The Brain Prize 2024

INFORMATION PACK

THE
BRAIN
PRIZE

Founded by the
Lundbeck Foundation

Foundational research in computational and theoretical neuroscience has been recognised with the award of The Brain Prize 2024

This year The Brain Prize worth DKK 10 million (€1.3 million) is awarded to:

Larry Abbott (USA)

Terrence Sejnowski (USA)

Haim Sompolinsky (Israel/USA)

Professor Richard Morris, Chair of The Brain Prize selection committee explains the reasoning behind this year's award.

"Theoretical and computational neuroscience permeates neuroscience today and is of increasingly growing importance. Larry Abbott, Haim Sompolinsky and Terrence Sejnowski have made pioneering contributions to the field and have made seminal discoveries in our understanding of the principles that govern the brain's structure, dynamics and the emergence of cognition and behaviour.

All three candidates originally trained in physics, but they have long worked in neuroscience and have applied novel and sophisticated approaches from physics, mathematics, and statistics to the analysis of highly complex datasets acquired by experimental neuroscientists. They have also proposed conceptual frameworks for understanding some of the brain's most fundamental processes and how these may go awry in some of the most devastating disorders of the nervous system.

Their scientific achievements have also paved the way for the development of brain-inspired artificial intelligence, one of the emerging and transformational technologies of our time."

About The Brain Prize

Scope

The Brain Prize is awarded each year by the Lundbeck Foundation. The Lundbeck Foundation is one of Denmark's largest foundations encompassing a comprehensive range of commercial and philanthropic activities – all united by its strong purpose; Bringing Discoveries to Lives. The Foundation's philanthropic grants amount to more than DKK 500m annually and primarily focus on the brain – including the world's largest personal prize for neuroscience, The Brain Prize. The Brain Prize recognises highly original and influential advances in any area of brain research, from basic neuroscience to applied clinical research. Recipients of The Brain Prize may be of any nationality and work in any country in the world. Since it was first awarded in 2011 The Brain Prize has been awarded to 47 scientists from 10 different countries.

Selection and award

Only candidates who are nominated by others will be considered for The Brain Prize. Each year, the Lundbeck Foundation receives many outstanding nominations from all over the world. Recipients of The Brain Prize are chosen from the pool of nominees by The Brain Prize selection committee which consists of 10 leading neuroscientists from all over the world, and from diverse disciplines within neuroscience. Brain Prize recipients are presented with their medals by His Royal Highness, King Frederik of Denmark, at a ceremony in the Danish capital, Copenhagen.

Purpose

The Brain Prize is first and foremost a celebration of outstanding science and outstanding scientists, but it is also an opportunity to raise awareness of the winners, their science, and their field. Following the award of The Brain Prize, recipients engage in a series of seminars, lectures, and conferences, organised by the Lundbeck Foundation. These activities celebrate the achievements of The Brain Prize winners and help raise awareness of their work and their field amongst the global neuroscience community. The Brain Prize is also used as a platform to engage with and educate the public about the importance of brain research, its challenges, and breakthroughs. The Brain Prize also serves to highlight the Lundbeck Foundation's vision of making Denmark a leading neuroscience nation.

More information about The Brain Prize, Brain Prize Laureates and the nomination and selection process can be found here. Here you will also be able to access educational material and documentary films about Brain Prize winners and their science.

Larry Abbott

William Bloor Professor of Theoretical Neuroscience at Columbia's Zuckerman Mind, Brain, Behavior Institute

Larry Abbott is a physicist-turned-neuroscientist who uses mathematical modeling to study neural circuits responsible for sensation, action and behavior. Abbott's PhD from Brandeis University and his postdoc at the Stanford Linear Accelerator Center were in theoretical particle physics.

He was a professor of physics at Brandeis when, in 1989, he transitioned to neuroscience research, joining the Biology Department in 1993.

In 2005, he moved to Columbia University, and he is currently the William Bloor Professor of Theoretical Neuroscience at Columbia's Zuckerman Mind, Brain, Behavior Institute.

In addition to theoretical work on neural network and synaptic dynamics, Abbott has collaborated with numerous experimental colleagues on a variety of topics and systems, including vision, olfaction, electrosensing, motor control, memory and navigation.

In collaboration with Eve Marder, he developed the dynamic clamp, a tool of experimental electrophysiology, and he is the co-author with Peter Dayan of a widely used textbook on theoretical neuroscience.

His current work includes connectome-based circuit modeling of neural circuits in Drosophila.

Terrence Sejnowski

Francis Crick Professor, The Salk Institute for Biological Studies & Distinguished Professor of Neurobiology, UC San Diego

Terrence Sejnowski received a Ph.D. in Physics from Princeton University. He was a postdoctoral fellow at Princeton University and Harvard Medical School before being appointed to a faculty position in the Department of Biophysics at Johns Hopkins University in 1981.

He moved to La Jolla in 1989 and is currently the Francis Crick Professor at The Salk Institute for Biological Studies and a Distinguished Professor of Neurobiology at UC San Diego. Terrence Sejnowski was also an Investigator with the Howard Hughes Medical Institute from 1991 to 2017.

He is a member of the National Academy of Sciences, the National Academy of Engineering, the National Academy of Medicine, the National Academy of Inventors, and the American Academy of Arts and Sciences.

Dr. Sejnowski's research in neural networks and computational neuroscience has been pioneering. His research aims to understand the computational resources of brains and build linking principles from brains to behavior using computational models. He pursued this goal with both theoretical and experimental approaches at multiple levels of investigation ranging from biophysical to systems levels.

