes these problems so hard is the number of steps from A to B. Whereas an average game of chess lasts about 30 to 40 moves, these problems require solutions that take a million or more steps, or moves. After studying neural networks, Gukov realized he could build an AI algorithm that learns to play the game better, or solve a particular problem, as it competes against itself. The program starts

Seismologist **Zachary Ross** (left) works with computer scientist **Yisong Yue** to develop AI visualizations of underground seismic activity.



with knowledge of its existing conditions, a set of rules about what moves it can make and a definition of what it means to win. Then it uses a machine-learning technique called reinforcement learning, similar to the way a person might train a dog to sit, in which the computer tries a move and gets feedback on whether it is closer to its goal.

Gukov and his colleagues recently used this approach on a decades-old problem known as the Andrews-Curtis conjecture. They did not actually solve that problem, but they managed to disprove two sets of potential counterexamples that, had they been true, would have disproved Andrews-Curtis. Though many of these math problems seem intractable now, Gukov says that in 10 years finding those paths from A to B could seem as simple as a computer winning a game of chess, a feat once considered nearly impossible. "I don't know how likely or unlikely it is, but there is definitely a chance that AI can master that kind of problem," he says. "It develops new strategies that are better than what humans can do."

Across the Institute, researchers are learning that AI can help them to think bigger and do more. They are confident about the varied roles the technology can play in scientific research, whether that involves crystallizing mountains of data into new and useful insights, uncovering patterns in data too subtle for humans to notice, or using the power of neural networks or machine learning to streamline experimentation and create new knowledge, develop new therapies, or understand the world and its complex systems in new ways.

Anima Anandkumar, Caltech's Bren Professor of Computing and Mathematical Sciences, says Caltech has taken a leading role in imagining and implementing new uses for AI. "I've been at Caltech since 2017, and since then, I've had so many collaborations, including with other faculty who work on AI, and interdisciplinary collaborations across campus," she says. "The small size of Caltech and its open mindset has made this possible much more than in other, bigger universities. The impact of it can be directly seen in the work that has happened since."

The Institute has infrastructure in place to assist labs as they make this transition. Anandkumar, along with Yisong Yue, professor of computing and mathematical sciences, run AI4Science, an initiative launched in 2018 that helps scientists across the Institute discover which AI tools and resources might help advance their work. The program is a collaborative effort with the University of Chicago and has support from the Margot and

Tom Pritzker Foundation. Seven years on, Anandkumar and Yue have partnered with principal investigators from a wide variety of disciplines on a dozen projects in which AI has helped to make significant advances. "Fundamentally, AI is already transforming the whole scientific method," Anandkumar says.

As part of AI4Science, Yue teaches a class that provides PhD students with a basic understanding of how AI works and what kinds of AI tools they can incorporate into their research. One point he stresses is the need to start with high-quality data. "Data is the fuel for AI," Yue says. "AI converts that data into this model from which you can extract knowledge." If researchers lack good data, he helps them find it, whether that entails digging it out of the scientific literature, using a computer to generate simulated data, or more effectively leveraging the collected experimental measurements.

AI to Model the Physical World

The wildfires that swept through the Los Angeles area in January, touched off by an extreme wind event, reinforced the importance of timelier weather forecasts. Anandkumar and her colleagues are working on AI technologies to create those forecasts and potentially save lives in future natural disasters.

Existing weather-forecasting models are based on complex mathematical equations that describe the physical processes governing how Earth's atmosphere and oceans behave. They are fed by observations of current conditions, such as temperature and humidity, and are run on enormous supercomputers that cost hundreds of millions of dollars. Anandkumar runs simulations based on the same observations, but she skips all the math and instead trains her neural network using historical weather forecasts. The AI then looks for patterns in how those old forecasts played out and applies those patterns to new weather measurements to make its prediction. Anandkumar's AI can run on a single graphics processing unit like those in a home gaming PC, but her results are just as accurate as those generated by a supercomputer.

Anandkumar's system uses 50,000 samples of historical weather data gathered at six-hour intervals over the past four decades. For a neural network, that's not a lot of data, so to extract more value from it, Anandkumar uses neural operators, which are tools developed by her lab. Based on rules about physical processes, such as fluid flows and the conservation of mass, neural operators

Decoding AI Lingo

take discrete data points and create continuous mathematical functions, allowing researchers to examine factors in a system, such as its fluid dynamics, at varying scales to provide a wider view of what is happening than would be available with discrete data points.

For extreme weather events, the system can make accurate predictions days earlier than standard forecasts. For instance, when Hurricane Lee was brewing in the Atlantic in September 2023, Anandkumar used her test model to create a forecast, 10 days in advance, of when the storm would make landfall in Nova Scotia. Meanwhile, the standard European and US models were producing plots that had it heading out to sea.

