

Family Membership Through Generations

The benefits of belonging to CEFCU last a lifetime!



Putting members first for 75 years.


CEFCU exclusively serves the extended Caltech community and their families. Membership is open to Caltech, Jet Propulsion Laboratory, and The Huntington employees, contractors and volunteers, as well as Caltech undergraduate and graduate students, the Caltech Associates, Caltech Alumni Association, CARA, TMT, Polytechnic School, and immediate family (by blood or marriage) of current members.

800/592-3328 • www.cefcu.org



Must qualify for CEFCU membership to join. Minimum \$5 deposit and one-time \$5 membership fee due upon opening any CEFCU share account. **Federally insured by NCUA.**


The Many Facets of AI





When you look into a kaleidoscope, play a game of chess, or open an Erector Set, you are unlocking new perspectives and possibilities. Each twist of a hand, move on a board, and connected piece represents a single step in a journey that could produce any number of different outcomes. In much the same way, artificial intelligence draws on and reconfigures data, **generating seemingly limitless insights** while simultaneously **mirroring the biases** we put into the system. No longer a mere curiosity, AI has emerged as a **foundational tool of the educational, research, and societal landscape.**




Like the rest of society, Caltech's scientists, educators, and administrators are navigating this new reality. But they are not passive observers; they are **pioneering efforts to harness the technology in innovative ways** as they seek to **solve fundamental scientific mysteries** and **forge new lines of inquiry.** Meanwhile, they are **weighing the implications** of AI's increasingly ubiquitous presence. The



three features that follow seek to illuminate this evolving journey on campus, while also tracing how the Institute helped enable this paradigm shift.





It Together

Tracing the roots of neural networks, the building blocks of modern AI, at Caltech.

By Whitney Clavin

In the early 1980s, three giants of Caltech's faculty—Carver Mead (BS '56, PhD '60), now the Gordon and Betty Moore Professor of Engineering and Applied Science, Emeritus; the late Nobel Laureate in Physics Richard Feynman; and John Hopfield, then a professor of biology and chemistry, who would go on to win a Nobel Prize as well—became intrigued with the connections between brains and computers. The trio would gather for lunch at Caltech's Athenaeum and wonder: How do our brains, with their billions of interconnected neurons, process information? And can computers, which work in more straightforward number-crunching ways, mimic the brain's ability to, in essence, think?

These conversations ultimately led to a new graduate-level course, The Physics of Computation, intermittently co-taught by all three professors from 1981 to 1983. Hopfield recalls that there were obstacles to getting the course off the ground. "Back then, there was very little interaction between computer science and other fields,"

says Hopfield, now an emeritus professor at Princeton and Caltech's Roscoe G. Dickinson Professor of Chemistry and Biology, Emeritus. "We had a diffuse mixture of ideas we wanted to present, and it took a while for us to convince the Institute to approve the course. Still, it was an intellectually exciting period and brought great new students and many guest lecturers to Caltech."

Alongside these conversations, Hopfield began to formulate ideas for creating simple networks that mirror the way human memory works. In 1982, he published a theoretical paper describing how an artificial neural network, modeled after the structure of the human brain, could be programmed to learn and to recall. Though other researchers would later build these networks using computer chips, Hopfield's research used math to describe a new biology-inspired scheme that could be trained to "remember" stored patterns, such as images. The computers could recall the images even when only an incomplete or fuzzy version of the same image was available. The system is akin to someone remembering the full experience of



Clockwise from top: Carver Mead (BS '56, PhD '60), John Hopfield, and Richard Feynman.

hearing a particular song after catching a snippet of the tune on the radio.

The roots of modern AI programs like ChatGPT can be traced to biology-inspired models similar to the Hopfield network, as it is now known. For this seminal research, Hopfield was awarded the 2024 Nobel Prize in Physics, together with Geoffrey Hinton of the University of Toronto.

Hopfield's breakthrough came during a pivotal time in Caltech's history when ideas had just begun to flow between neuroscience and computer science. "AI research was developing very slowly and still had many doubters," Hopfield says. Despite the challenges, Hopfield sought to make the movement more official:

In addition to the Physics of Computation course he helped run, he sought to organize a new interdisciplinary program offering graduate-level degrees.

Caltech's then-provost Robbie Vogt, now the R. Stanton Avery Distinguished Service Professor and Professor of Physics, Emeritus, supported the idea, and, in 1986, the Institute's Computation and Neural Systems (CNS) program was born, with Hopfield as its first chair. Today, CNS comprises a vibrant group of scholars that has produced more than 150 PhD graduates.

"This program was the first of its kind that took in highly quantitative students from physics, engineering, and mathematics who were interested in both brains and computers," says Christof Koch, who served as the first faculty member hired in CNS and later on as a chair of CNS, before leaving Caltech in 2013 to become the chief scientific officer and president of the Allen Institute for Brain Science. "Now there are many other places that similarly look at brains as computational systems, but we spearheaded the effort."

