How the program in Computation and Neural Systems at Caltech began

John J. Hopfield December 2013

Computation and Neural Systems (CNS) has become a successful and evolving part of Caltech and its intellectual landscape, and now seems an obvious thing to have done. This is the story of how it began and how its culture originated. For those who are or who will be part of university faculties, there are a few lessons here about how not to get something started. Obviously this is a very personal view, and in common with all histories written by participants, it should be regarded as one version of the truth.

The story begins in 1977. At that time I was a physics theorist who, after 15 years of research in conventional solid-state physics, had recently taken up questions arising in the physics of biological molecules and processes. I spent the winter of 1977 as a visitor at the Bohr Institute/Nordita in Copenhagen. This physics institute made a sporadic but continuing outreach toward biology. I arranged many broadening interfacial seminars for them, but disappointingly found no new science problem for myself, although that was for me the real purpose of my visit.

Shortly after my return to Princeton without a new problem to work on, Francis O. Schmidt descended on me. He ran an entity called the Neuroscience Research Program (NRP) which chiefly held small meetings attended by 20 regular members of the program and 20 outsiders chosen for the special topic under consideration at that meeting. He apparently had gotten my name from relativist John A. Wheeler, who (for reasons that I have never grasped) had always been one of my staunch supporters. Schmidt invited me to talk at the next meeting. I told him I knew nothing of neuroscience. He said that it didn't matter, just speak on what interests you, so I talked about biomolecular accuracy. The audience with whom I was supposed to engage was a superb international set of neurologists, neuroendocrinologists, psychologists, immunologists, electrophysiologists, neuroanatomists, and biochemists. Apparently Frank wanted to add a theoretical physicist to the group in hope that this might help his Program become deeper and more complete. The audience who voted on my membership understood little of what I said, but it didn't matter. The nomination was a put-up job, for Frank controlled everything.

By the end of the three day NRP meeting I had begun to realize that how mind emerges from brain is the deepest question posed by our humanity. It was nominally being pursued by the NRP 'club' of diverse talents and great abilities, but each scientist looked very narrowly at his research area and never discussed broader issues. It appeared to me that this group of scientists could never really engage with the problem of mind and emergence, because the issues can only be expressed in an appropriate mathematical language and structure. None involved with the NRP at this time moved easily in this sphere. So I leaped on Frank's invitation in the spring of 1978 to become a member of the group. My own basic education in neurobiology ensued through attending the semi-annual NRP meetings, sitting next to world experts in their fields who would patiently explain to me what each speaker was really talking about.

I moved to Caltech in January 1980, and a year later wrote a memo to the president of Caltech about the potential relationship between "higher nervous function", described as observable normal behaviors of humans and other animals, and the kind of neurobiology then going on at Caltech. The memo described the gap as something like having a set of people working on weather, another set of people working on molecular physics and chemistry, but having no one asking what was the relationship between weather and the molecular collisions which were obviously at the bottom of it. The memo concluded

> "It strikes me that this direction has as great intellectual impact as any area of science I can name; a potential for interaction with engineers; that its time is rapidly approaching; that the area is not yet 'owned' by others; (in greatest part because of an anti-theory stance by the majority of biology experimentalists); and that Caltech has most of the necessary support structure for such an effort." (full memo appended, which is a shortened version of a considerably longer memo still in my files. The latter includes some of the political/science problems presented by the very narrow view of major Caltech biology faculty.)

You might think that CNS could follow rapidly on the heels of such a memo, particularly since I had written it to Murph Goldberger, a good friend of mine (from the time when we were both in the physics department at Princeton) who was then the president of Caltech. What actually happened?

The result of that memo was a committee that I chaired, set up by the president in March 1982, to look into this direction. First mistake: the committee was appointed by, and reported to, the president rather than the provost. The provost, Jack Roberts, under whom such questions really belonged, had been bypassed and as a result had no interest in the success of the committee. He did not think highly of interdisciplinary thrusts in any event.