It is important to not only have accurate early predictions for these severe and dangerous weather conditions but to understand the uncertainty in those predictions. Officials can decide what actions to take based on the level

of certainty; they might react differently to, say, a prediction of a hurricane's landfall with a 90 percent confidence score than to one with 60 percent confidence. Because AI can produce predictions fast and cheaply, Anandkumar's team can create thousands or even millions of simulations, each with slightly different conditions, while the traditional forecasts on supercomputers produce only a few dozen.

With those large numbers, she can average the forecasts to see which outcome appears most often, providing calibrated forecasts necessary for early intervention in extreme situations. She hopes AI can do the same for other weather conditions like the winds that drove the

LA fires. "If predictions had been done even earlier than they were, with confidence levels conveyed to the public, people perhaps could have started fireproofing sooner," Anandkumar says.

The same approach can be used to model other turbulent systems such as plasma flow inside a nuclear fusion reactor, allowing for real-time predictions of whether the flow could damage the reactor or continue toward fusion. This could

give scientists the ability to make on-the-fly adjustments to achieve successful ignition. Elsewhere, it could permit fast adaptations to air turbulence, enabling firefighters to use a drone to monitor and even combat a blaze in conditions in which human pilots would be grounded.

While Anandkumar focuses on weather forecasts spanning days, Tapio Schneider,

Artificial intelligence:

An application or machine that mimics human intelligence.

Algorithm:

A set of instructions or sequence of steps that tells a computer how to perform a task or calculation. In some AI applications, algorithms tell computers how to adapt and refine processes in response to data, without a human supplying new instructions.

Machine learning system:

A machine or program that is fed and trained on existing data and then is able to find patterns, make predictions, or perform tasks when it encounters data it has never seen before.

Neural network:

An interconnected set of processing units, or nodes, modeled on the human brain, that is used in deep learning to identify patterns in data and, on the basis of those patterns, make predictions in response to new data. Deep learning involves applications or machines that use complex networks of instructions to carry out sophisticated tasks.

Caltech’s Theodore Y. Wu Professor of Environmental Science and Engineering, uses AI to tackle climate modeling, which covers centuries and includes scenarios that do not exist in historical data. Current climate models do not accurately capture the way ocean turbulence distributes heat or the effects of turbulence within clouds, both of which affect the climate system. Clouds, in fact, account for more than half the uncertainty in existing climate models, Schneider says, because there is no observational technology that can directly measure what is going on within the clouds, such as how surrounding air mixes in. That means there is insufficient data on which to train an AI model directly.

To deal with the gaps in the data, the researchers are developing individual models of those small-scale processes that they can add to existing models to create the big picture. Schneider and his team understand the physics of those processes and can use the data they do have—on temperature, humidity, cloud cover, and the like—to create simulations of these inner processes. They use those simulations to pretrain a physics/AI hybrid model and then feed actual observational data gathered from satellites, ground sensors, and ocean buoys to fine-tune the pretrained models, making them more accurate. “If you use Earth observations alone, there isn’t quite enough information in them to learn about the turbulent processes directly,” Schneider says. “But if you have a good pretrained model, then fine-tuning with Earth observations seems to work.”

Zhaoyi Shen, a senior research scientist in Schneider’s group, has created a library of about 500 such cloud simulations for different climates around the world, along with varying versions of climate models based on separate assumptions. The lab is collaborating with Google to expand the database to thousands of simulations for potential use by other climate modelers. Meanwhile, a graduate student in Schneider’s group, Andrew Charbonneau, has built a model that uses AI alone to predict snow thickness based on environmental parameters like humidity, further refining the larger climate models.

Limited data is not a problem for seismologists like Zachary Ross, professor of geophysics, whose earthquake models can take advantage of data going back decades as well as so much new data that researchers cannot analyze it all on their own. “We have hundreds and hundreds of sensors across California that are sending back data every second,” Ross says. “It would be totally impossible for humans to do this kind of work entirely by hand.”

Ross and his colleagues use that wealth of data to generate AI-based computer models of what is happening underground. They can even visualize a network of faults based on readings of how seismic waves spread through the ground—a technique the researchers used to discover

Chemical engineer and Nobel laureate **Frances Arnold** (below) uses generative AI to create new gene sequences for enzymes she and her team engineer in the lab.



that the southern part of the San Andreas Fault is tilted rather than vertical. Fault orientation affects the pattern of shaking seismic waves can produce. This knowledge, in turn, can be applied to building codes, so that homes can be constructed to withstand expected tremors without specific requirements that lead to expensive overbuilding.

Those same algorithms can apply to other technologies that are informed by wave mechanics, such as radio, optics, and imaging inside the human body. “Today, nearly every step of what my research group does has AI components in it at some level,” Ross says. “AI has changed almost every aspect of our work.”