The central issues he has explored are how synaptic strength is regulated, how dendrites integrate synaptic signals in neurons, how networks of neurons generate dynamical patterns of activity, how sensory information is represented in the cerebral cortex, how memory representations are formed and consolidated during sleep, and how distributed sensorimotor systems are coordinated.

His laboratory developed Independent Component Analysis (ICA) for blind source separation, which is universally used for analyzing EEG from the scalp and brain imaging by functional magnetic imaging (fMRI).

Sejnowski was also a pioneer in developing learning algorithms for neural networks in the 1980s, inventing the Boltzmann machine with Geoffrey Hinton; this was the first learning algorithm for multilayer neural networks and laid the foundation for deep learning.

He is the President of the Neural Information Processing Systems (NeurIPS) Foundation, which organizes the largest AI conference, and he is a leader in the recent convergence between neuroscience and AI.

Haim Sompolinsky

Professor, Harvard University and Hebrew University

Haim Sompolinsky earned his PhD in Physics from Bar-Ilan University, Israel. Currently, he holds positions as Professor of Physics and Neuroscience (Emeritus) at Hebrew University, Israel, and as Professor of Molecular and Cellular Biology and of Physics (in Residence) at Harvard University, USA.

The laboratory led by Haim Sompolinsky employs statistical physics methods to investigate the emergent dynamics and collective behavior of complex neuronal circuits and their relationship to critical brain functions, including learning, memory, perception, and cognition.

His theoretical predictions have received experimental support from the study of navigational circuits in fly and rodents. His work has elucidated how the dynamic balance between neuronal excitation and inhibition leads to chaotic yet stable patterns of brain activity. This has influenced our understanding of the origins of variability in neuronal activity, the mechanisms underpinning the stability of neuronal dynamics, and the impact of the disruption of excitation-inhibition balance in neurological diseases.

More recently, Sompolinsky has developed geometric methods that provide a principled approach to the study of information processing in vision and language, in both artificial neural networks and brain circuits. This work has revealed surprising similarities between the two systems and opens a new path for synergetic investigations of intelligence in natural and artificial systems.

Brain Prize 2024 Commentary

Changing how we think we think – Transforming neuroscience with theory

Tim P. Vogels

Professor of Neurotheory and Computational Neuroscience
Institute of Science and Technology, Austria

Γνῶθι σαυτόν, know thyself! –Apollo's instruction to glance inward, at our own mechanistic origins, is perhaps what most distinguishes us from both our biological ancestry, as well as our most recent creation, artificial intelligence. In between the biology and the artifice sit squarely Haim Sompolinsky, Terry Sejnowski, and Larry Abbott, whose contributions have been recognized with The Brain Prize 2024.

As old as humankind–maybe it's defining quality–is the quest to understand what makes us "tick". Unraveling the complexity of the brain with its intricate neural processes has accelerated over the past century. Generously sponsored, neuroscience has seen breakthroughs in ever more rapid succession, in ever more quantitative detail and relief. This acceleration was marked by the entry and coalescence of biology-minded physicists into an emerging subfield of Theoretical and Computational Neuroscience.

Theory in neuroscience can be traced back to the work of Galvani, Lovelace, Lapique and Sherrington, who speculated about the calculus of the nervous system (Ferry, 2015), the nature of integrability (LaPique, 1907) and transmission (Sherrington, 1897) of its unit. Even Cajal, despite his low opinion of theorists (Cajal, 1897), theorized frequently and famously–often wrongly, for example about theorists (Cajal, 1897).

However it was the foundational work of McCulloch and Pitts in the mid-20th century that presented the first analytical model to express neurons as binary logical operators, setting the stage for neural network development and computational mind theories (McCulloch & Pitts, 1943). Hodgkin and Huxley followed this quantitative spirit, detailing the biophysical underpinnings of nerve impulse transmission in their now iconic description of ion channel interaction (Hodgkin & Huxley, 1952). Collectively, these early studies established the groundwork for all contemporary models of neuronal and neural networks, i.e. those geared towards biological plausibility, as well as those geared towards efficient computing and artificial intelligence (AI), respectively. Towards the former, biological faithfulness, Donald Sholl and Wilfrid Rall (soon after with John Rinzel) mathematically elevated our understanding of dendritic structure and signal integration (Sholl, 1953, Rall, 1959). Quantitative tools of measuring brain activity–see, e.g. Brazier, 1962 for an early version of computer-aided EEG (Brazier, 1962)–were developed on these foundations, and Wilson and Cowan soon honed in on the dynamics of populations of neurons (Wilson & Cowan, 1972), within the compute constraints of their days. Soon after, John Hopfield introduced a neural network model capable of storing and retrieving information (Hopfield, 1982), and carrying an explanatory–and inspirational–power for many experimentalists, creating a formidable articulation of the memory engram (Semon, 1921). More importantly, it was a model around which physicists, psychologists and biologists alike gathered, conversing and exchanging ideas in a common language.

Three of the physicists who came to Hopfield's roundtable were this year's Brain Prize winners. They were drawn in by the air of excitement and discovery; by the promise that, having understood mechanistically individual cells thanks to Hodgkin and Huxley, the future was wide open. The early work on neural networks was just the beginning. They were also drawn in by their experimental colleagues who saw the synergy that quantitative and qualitative approaches could bring. Together with Eve Marder, Moshe Abeles, Ad Aertsen, Peter Dayan and Li Xiaoping, who were already then particularly unafraid of combining numbers and wetware, they recognised the opportunity for genuine symbiosis to solve the riddles of the brain.