The second problem was having a committee chairman without credentials. My interest in information processing and computation at the level of the nervous system was recent and entirely accidental. As of spring1982 I had published absolutely nothing in this area. I was a committee chair with a stirring (to me) vision—but no scientific credentials in the area being advocated. At Caltech this matters. Caltech faculty do not 'do the vision thing'. Only accomplishments have significance.

The biggest error was the failure of Goldberger and me to recognize an essential principle of university management, namely that *the normal function of a broadly representative university committee is to prevent change.* The committee chosen was indeed broadly representative, comprising Seymour Benzer (molecular biology), Richard Feynman (physics), Roger Noll (economics), Jim Hudspeth (auditory hair cells), Mark Konishi (birdsong learning), David Van Essen (primate neuroanatomy), Lou Breger (philosophy) and me (theoretical physics and biology). It was therefore certain to advocate nothing of interest.

The report of the committee failed to make real connections with Caltech's Engineering and Applied Science Division. Carver Mead had declined to serve on the committee. It did not help that its budding Computer Science group in EAS was at war with a group in Physics over issues of parallel computing machines and was opposed to anything that sounded like cooperation with the Physics, Mathematics and Astronomy (PMA) Division. The biological members of the committee really wanted only more support and companions for the biological directions they were already pursuing. Humanities, led by Roger Noll and Lou Breger, had no interest whatsoever into going into psychophysics or modern psychology, but naturally wanted more money and positions. The end of my cover letter on the report despairingly reads

> 'Obviously, a proposal to add faculty in an area must ultimately come from a division or divisions. Division chairs and executive officers must be the major and enthusiastic advocates of such propositions in order for them to be implemented without perversion.'

Sadly, not one of them was interested in leading or even participating in such a direction. (In a separate note to Goldberger necessitated by a failure of the committee to support the kind of enterprise we had been charged with developing, I even facetiously suggested that Murph should appoint Feynman executive officer for Computer Science to provide visionary leadership. Meanwhile the provost—although a tennis buddy of mine—sat on the sidelines and smiled. (He was at the time battling with Wolfram in PMA over copyright and royalty issues, which didn't help his view on the compatibility of computer science and faculty cooperation.) Roger Sperry received a Nobel Prize in 1981 for his work on the brain was a professor in the Biology Division, and might have weighted in, but at this point in his career was utterly non-reductionist and saw no point in studying the properties of neurons. (I happened to be sitting next to him in a Biology faculty meeting at which a new assistant professor in cellular neurobiology was being considered. Sperry abstained from voting, then turned toward me and said very quietly 'I am sure this young man is an outstanding scientist. But he will never discover anything that I want to know'.) So this initial effort to generate an intellectual thrust in the spirit of the later CNS program completely collapsed, disappearing without a trace in late 1982. A few additional sordid tales of this collapse are available in my files.

For 1981-82 Feynman, Mead, and I planned out a yearlong course cross-listed between three Divisions. It was called 'The Physics of Computation'. It was to unite in understanding how real brains function, how VLSI and modern computers are designed and built, the essence of information, coding, and computation, and the physics that provides the underlying microscopic basis for both engineering and biological computers. Carver and I were a bit worried, but felt that with Feynman it could all be carried off in a grand and unified manner. With Feynman's name associated with it, no one could publicly challenge the grandiosity of such an attempt at synthesis, so it was duly approved as a new course designated as PMA/EAS/Bio 250.

Unfortunately, that was a year in which Feynman had one of his heroic struggles with cancer, and he did not participate at all. Most aspects of the subject were far from what

Carver and I knew, and most of the lectures were given by too wide an assortment of excellent invitees. Without Feynman there was little intellectual linkage to connect the diverse topics into a meaningful unity. The kindest term which comes to mind when I remember that course is 'disaster'. Understandably, few students persisted through the entire year. (One of these few students was Markus Meister, who then voted with his feet and did his PhD thesis in bacterial chemotaxis, about as far from a nervous system as you can get. He recently returned to Caltech as professor of biology, and is a computational neurobiologist affiliated with CNS.) Carver and I vowed 'never again', though I did myself learn a huge amount from attending all the lectures.