Ross’s seismic studies also feed into Yue’s research, which focuses on understanding and improving AI itself. By collaborating with Ross to create models that visualize underground seismic activity, Yue can look at how well the AI system lives up to expectations and where it falls short. Yue has also collaborated with Katie Bouman, associate professor of computing and mathematical sciences, electrical engineering and astronomy, to create an AI system that turned astronomical observations into the first image of the supermassive black hole at the center of the Milky Way galaxy. “Being able to work on these projects gives you a sense of what the fundamental challenges in AI are,” Yue says.

AI at the Molecular Level

Frances Arnold, the Linus Pauling Professor of Chemical Engineering, Bioengineering and Biochemistry, and director of the Donna and Benjamin M. Rosen Bioengineering Center, now uses neural networks to assist in her work on

directed evolution, an enzyme-creation process that won her the Nobel Prize in Chemistry in 2018.

To generate new enzymes—proteins that can build new chemicals or break down others—Arnold and other protein engineers select gene sequences that encode for enzymes. By mutating and recombining the gene that encodes the enzyme followed by artificial selection for the desired traits, she can use her approach to “breed” biomolecules. Just as in natural evolution, only the fittest versions live to reproduce, but it is Arnold and her team, not the environment, who determine that fitness. If the enzyme is closer to her goal—such as being able to break down plastic—than its ancestor, she continues introducing mutations and searching for the most successful offspring in each generation.

More than a dozen years ago, with the help of then Caltech computer scientist Andreas Krause, Arnold started using machine-learning tools—statistical approaches that do not necessarily involve neural networks—to figure out which gene sequences were most likely to produce the next generation of “fit” enzymes. These days, working with Anandkumar, Arnold relies on generative AI to come up with new sequences.

As it turns out, the same neural-network-powered large language models (LLMs) that enable ChatGPT can also work on other material besides text. These models were originally trained by being fed billions of words—or fractions of words, known as tokens—and then tasked with figuring out the relationships among them. The LLMs use what they learn to then predict the most likely word to follow what comes before. This process can also work on computer code or, in this case, DNA. “Large language models are very obvious to use given all the sequences, the library of evolution that’s collected in DNA databases,” Arnold says.

Using this tool, Arnold says she can envision a day when, at the push of a button, a computer could generate a gene sequence to create an enzyme that performs a desired task without going through the iterative evolutionary process, and that enzyme could then be quickly synthesized by a robot. “I’ve been wonderfully surprised at the insights that machine-learning algorithms get from mutational data that were not obvious to me with a human brain,” Arnold says.

Arnold has also become aware of how optimizing these technologies requires changes in lab and data-collection methodologies. “You can’t get the right data for training models without improvements in the experimental method,”

Autonomous system:

A system in which a machine makes independent, real-time decisions based on human-supplied rules and goals.

Automated system:

A system in which machines execute repeated tasks based on a fixed set of human-supplied instructions.

Model:

A computer-generated simplification of something that exists in the real world, such as climate change, disease spread, or earthquakes. Machine learning systems develop models by analyzing patterns in large datasets. Models can be used to simulate natural processes and make predictions.

she says. To that end, her lab has developed a method for sequencing the genes that code for the proteins they are studying, matching the sequences to the functions of those proteins. This data is then labeled with a type of barcode optimized for use by computers. “We’re now changing the way we do the experiments in order to make use of the power of these data-driven methods.”

Similar approaches may be relevant to the pharmaceutical industry, where figuring out how to make, for instance, a particular cancer drug is one of the biggest expenses driving up costs. Hosea Nelson, a professor of chemistry at Caltech and a principal investigator in the National Science Foundation’s Center for Computer Assisted Synthesis (C-CAS), wants to go beyond proteins to figure out how to synthesize any chemical imaginable—a feat that could drastically decrease the amount of money required to find an effective and safe medication.

To make new drugs, or any other chemicals, chemists need to figure out the right reactions to use. In essence, developing any chemical reaction is like creating a recipe from scratch. But with all the variables involved—different ingredients (and quantities of each), the sequence in which to add them, how long to cook them and at what temperature—millions of possible recipes could exist, and it could take human chemists hundreds of years to explore them.

Instead, C-CAS principal investigators such as Nelson and Sarah Reisman, Bren Professor of Chemistry and the Norman Davidson Leadership Chair of the Division of Chemistry and Chemical Engineering, spend a week or two in the lab physically cooking up a more manageable number of possible reactions to create a molecule with specific attributes. They then measure various features in those reactions, such as how much of the molecule each one yields. They feed about 70 percent of the reactions into a neural network, which uses the data to determine the pattern of features most likely to produce the desired result. “What we’re interested in is using AI to uncover relationships that allow us to have a better understanding of a chemical reaction,” Nelson says.