And solve riddles, they did. The three prizewinners soon became trailblazers of theoretical and computational neuroscience, exploding onto the stage with important contributions almost immediately; Between the three of them, they formalized the performance of Hopfield's network (Amit et al., 1985), as well as developed novel experimental methods, such as the dynamic clamp technique (Sharp et al., 1993), and they made extensive contributions on neural network models like the Boltzmann machine (Ackley et al., 1985) that have proven essential to the fields of neuroscience and AI. The last five decades since the publication of these papers have been dizzyingly productive for two reasons: The ever accelerating development of new technologies and methods, and the principled, quantitative interpretation, extrapolation, and reproduction of experimental results in theoretical models and frameworks.

Arguably, the former has received a fair share of attention in the past, also from the Brain Prize, for the multiple technological and experimental revolutions of the field. The latter has been less recognised by broad neuroscience awards, and it remains a party favorite to ask "what has theory ever done for neuroscience?" – A lot, as it turns out, and a lot from Sejnowski, Abbott and Sompolinsky.

Here, we celebrate three of the most prominent figures in computational and theoretical neuroscience, maybe to be recognised on behalf of the field that they helped to shape so thoroughly. Together they have contributed to neuroscience such that their absence is unfathomable, in hindsight. The main directions of their research, representational learning, statistical population-level dynamics, and implementational circuit and synapse models shine like bright and steady stars–navigational markers for the rest of us.

Like all three of this year's winners, Terry Sejnowski's career began in physics. He earned a Bachelor's degree from Case Western Reserve University in 1968 and continued at Princeton, where he completed his Master's and Ph.D. in physics in 1978, in John Hopfield's lab, who must have inspired him early to take up postdoctoral fellowships in biology at Princeton (1978-1979) and Harvard Medical School (1979-1981) under John Archibald Wheeler's and Stephen Kuffler's supervision, respectively. He joined Johns Hopkins University's Department of Biophysics as faculty in 1982 and since 1988 Sejnowski has been at the Salk Institute.

Terry Sejnowski's contributions to neuroscience began with his work on the Boltzmann machine, exploring mechanisms of local synaptic learning and bridging the gap between computational theories and biological realities (Ackley et al., 1985; Sejnowski & Rosenberg, 1987). The machine, a type of stochastic recurrent neural network, could learn deep data representations in an unsupervised manner, through local learning rules that activated neurons probabilistically without being primed with a solution. The Boltzmann machine is an example of unsupervised learning–an algorithm that can compress the statistics of data. Sejnowski's search for related algorithms had major practical implications. For example, his invention of Independent Component Analysis (ICA) (Bell & Sejnowski, 1995) gave us our most practical way to separate mixed signals into their constituent parts. Sejnowski pioneered its application in brain imaging, but others have applied it to deconstruct data across a myriad of scientific disciplines. It is most famous for being the first practical solution to the "cocktail party problem", where listeners can hear multiple conversations in parallel and must discern who is saying what.

More broadly, Sejnowski should take a large portion of the credit for the power of modern representation learning. In the 80s and 90s, when AI research was focused on symbolic solutions, Sejnowski's lab was focused on how to make machines that learn from data. His lab members at the time went on to be pioneers in deep learning (Hinton), reinforcement learning (Montague, Dayan), and predictive coding (Rao) amongst others. Many of the key ideas that have led to the powerful representation learning systems of today have had their seeds in these creative times. Sejnowski has remained at the forefront of computational neuroscience, also through his mentorship of students and postdocs, and custodianship of massive open online courses and conferences such as NeurIPS (for which he has served as president since 1993), and journals like Neural Computation (where he is the editor-in-chief). His contributions will continue to bear fruit through the people and institutions he helped to build.

Larry Abbott, too, began his career in physics, at Oberlin College and later at Brandeis University, where he received his PhD under the supervision of Howard Schnitzer in 1977. After postdocs in theoretical particle physics and cosmology at the Stanford Linear Accelerator Center (1977-1979), and CERN (1980-1981), he returned to Brandeis for his first position. Abbott's career pivoted towards neuroscience in 1989 when he was enchanted by the sirens' (read: Lobsters') spiking song in Eve Marder's lab. Since, Abbott became a co-founder of Columbia University's Center for Theoretical Neuroscience in New York.

Abbott's entry into neuroscience was marked by the development of the so-called dynamic clamp with Eve Marder: a computational tour-de-force to simulate channel action in real time, allowing closed loop patch clamp studies on channel function. Soon after, Abbott delved into synaptic plasticity, and later network dynamics, developing theoretical and numerical models that confirmed and predicted experimental results and neatly recapitulated what we knew about plasticity and depression, cortical gain modulation and more generally how the brain adapts to a changing environment (Song et al., 2000, Abbott et al., 1997, Chance et al., 2002). His work was often groundbreaking, owing to his ability to intensely focus on and isolate a specific process of interest, showing a rare talent to make complicated biological interactions accessible by reducing them to their relevant components, following Einstein's directive to make things "as simple as possible, but not simpler". The overarching spirit of Abbott's work is a deep respect for the data, and a constant striving to reproduce and explain the mechanisms by which the data originated. He, like none other, connected the biology to algorithmic and implementational understanding. I, and my peers in the lab have benefitted from this talent, when it came to receiving advice and guidance on how to model the interaction of external and internal cortical activity in recurrent and feed forward networks that aimed to explain the production of sensory and motor activity (Vogels & Abbott 2005, Rajan & Abbott 2006, Sussillo & Abbott, 2009). In the more recent past, Abbott switched systems, to focus on the nervous system of drosophila melanogaster, helping to tie together–in models that deeply respect the biology–disparate experimental results and connect them to earlier theoretical work such as receptive-field models of his colleague Haim Sompolinsky (Ben-Yishai et al., 1995). Through his work, Abbott continues to advance our understanding of neural mechanisms and dynamics at the micro and macro level.