However, there was some fallout. In the spring of 1983, I helped Feynman give a one-term course relating to the physical limits of computational 'machines'. And by 1985, Mead, Feynman, and I were all giving independent courses containing different fragments of the original Physics of Computation course. Mead had developed a research interest and course in network analog VLSI, Feynman focused on the physics/computer intersection, and I was developing the network computation course that was later to become CNS 185.

In the fall of 1984, a campus+JPL afternoon on 'neural networks' was organized, thanks in greatest part to promotional efforts by (as I recall) Terry Cole, Ed Posner and Dimitri Psaltis. It was repeated in November1985 (see attached program). Having seen Murph Goldberger at lunch at the Athenaeum that day, I invited him to come to the meeting, and he was most impressed by the liveliness and breadth of the group of talks. Two weeks after that, I ran into Rochus (Robbie) Vogt, then provost. I lamented to him the fact that there was a brilliant young man with whom I had shared an office at MIT in the spring of 1984 who would make a wonderful assistant professor at Caltech. But alas, since his research interests lay somewhere between neurobiology and engineering, neither Biology nor Engineering and Applied Science (EAS) would possibly appoint him.

Robbie's reply was roughly this: 'If you can convince me of the quality of this man and the significance that his appointment would have for Caltech, I will solve the political problem of how to get him appointed.' That task proved easy because Koch was in fact stellar—and in addition was a member of the *Studienstiftung*, an elite German student society to which Vogt had also been elected years earlier. There was now a shared enthusiasm between president and provost.

But such an appointment required a home base, a PhD option to which Koch would naturally affiliate. Thus the appointment of Christof Koch became the catalyst for defining a new interdivisional educational program. Mead knew Tommy Poggio, a mentor of Koch, and through that connection also was enthusiastic about getting Koch to Caltech and EAS. Vogt appointed a study committee and stacked the deck, selecting only professors who could by now see the educational and student recruitment benefit of a new PhD option for their own students. The committee members were Geoffrey Fox (physics), Carver Mead (EAS), David Van Essen (biology), and me (biology and chemistry). We of course enthusiastically proposed CNS as a new interdivisional option and thus a home to which Koch could be affiliated.

Easy to get the proposal passed? Biology as a whole was rather dubious; and few of even the neurobiologists were supportive. Engineering was somewhat more supportive, because of the high esteem with which Carver was held. In Physics Feynman's implicit support was important, though of course he had no real interest in institutional or educational structure. It helped also that I had by now written my early 'neural' papers and had become a legitimate advocate in all three Divisions involved.

To illustrate campus views, I quote from physics professor Steve Frautschi. The only reason that I single out Steve, someone whom I greatly respect, is that I happen to possess written documentation of his attitudes. Frautschi was the chairman of the Graduate Study Committee, whose approval was necessary to start the new option. He sent me a list of seven questions. I quote a few from a hand-written memo, which is still in my files.

> 1) Is it really necessary to form a new option? Can't the same goals be achieved within existing options, by redefining their requirements somewhat?
> 2) Fads come and go. Is it in the best long-term interests of students to have a degree with this unusual option label, rather than the well-understood traditional labels?
> 3) We already have a lot of options, and this new one seems rather new and specialized.
> 4) Isn't the organizational device of an option that is not in a single division troublesome and ultimately unstable?

And so on.

There was not one positive word in the memo. Ironically, when acting as chairman of the graduate course of study committee, Frautschi spoke in favor of flexible requirements for existing options as a route for not needing a new CNS option. But when Frautschi sat on a PhD committee as a representative of physics, he became a staunch defender of the purity of conventionally defined physics. When students such as Dawei Dong or Eric Mjolsness did CNS-type thesis projects for a degree in physics, Frautschi strongly questioned whether theirs was really research for which a physics degree should be granted.