Cancer research can also benefit from the support of AI. For instance, a number of efforts in the field are focused on coaxing the body’s own immune system to attack tumors. This works, says Matt Thomson, professor of computational biology, when T-cells—one component of the immune response—are able to infiltrate the tumors. Some tumors, however, manage to evade T-cells. Thomson is looking for ways, perhaps with drugs or gene editing, to reprogram the tumors so the immune system can exclusively target the cancer.

In this work, Thomson employs a technique called seqFISH (sequential Fluorescence In Situ Hybridization), developed in the lab of Long Cai, a Caltech professor of biology and biological engineering. SeqFISH uses fluorescent probes that attach to and illuminate DNA, mRNA, and proteins in cells, providing a detailed readout of their makeup. Once he knows exactly what is in the cells, Thomson asks a neural network to predict how altering the DNA or proteins would change the behavior of the cells or of the tissues they make up.

To that end, he and his colleagues recently unveiled Morpheus, a deep-learning neural network that predicts how to alter a tumor to make it more susceptible to immune therapy. One strategy identified by Morpheus involves altering the amount of particular proteins that were expressed by three different genes, turning up the expression in two while turning it down in one. The AI predicted this would allow T-cells to enter tumors they could not previously penetrate. Morpheus has suggested alterations for tumor cells in both melanoma and colorectal cancer, and Thomson's group is seeking funding to work with a clinical partner to apply the computer's results in clinical research. A similar approach could lead to treatments for other diseases as well.

"The real advance is that the AI system can look at lots of data from human tumor samples, and then it can integrate that information to make coherent and very specific predictions about therapies," Thomson says. It would be hard enough for humans to figure out what reprogramming just one of each of the 30,000 genes in a cell would accomplish. Looking for all combinations of three genes would entail sorting through 27 trillion possibilities. "How would a human ever look at that data to get a picture of what's going on and design therapies?" he says. "It's impossible, but we can develop AI systems that can do the job in about a day."

AI to Understand the Brain

Colin Camerer, the Robert Kirby Professor of Behavioral Economics and Leadership Chair and director of Caltech's Tianqiao and Chrissy Chen Center for Social and Decision Neuroscience, uses AI to garner insights about how people make decisions and form or break habits. The field of economics has traditionally tackled those questions by watching what people buy or having them respond to questionnaires. Camerer enhances these techniques by adding in more-objective measures, such as eye tracking to see what people are actually paying attention to, and functional magnetic resonance imaging (fMRI) to see which parts of the brain light up when people focus on a

Big data:

The massive amounts of data that come in quickly and from a variety of sources, such as internet-connected devices, sensors, and social platforms. In some cases, using or learning from big data requires AI methods. Big data also can enhance the ability to create new AI applications.

Generative AI:

Deep learning networks, such as large language models (LLMs), that can recognize, summarize, translate, predict, and generate content using large datasets.

particular choice. The latter effort got a boost in 2003 with the launch of the Caltech Brain Imaging Center. "The idea has been to take a very central thing that economists have studied in a certain way and try to study it with a fresh eye and with better machinery," Camerer says.

By mapping what is happening in literal neural networks as people play the standard economic games used to discover how subjects make choices, the researchers can analyze objective measurements instead of relying on subjective reports. But it can be difficult to sort out good hypotheses from spurious ones without the help of AI. "What machine learning is really good at is taking a lot of possible predictor variables and winnowing down the ones that really are solid to make good predictions," Camerer says.

Recently, Camerer and his team created a machine-learning algorithm to see if they could tell how long it might take someone to develop a habit of going to the gym or for a health care worker to get in the habit of

handwashing. Although they found there was no magic number, they discovered that gym attendance took about six months to become habitual whereas handwashing took only about six weeks. The algorithm sorted out which variables were important: Most months had no predictive value for someone going to the gym, although there was a decrease in December and an increase in January. But the day of

the week did, with Monday and Tuesday being the likeliest days. The best predictor was how many days had elapsed since someone had gone to the gym.

What Comes Next?

While Caltech scientists recognize the promise of AI in reshaping how they do their research—and the questions they are able to answer—they caution the public about assuming that they are simply turning over their labs to a computer. To take full advantage of the promise of AI, Nelson says, requires researchers and students who are willing and able to explore what works and, more importantly, what does not. "There's a lot of problem-solving and technical skills that go into what we do," he says. "It's very physical."

Arnold adds that a primary benefit of AI is that it allows researchers the freedom to explore and imagine. It then provides support to fill in the more data-driven details. "It's a new tool that makes much of our work easier," Arnold says, "and I hope in the future will make it very straightforward to design these new catalysts that evolution hasn't cared about but would be useful to us." 📺



Caltech researchers navigate AI's shifting and multifaceted future, charting a course for its ethical development and application.

By Julia Ehlert Nair