Haim Sompolinsky also aimed to become a physicist, and remains a Professor of Physics today. He obtained his Ph.D. from Bar-Ilan University in Israel (1980). Following that, he pursued postdoctoral research at Harvard University's physics department (1980-1982) under the mentorship of Prof. Halperin. Sompolinsky originally focused on theoretical physics as an associate professor at Bar-Ilan University before his transition to a professorship in physics at the Hebrew University of Jerusalem. His research interests, now as ever, include phase transitions, critical phenomena, nonlinear dynamics, and the statistical mechanics of spin glasses, but the focus of where he sought out these phenomena drifted.

A distinct phase transition to computational neuroscience occurred in the mid-1980s when he was recruited by his colleagues Daniel Amit and Hanoch Gutfreund, to lend a helping brain with the analytical considerations of applying spin glass theory to the newly published work of John Hopfield. I was told it was then, in Amit's office that he developed his extraordinary strength to recognise the essence of the théorie-du-jour, and express its importance in uncompromising rigor. It was this strength that allowed him to distill how large populations of neurons interact to produce coherent activity patterns (Ginzburg & Sompolinsky 1994), how the seemingly chaotic nature of cortical neuron spiking could stem from an intricate balance between excitatory and inhibitory inputs (van Vreeswijk & Sompolinsky, 1996), and how contrast invariant orientation selectivity–and as a byproduct, self-sustained working memory–can originate from ubiquitously found center-surround “Mexican hat” wiring (Ben-Yishai et al., 1995). These insights are some of the deepest in neuroscience in the last 50 years, giving us a whole new regime of cortical function, explaining how synapses can be strong, without overwhelming the system constantly. They are now the dominant theories of how the cortex operates. If the above milestones seem self-explanatory, it is thanks to Sompolinsky and colleagues (one of whom, Carl van Vreesvijk, recently passed prematurely, but left us with a beautiful legacy of studies), pushing theoretical concepts into the community’s focus of interest, transforming them into household names. The longest theme of Sompolinsky’s research, the beginning of which coincided with his co-founding of one of the world’s earliest Interdisciplinary Centers for Neural Computation at the Hebrew University of Jerusalem remains his perceptron (Barkai et al, 1992) and tempotron (Gütig & Sompolinsky, 2006) work that continues to surprise with new insights on the balance of stability and sensitivity to input correlations dominates and facilitates synaptic learning.

Sompolinsky, Sejnowski and Abbott entered the field of neuroscience mid-career, as principal researchers with their own groups. They were taken by the beauty of what has been called the most complex system of the universe. They could not help but contribute to describing its wonders in the language of maths, rigorously but also accessibly. If one has trouble pinpointing what was the main contribution of any of this year’s winners, because one cannot identify the one seminal line of work, it is–in my opinion–testament to the breadth of their work. Representational learning, statistical population level dynamics, and implementational circuit and synapse models by Sejnowski, Sompolinsky and Abbott, respectively, represent the three possibly most important axes of how we approach understanding the brain today.

Acknowledgements

Thanks to Panos Bozelos and Tim Behrens for help with writing this article.

References

Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science*, 275(5297), 221-224.

Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive science*, 9(1), 147-169.

Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1985). Spin-glass models of neural networks. *Physical Review A*, 32(2), 1007.

Barkai, E., Hansel, D., & Sompolinsky, H. (1992). Broken symmetries in multilayered perceptrons. *Physical Review A*, 45(6), 4146.

Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6), 1129-1159.

Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9), 3844-3848.

Brazier, M. A. (1962). The analysis of brain waves. *Scientific American*, 206(6), 142-155.

Cajal, S. R. Y. (1897). *Advice for a young investigator.*

Chance, F. S., Abbott, L. F., & Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron*, 35(4), 773-782.

Ferry, G. (2015). Ada Lovelace: in search of “a calculus of the nervous system”. *The Lancet*, 386(10005), 1731.

Gat, I., Tishby, N., & Abeles, M. (1997). Hidden Markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. *Network: Computation in neural systems*, 8(3), 297-322.

Ginzburg, I., & Sompolinsky, H. (1994). Theory of correlations in stochastic neural networks. *Physical review E*, 50(4), 3171.

Gütig, R., & Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing–based decisions. Nature neuroscience, 9(3), 420-428.

Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4), 500.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8), 2554-2558.

Lapique, L. (1907). Researches quantatives sur l’excitation electrique des nerfs traitee comme une polarization. Journal of Physiology, *Pathology and Genetics*, 9, 620-635.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5, 115-133.

Rajan, K., & Abbott, L. F. (2006). Eigenvalue spectra of random matrices for neural networks. *Physical review letters*, 97(18), 188104.

Rall, W. (1959). Branching dendritic trees and motoneuron membrane resistivity. *Experimental neurology*, 1(5), 491-527.

Sejnowski, T. J., & Rosenberg, C. R. (1987). Parallel networks that learn to pronounce English text. *Complex systems*, 1(1), 145-168.

Semon, R. W. (1921). The mneme. G. Allen & Unwin Limited.

Sharp, A. A., O’Neil, M. B., Abbott, L. F., & Marder, E. (1993). Dynamic clamp: computer-generated conductances in real neurons. *Journal of neurophysiology*, 69(3), 992-995.

Sherrington, C.S. (1897) in *Textbook of Physiology* (Foster, M., ed.), p. 60

Sholl, D. (1953). Dendritic organization in the neurons of the visual and motor cortices of the cat. *Journal of anatomy*, 87(Pt 4), 387.

Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature neuroscience*, 3(9), 919-926.

Sussillo, D., & Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4), 544-557.

Van Vreeswijk, C., & Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293), 1724-1726.

Vogels, T. P., & Abbott, L. F. (2005). Signal propagation and logic gating in networks of integrate-and-fire neurons. *Journal of neuroscience*, 25(46), 10786-10795.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1), 1-24.