Frautschi's words and attitudes were a reflection of the views of the silent majority of Caltech faculty. Happily, Provost Vogt was of the Chuck Colson school of administration (Nixon White House era). 'Once you have them by the balls, their hearts and minds will follow'. And so CNS did come to be founded with 'enthusiastic' support of Division Chairs afraid to oppose Vogt lest it cost them resources, and an appropriate home base for Christof Koch was thereby created.

I volunteered to become its first chair. The story is more complex than that. A few years before I had witnessed an appalling dressing-down of a subordinate by Vogt while he was

PMA Division chair. I remarked on this behavior to a faculty colleague who assured me that it was SOP for Vogt. He apparently viewed humiliation, invective, and emotional tirades as legitimate management tools. I resolved to myself *never* to take an administrative position under him. In one of the committee's last meetings with Vogt, he remarked that a designated chair of CNS would be necessary for Caltech politics, and for the short term the chair would go to Vogt for CNS resources in order that CNS support did not seem to come from existing Divisional budgets. There was one volunteer, Geoffrey Fox. Geoffrey knew from experience that Vogt required a name *now*, and seeing no other volunteers he filled the vacuum. But Fox's background and interests were so far from the envisioned center of gravity of the new option that it was in my view totally unsuitable. I had to prevent a disaster. So I told Vogt and the committee of the episode I had witnessed between Vogt and his staffer and of my personal resolution. I then volunteered to take the position, but added that if Vogt ever raised his voice at me, he had my immediate resignation. He glowered, appointed me, remained a fervent supporter of CNS, and we had civil interactions ever after.

After that, there were two critical items for success. First, Murph Goldberger, being firmly behind the idea, actively helped raise important seed funds from the Parsons Foundation. Tom Everhart when he became president was equally supportive with respect to the Pew Charitable Trust. Second, and *the* most important, there was tremendous interest by even the first group of graduate student applicants. 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.

A Pasadena Star-News editorial written when Goldberger retired as president remarks "Goldberger's proudest achievements at Caltech include strengthening the Institute-JPL relationship and overseeing the extraordinary cooperative development of new computer concepts based on neural networks." This is the usual hash that newspapers frequently make in sci/tech reporting, but clearly Murph was very proud of the innovation that CNS represented.

Not all histories have useful lessons. There might be two in this story.

First, the way to attract good graduate students and good faculty is to produce an educational program with a broad base, but whose center of gravity is located at the center of interest of the students and faculty being recruited. There are many ways to dice and slice science, and sci/tech education should not be heavily reliant on the way that the courses, departments, and programs were organized in 1940.

The other lesson? Moving the intellectual focus of a university is somewhat like moving a graveyard. In both cases you should expect little help from the inhabitants.

TO: M. L. Goldberger ~ December 1980

FROM: J. J. Hopfield

The Mechanism of Higher Nervous Function

"Higher nervous function" describes behaviors measurable in the psychological realm. Examples of such behaviors or psychological properties include:

The abstraction of relevant information from complex patterns of stimuli.
Memory (physically distributed) and recall of experience
Attention

The above properties are emergent and chiefly collective properties of a large number of relatively simple neurons. In the same sense all the phenomena of weather - wind, snowflakes, lightning, clouds, surf - are emergent properties of about 10^{44} simple molecules of oxygen, nitrogen, and water.

There are several essential parts to making progress in this science. First, studies of functional neuroanatomy are very important. To what extent are the detailed connections between neurons exactly determined or statistically determined? What is the relation between simple processing and the connectivity of the neurons? There are several members of the Biology Department whose research has general connections with this area, but tends to be oriented toward qualitative "mapping" and comparitive biology rather than the central problems of higher nervous function.

Second, psychophysics studies at various levels provide essential detailed information on the nature of higher processing. Psychophysics is scarcely new to Caltech, but what must be emphasized here is the attempt to push psychophysics to a place where it can be mathematically modelled and the neural structure responsible for implementing the mode sought.