Yaser Abu-Mostafa (PhD '83), a professor of electrical engineering and computer science at Caltech who did theoretical work on Hopfield networks in the 1980s, recalls that by the middle of that decade, more and more people were joining the growing

AI community worldwide thanks to the innovative work being done on campus. "What Hopfield did was very inspirational," he says. "It established in people's minds that this can be done." Abu-Mostafa initiated an AI-themed workshop, which later led to the creation of the Neural Information Processing Systems conference in 1987. Now known as NeurIPS, the gathering has become the largest AI conference in the world. (See page 13.) "It has been very rewarding to watch a field forming from scratch," Abu-Mostafa says.

Built on Physics

In the late 1970s, Hopfield, then a biophysics professor at Princeton University, attended a series of neuroscience lectures in Boston and quickly became fascinated with the topic. As a condensed matter physicist by training and the son of two physicist parents, he wanted to understand how our minds emerge from the complex network of neurons that make up human brains. "I was very interested in the interface of physics and living matter," he says.

In 1980, Hopfield left Princeton for Caltech in part due to the Institute's "splendid computing facilities," which he would use to test and develop his ideas for neural networks. However, Hopfield did not set out to create an artificial intelligence. "I was hoping the networks would tell us how the brain works," he says.

His idea was to build a simple computer program based on the vast network of billions of neurons in our brain and the trillions of connections among them. Computers of the 1980s were used to execute

long sequences of commands and search databases for information, but that process took time and required increasingly large amounts of storage space. Imagine trying to remember the name of a singer and having to comb through a catalog of all the singer names in your head one by one—it might take a while.

Instead, our brain has a more efficient system of retrieving information that relies on neurons changing their architecture as they learn new connections. Memories are encoded in different patterns of neural activity; as Hopfield says, the brain is a dynamic biological system.

He decided to model his neural network on another dynamic system in nature involving magnetism. Called the Ising model, the system describes how the up or down spins of electrons in a material can influence one another and spread magnetized states. When this occurs, the system evolves toward the lowest-energy state, like a ball rolling down a hill.

Hopfield networks also evolve toward low-energy states in a mathematical sense. These neural networks are composed of artificial neurons connected via nodes, with each connection having a different strength, or weight. A set of computer codes, known as an algorithm, directs the network to tune the connection strengths between these neurons such that a stored image, like that of a spider, becomes linked to a particular low-energy state. When a fuzzy image of a spider is fed into the Hopfield network, the network's artificial neurons assess the available information and then adjust their activity levels by evolving toward the low-energy state matching the stored image. In this way, the system learns to recognize images of objects.

The backbone of any neural network is an algorithm (or learning rule); a key feature of Hopfield's algorithm, says Abu-Mostafa, is that it allowed the system to learn and grow increasingly smart. "Learning is absolutely essential to intelligence," he says. "Hopfield extracted the essence of neurons." Abu-Mostafa notes that the theoretical paper Hopfield published in 1982, "Neural networks and physical systems with emergent collective computational abilities," is the fifth-most-cited Caltech paper of all time.

Physics played a key role in Hopfield's success, Koch says, and this "led to a massive influx of physicists into the field."

"Hopfield figured out how to mold the energy landscape [a map of the possible energy states of a system]. His network was trained to dig a hole in the landscape corresponding to the image pattern being trained," adds Erik Winfree (PhD '98), professor of computer science, computation and neural systems, and bioengineering at Caltech, and a former CNS student of Hopfield's. "He brought physics to the networks."

In Hopfield's Nobel Prize lecture in December 2024, he explained how the Ising model of magnetism could be generalized to replicate a biological system like the brain. "Everything really came together when I saw these two parts of science are really described by the same set of mathematics," Hopfield said.

Mead adds that others had attempted to build artificial neural networks before, but few could envision them scaling up to the sizes needed to perform interesting tasks. "Hopfield showed that they were possible," he explains. "This was the first time people started thinking that neural networks might be useful."

How Computers Caught Up



Around the time that Hopfield was working on the theory behind his neural networks, Mead and his collaborators had begun to transform the computer industry by inventing a new way to pack more of the tiny semiconductors known as transistors onto computer chips, a process called very large-scale integration (VLSI). VLSI allowed millions, and now billions, of transistors to squeeze onto single chips, a feat that enabled the development of desktop computers, cell phones, and myriad other computing gadgets.

In the early 2010s, researchers realized they could use a type of VLSI chip employed in video games, called graphics processing units (GPUs),

to handle the huge computational demands of AI networks. Though GPU chips were not invented at Caltech, some aspects of their origin can be traced back to the early VLSI research on campus.

A key feature of GPUs, which makes them critical for large AI neural networks, is a type of computing called parallel processing. Essentially, this means they can

From left to right: Robbie Vogt, Christof Koch, Yaser Abu-Mostafa (PhD '83), and Erik Winfree (PhD '98).





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.