Autobiographies of the 2024 Brain Prize winners

Larry Abbott

William Bloor Professor of Theoretical Neuroscience at Columbia's Zuckerman Mind, Brain, Behavior Institute

When I was around 5 years old, my father bought a new amplifier for his record player. The first thing he did was not to hook it up and listen to it but, instead, to unscrewed the cover so he and I could peer in and see the electronics. Ever since, I have been driven to ask and wonder about how things work.

I grew up in Toronto where my father was born and where my mother and her family ended up after fleeing Nazi Germany. I am a middle child with an older and a younger sister. When I was 7, my father gathered us together to announce that he was getting a new job. I remember hoping that he was going to join the Royal Canadian Mounted Police (although he was a scientist-engineer), but he said we were moving to Boston. We actually moved into a Boston suburb, Needham, where I graduated from high school in 1968.

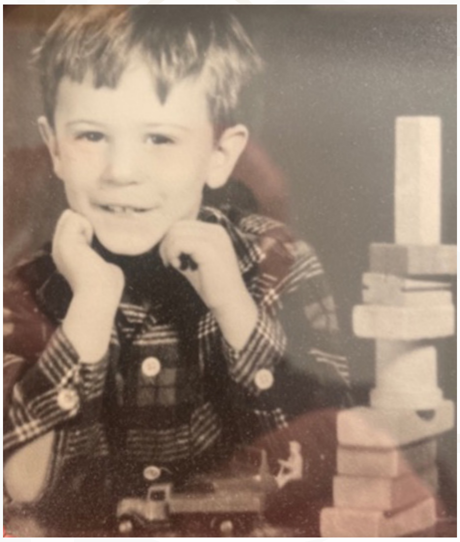
A younger me.

I attended Oberlin College. I was a distinctly mediocre student; my favorite school-related activity was playing on the hockey team. The Vietnam War caused great unrest on campus. My one attempt to take a biology course was thwarted when the college was shut down by the tragic Kent State University shootings, and we all went off to Washington D.C. to protest. During the summer of my sophomore year, I worked at Honeywell and was given the task of studying a new design for an MOS-based computer memory chip being developed. I realized that a simple trick - allowing the memory bits to flip between 0 and 1 values while keeping track of the flipping - would double the chip's speed. I received a patent for this, and the idea was used in a memory chip built by Intel.

I decide to be a physics major in college when I was shown how the speed of light can be derived from Maxwell's equations. Despite a strong commitment to physics, at some point I grew disenchanted with the college experience and dropped out. As a result, I do not have an undergraduate degree. Rather than studying physics the sensible way, at college, I studied it on my own while painting houses until I realized that there must be a better way to proceed, and I applied to graduate school. I was lucky that Brandeis University agreed to admit me although, due to my clear lack of qualifications, without a stipend (they generously fixed this after one semester). At Brandeis, I finally became a good student and was inspired by what I was learning. My life was further enriched when I married my wife, Cathy, in my first year as a graduate student, and we had our first child, Paul, in my last.

With Howard Schnitzer in 2003.

My thesis advisor was Howard Schnitzer (I published my first paper with Howard in 1976 and, 45 years later, I published a paper with his son Mark, a neuroscientist at Stanford). My thesis work was in theoretical particle physics. In particular, Howard and I developed a novel calculation method for quantum field theory. I then did postdoctoral work at the Stanford Linear Accelerator Center before returning to Brandeis as a faculty member in the physics department. Along the way, I also spend a year working at CERN, the European accelerator center. During my last year as a postdoc, our daughter Karen was born.

With Paul, Cathy and Karen.

My graduate and postdoctoral years coincided with a remarkable time in particle physics, when it became apparent that the so-called "standard model" provided an accurate description of all the known elementary particles and their interactions. During my postdoc, I published a study that provided some of the evidence for this. My faculty-level research in particle physics included work on the cosmological constant, an axion theory of dark matter, development of the background field method, calculations of the microwave background anisotropy, and work in general relativity and gauge field theory.

I progressed along the path from assistant to full professor of physics over the course of around 10 years. At that point, it felt to me that theoretical particle physics had become the victim of its own success; it was clear that the theory worked and that new data to ponder would not be available for at least 20 years. I did not want to wait that long, so I started looking for other research areas. I was drawn to the beautiful work on spin glasses done by Georgio Parisi and others, and this led me to a remarkable paper by Dani Amit, Hanoch Gutfreund and Haim Sompolinsky on the Hopfield model. I also saw an exciting talk by Terry Sejnowski on the re-emerging field of artificial neural networks. These influences led me to become a neural network physicist - it was clear that someone who knew what an axion was but not an axon was not a neuroscientist.

And then I walked into Eve Marder's laboratory at Brandeis. I went there solely out of curiosity, and then postdoc Michael Nusbaum kindly took me into the rig room and showed me the oscillating neuronal circuit (from a lobster) that he was working on. My fate was sealed by the mesmerizing rhythmic sound of action potentials (not that I knew what they were) on the audio monitor. I walked out of the lab in a daze and, frankly, terrified because I knew that I was going to switch my research to neuroscience, a field that was as ignorant of me as I was of it, and that this was going to be a disaster. The reason it wasn't was Eve.

With Eve in her office.

I went back to see Eve the next day and confessed my infatuation with neuroscience, but perhaps not the full extent of my ignorance. Rather than throwing me out, she talked to me about neuroscience day after day until we each finally started understanding what the other was saying. Once this happened, Eve and I could combine our skills, and one result of this was the dynamic clamp, now a widely used method to mimic conductances in recorded neurons. Eve and I went on to do a number of studies including developing models of the homeostatic regulation of intrinsic neuronal conductances with Gwendal LeMasson and testing them experimentally with Gina Turrigiano.