Third and central, the understanding of higher nervous function requires mathematical modelling and a conceptual theoretical framework with which to examine the experiments and the structures, and from which new important experiments can emerge. This area is now developing, with various interesting and testable models of simple features of higher nervous function now available.

For example, what information is abstracted from a visual pattern? Francis Crick and collaborators would like to claim that a scene is Fourier filtered through a spatial filter less than an octave wide, and that the zero crossings of this filtered image (which, by Logan's theorem, contain all the information of the filtered signal except a normalizing constant) are the fundamental symbols of a visual pattern. There are both anatomical and psychophysical tests of this idea.

One of the measures of the challenge and excitement of an area is the caliber of the men and institutions attracted to it. It is significant that such Nobelists as Francis Crick, Leon Cooper, and Gerald Edelman are all basically involved in working on the conceptual basis of higher nervous function, with strong interactions with (or their own) experimental programs. Both the Salk and MIT appear to be moving toward this direction.

Caltech neuroscience peripheral to these matters is very often excellent, and would complement well a new kernel of interest in the mechanism of higher nervous function. A significant interaction in the computer area also seems possible. The brain is crudely analog; the detailed circuitry of every brain is very different from every other one and changes as a function of time; memory is distributed rather than by location; ambiguity abounds and is arbitrarily resolved. Are there some lessons here for the designer of huge chips or giant computers? And in the other direction, the mathematics of a large ensemble of interconnected highly non-linear neurons becomes intractable. The construction of special VLSI chips could permit the study of the behavior of well-defined models whose mathematics is too tough to approach without the physical insight coming from seeing a behavior.

It strikes me that this direction has as great intellectual impact as any area of science I can name; as well as a potential for interacting with engineers that its time is rapidly approaching; that the area is not yet "owned" by others; (in greatest part because of an anti-theory stance by the majority of experimenters); and that Caltech has most of the necessary support structure for such an effort.

Meeting Announcement

HOPFEST

Friday, November 15, 1985

1:10 p.m. - 5:15 p.m., 24 Beckman, Caltech

1:10 - 1:15	Introduction	E.C. Posner
1:15 - 2:15	Session I - Foundations	R.J. McEliece, Chairman
1:15 - 1:30	"Log and Linear Capacity Loss"	E.C. Posner
1:30 - 1:45	"Capacity with Quantized Connections"	S. Venkatesh
1:45 - 2:00	"Combinatorial Complexity of Neural Networks"	P. Baldi
2:00 - 2:15	"Integrating Storage Capacity in a Content - Addressable Memory"	T. Potter (DEC/SUNY)
2:15 - 2:20	Short Break	
2:20 - 3:20	Session II - Neurobiology	J. Bower, Chairman
2:20 - 2:35	"A Molecular Bistable Switch Present in Synapses in the Telencephalon"	M.B. Kennedy
2:35 - 2:50	"Association Fiber Systems in the Piriform Cortex and Possible Role in Content-Addressable Memory"	L. Haberly (Wisconsin)

ECP 10/31/85

2:50 - 3:05	"Neural Processing in the Cerebellar Cortex"	J. Bower
3:05 - 3:20	"Parallel and Hierarchical Processing in the Primate Visual Cortex"	D. Van Essen
3:20 - 3:40	Break	
3:40 - 4:25	Session III - Implementations	J. Lambe, Chairman
3:40 - 3:55	"Asymmentry and Error Correction"	J. Lambe
3:55 - 4:10	"Holographic Implementation of the Hopfield Model"	D. Psaltis
4:10 - 4:25	"Frequency Multiplexing and Space Invariance in the Hopfield Model"	G. Sirat
4:25 - 4:30	Short Break	
4:30 - 5:15	Session IV - Simulation & Computation	Y. Abu-Mostafa, Chairman
4:30 - 4:45	"A Tale of Three Neurons"	J. Chover (Wisconsin)
4:45 - 5:00	"Escaping Local Minima in Neural Net Optimizations"	E. Baum
5:00 - 5:15	"Decoding at the Analog Level"	J. Hopfield