I moved from Brandeis to Columbia University in 2005 to set up a theory center with Ken Miller. The center has grown over the years and now has a faculty of 9 and over 50 postdocs and graduate students. All of my colleagues at Columbia have made me a better scientist, and I have especially benefitted from the wisdom, inspiration and friendship of Richard Axel.

My work is intensely and extensively collaborative. I believe that theorists should be adventurous because, unlike experimentalists who are often tied to a single species and subset of brain areas, theorists can roam, thereby carrying ideas across sub-fields. In this spirit, I have worked on a variety of topics with a large number of wonderful collaborators – unfortunately too many for a complete list. The point of applying mathematics and computer simulation, the tools of theoretical neuroscience, to neural systems is to reveal aspects of how things work that would be difficult to see without an overarching, even if abstract, framework. I will try to illustrate the two principles outlined in this paragraph with examples from my work.

Synaptic plasticity has been a rich area for theoretical neuroscientists. By combining a model for the dependence of synaptic plasticity on spike-timing with a model of place-cell activity, Kenneth Blum and I showed how a predictive representation could develop in the hippocampus, which led to experimental work revealing the predicted predictive shifts. Later, Sen Song, Ken Miller and I showed how spike-timing dependent plasticity (a name we coined) could induce competition between different synapses that was beneficial for learning. Using a model allowed us to infer what would happen across many synapses from data characterizing plasticity at a single synapse. Timing-dependent plasticity is an important element in a circuit of mormyrid electric fish studied by my colleague Nate Sawtell. In work with Nate and others, modeling has allowed us to pinpoint the physiological and anatomical features that allow this circuit to cancel the fish's self-generated electric fields so it can detect the much weaker fields produced by prey, for example. Physiology from Nate's lab and an ongoing connectomics study have allowed

Theory Center members acting out the fly ring attractor heading direction system.

us to test and refine our model over the years. Spike timing is also relevant for the dopamine-dependent plasticity that supports memory in the mushroom body of Drosophila, and models of this system, developed and tested in collaboration with Richard Axel and members of his laboratory, allowed us to compute the fly's capacity for olfactory recognition memory. Finally, dopamine-mediated plasticity is also important for linking the fly's "compass" system to the visual world. In a collaboration with Vivek Jayaraman and a group at Janelia, we used a plasticity model to predict how optogenetic stimulation could remap this system, a result that illustrated the remarkable plasticity in the link between what a fly sees and its directional map of the world.

In 1996, Emilio Salinas and I developed a model for how visual information is transformed from retinal to body-centered coordinates in the primate parietal cortex, based on work and ideas from Richard Andersen (Terry Sejnowski and Alex Pouget also worked on this system). Recently, a collaboration with Gaby Maimon and members of his laboratory has shown that a similar mechanism, although much more compactly and elegantly realized, transforms sensory information from egocentric to allocentric (world-based) coordinates and, conversely, goal information from allo- to egocentric coordinates in Drosophila. Among other things, modeling work helped reveal the functional beauty of the fly's navigational circuitry. Connectomics is providing exceptionally informative for model building, and I have been fortunate to be part of teams exploring these data led by Marta Zlatic and Albert Cardona, and Jerry Rubin. I have also collaborated with Rudy Behnia and her lab studying the implications of connectomics for the fly visual system.

Lab dinner, 2023.

Working in academia means that the joy and privilege of a life in research is combined with the joy and privilege of educating and being educated by trainees at all levels. Watching all of the students and postdocs I have worked with grow and thrive has been profoundly satisfying, and their scientific contributions have been essential. Of course, an ultimate satisfaction is watching your children grow and thrive (with the added advantage that they don't ask for letters of recommendation). Paul is an attorney working in telecommunications and Karen is a professor of theoretical ecology at Case-Western University.

I am very pleased to be receiving the Brain Prize with my friends and colleagues Haim Sompolinsky and Terry Sejnowski. While feeling greatly honored, I think that all three of us acknowledge that this is really a prize for the entire field of theoretical neuroscience. I want to say to everyone in the field that this prize honors the collective efforts of us all.

Terrence Sejnowski

Francis Crick Professor, The Salk Institute for Biological Studies & Distinguished Professor of Neurobiology, UC San Diego

Terrence Sejnowski was born in Cleveland, Ohio. After completing his B.S. in Physics at Case Western Reserve University, he received a Ph.D. in Physics from Princeton University. He was a postdoctoral fellow at Princeton University and Harvard Medical School before being appointed to a faculty position in the Department of Biophysics at Johns Hopkins University in 1981. He moved to La Jolla in 1989 and is currently the Francis Crick Professor at The Salk Institute for Biological Studies and a Distinguished Professor of Neurobiology at UC San Diego. Terrence Sejnowski was an Investigator with the Howard Hughes Medical Institute from 1991 to 2017.

I spoke Polish until I went to grade school. My parents said I gave Polish lessons to the kids I played with on the street. I was also fast at finding pieces in picture puzzles, to the amazement of my parents and their friends. These two early talents later became passions: Teaching and solving scientific puzzles. My father was an engineer who designed jet engine fan blades. My mother was a born organizer who ran a pool of typists during WWII and organized her five sons and one daughter. I was the oldest. I attribute my self-confidence to my parents' trust in my early exploration of the world and the freedom they gave me to pursue my interests while growing up. I did not then fully appreciate this, but I now know how much I owe them.

I was the science guy at school. I vividly remember building a working volcano made from Paper Mache and fueled by a mixture of black powder and aluminum powder I had concocted in my basement chemistry lab. When I tested it outdoors, it spouted an impressive sparkling flame and black smoke, but when I demonstrated it to my grade school class, I failed to realize that it would fill the classroom with smoke, set off a fire alarm, and evacuate the school. This event made my reputation.

I was President of the Radio Club in high school. I stayed after school every day to talk with ham radio enthusiasts around the world and build electronic equipment. Mike Stimac, who was the faculty advisor, had tracked the first Sputnik satellite in 1957 with "the boys," a story that was picked up by the national press. We obtained a commercial radio transmitter and assembled a huge Yagi antenna on the school's roof for Project Moonbounce. Mike once asked me: "What is your mission?" This question has been in the back of my mind ever since then.

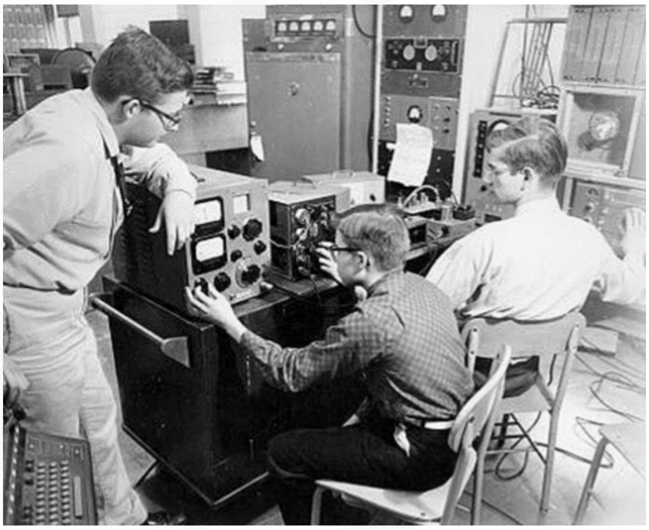

Terry at the St. Joseph High School Radio Club in 1964.
Top center: Tuning a receiver with the kilowatt transmitter behind him.
Bottom left: Preparing the antenna for Project Moonbounce.

I matriculated at the Case Institute of Technology and received a Physics B.S. from Case Western Reserve University. Physics was alluring, especially relativistic astrophysics, and I applied to graduate schools. My most memorable adventure in college was when Mike Stimac called and asked if I could help him co-pilot a Cessna 210 to the West Coast. When we flew into Los Angeles, I had a few days free to visit Caltech in Pasadena, a mecca for theoretical physics. Without having an appointment, I met with Carl Anderson, Chair of the Physics Department, who received a Nobel Prize for discovering the positron. He asked me whether I wanted to be an experimentalist or a theoretician. It had never occurred to me that I had to choose. He said there was someone at Caltech who did both, so there was an existence proof.

I went to Princeton for graduate school, where I fell under the spell of John Wheeler, a legendary physicist who coined the term black hole. I worked on sources for gravitational waves, like supernovae or colliding black holes, concluding that it would be decades before detectors could reach the super sensitivity needed for detection. Another Wheeler student, Kip Thorne, thought it would take seven years and come from a supernova in the Virgo cluster of galaxies, which I quoted in a paper I published in 1974. Gravitational waves were finally discovered in 2015, 41 years later. I sent Kip a congratulatory email reminding him of his predictions. He replied: “I was wrong about the Virgo cluster, but what’s the difference between 7 years and 40 years on a cosmic time scale?”

My neuroscience immersion

I have always been intrigued by brains, which have an inner complexity as mysterious as the cosmos. While working on the cosmos, I was taking courses from Charlie Gross, who worked on the visual systems of monkeys, Mark Konishi, a neuroethologist who worked on barn owls and songbirds, and worked in Alan Gelperin’s lab, who studied learning in the garden slug, Limax maximus. I was fortunate that John Hopfield, a biophysicist who previously had a distinguished career in condensed matter, was becoming interested in neuroscience. John’s mentorship gave me confidence that I could make a career by modeling neural networks in brains, which was not yet a traditional career path. My 1978 Ph.D. thesis on “A Stochastic Model of Nonlinearly Interacting Neurons” led to my first publications in a nascent field that would eventually become computational neuroscience.

After getting my Physics Ph.D., I took the Neurobiology Summer Course at the Marine Biological Laboratory at Woods Hole. This was an in-depth eleven-week immersion in ways to study brains: electrophysiology, neuroanatomy, biochemistry, neuropharmacology, and tissue culture. I loved every minute, recorded muscle action potentials, ran gels, and looked into neurons with electron microscopy. This led to my first publication in neuroscience on the freeze fracture of the ampullae of Lorenzini, ultra-sensitive electroreceptors in skates.

The course was taught by faculty in the Department of Neurobiology at Harvard Medical School. I stayed at Woods Hole for a few more weeks to finish the skate project. It was quite a surprise when Steve Kuffler, the father of modern neurobiology, called to ask if I was interested in working with him as a postdoc. The eleven-week immersion led to a two-year immersion at Harvard. Steve and I experimented daily, voltage clamping a peptidergic synapse and solving a mystery in the literature on the underlying conductance changes. I also kept up my interest in neural networks, attending a workshop at UC San Diego organized by Geoffrey Hinton. Geoff was part of the Parallel Distributed Processing (PDP) group, including Jay McClelland and Dave Rumelhart. This small group would produce a seminal two-volume book highly influential in the burgeoning field of neural networks. I had a chapter in the book entitled “Open Questions about Computation in the Cerebral Cortex.”

Steve died suddenly while we were writing up our experimental results, leaving me without a mentor. Fortunately, with help from Torsten Wiesel, who had become the department chair, I landed a job in the Thomas C. Jenkins Biophysics Department at Johns Hopkins University, where I set up a wet lab and taught a course on computational biophysics. This was a perfect launching pad for my career. Faculty in the department were highly supportive, and I was fortunate to attract exceptional students.

The 1986 Connectionist Models Summer School at the Carnegie Mellon University. Terry is in the first row next to Geoffrey Hinton and Jay McClelland.

On my way back from the West Coast in 1985, I was snowed in at the Denver Stapleton airport. Steve Zucker was also stranded, and we commiserated over hot chocolate. We discussed ways to model the visual cortex and the need for a small workshop to discuss modeling with experimentalists. The Woods Hole Workshop on Computational Neuroscience was born from this chance meeting and met annually with 10-15 junior researchers. Discussion usually began with the first slide and continued throughout. This workshop was a crucible for new ways to think about brain function. In 1988, a Summer School on Methods in Computational Neuroscience was inaugurated, which continues today. In 1997, the Woods Hole Workshop moved to Telluride, Colorado.

In 1994, Hirsch Cohen at the Sloan Foundation began supporting Centers for Theoretical Neuroscience. Jerry Swartz, who founded Symbol Technologies, a bar-code scanner company, continued this support in 2004 and expanded it to eleven Centers. At a time when grants from NIH for computational neuroscience projects were rare, the Centers helped train a generation of postdoctoral fellows and held annual meetings. The Sloan-Swartz Centers anchored the computational neuroscience community at the Salk Institute and UC San Diego. Jerry Swartz also funded a prize for computational and theoretical neuroscience under the aegis of the Society for Neuroscience, which increased its visibility. Jerry deserves recognition for his generous support over two decades.

Learning how to compute with neural networks

Geoff Hinton and I had started collaborating on what would become the Boltzmann machine. He had a faculty position at Carnegie Mellon University in Pittsburgh, and we would drive back and forth on weekends. I remember a phone call from Geoff, who announced that he had discovered how the brain works. A long-standing logjam was finding a learning algorithm for multilayer neural networks, and Geoff had an insight that led to a learning algorithm for training the weights in Boltzmann machines. This broke the logjam and soon led to backpropagation, which was much more efficient.

With learning algorithms, we could create small network models that solve computational problems, giving us insights into how brains might solve them. The puny computers at the time limited us to one layer of "hidden" units between the input and output layers. I used neural networks to ask questions about vision, which was the best-understood part of the cortex. In the 20th century, recordings from single neurons in the visual system revealed a multilayer hierarchy in which neurons responded to simple patterns, such as lines and edges in the first cortical layer, and more complex patterns, such as faces, in higher layers.

Terry at Johns Hopkins when NETtalk was invented in 1986.

In 1988, I trained a neural network model to compute the curvature of a shaded surface. I found that the hidden units had the properties of simple cells discovered by Hubel and Wiesel in the primary visual cortex of cats despite the absence of edges and lines in the images used to train the networks. Our paper in Nature introduced the term "projective field" of a neuron – its output targets – which was as important as the input receptive field in determining the function of a hidden unit in this simple neural network.

Of all my neural network projects from that era, the one that had the most impact was NETtalk, a network trained to translate text to speech. English phonology is notoriously difficult owing to many exceptions and influences from other languages. Remarkably, the learning process produced babbling sounds, then the regularities, and finally the exceptions, all of which were accommodated in a network with a few hundred units and twenty thousand weights, tiny by today's standards. This demonstrated that a small feedforward network could encompass the complexities found in phonology, a real-world problem not easily captured by rule-based systems.

In retrospect, NETtalk presaged the remarkable ability of much larger networks to perform many natural language tasks like language translation. In 1986, Geoff Hinton and I organized the first Connectionist Models Summer School at Carnegie Mellon, which brought together a group of young and enthusiastic researchers for intense immersion in the new neural network paradigm. I recall a skit where they formed three lines and simulated NETtalk, trying and failing to pronounce my last name.

These were the go-go years for neural networks. Conferences were spawned, journals were founded (including Neural Computation, which I founded at the MIT Press), and researchers in many fields of science, engineering, and mathematics made early contributions. The annual Neural Information Processing Systems (NeurIPS) Conference, founded by Ed Posner at Caltech, attracted researchers from all these areas. When Ed died in a tragic bicycle accident in 1993, I became the President of the NeurIPS Foundation. NeurIPS continued to grow and became a leading conference for machine learning and eventually the leading artificial intelligence conference, which in 2023 had 16,000 attendees.

A major career move

I moved to La Jolla in 1989, where I established the Computational Neurobiology Lab at the Salk Institute and the Institute for Neural Computation at UC San Diego, with faculty positions at both institutions. I also became an Investigator with the Howard Hughes Medical Institute, which provided 27 years of generous support, along with funding from the NSF, NIH, and Office of Naval Research, which made it possible for me to explore many new scientific, medical, and engineering directions. I felt right at home since the PDP group had laid the groundwork at UC San Diego, and many faculty were already working on brain modeling. Students joined my lab with backgrounds that ranged from biology, psychology, and bioengineering to physics and computer science. Discussions at daily tea in my lab at the Salk Institute meshed all these areas.

Francis Crick regularly joined us for tea and discussion. Francis had switched from molecular genetics to neuroscience when he moved to Salk in 1976 and was interested in consciousness, which biologists had long neglected. He focused on visual awareness. One day, when we were discussing brain modeling, Francis told us that a brain model was not an end in itself, and its purpose was to design a nonobvious killer experiment that would give the game away. His famous model of DNA came to mind. Francis knew my future wife, Beatrice Golomb, before I moved to La Jolla and gave a blessing when we were married at the Caltech Athenaeum in 1990. She is now a Professor of Medicine at UC San Diego and a medical detective.

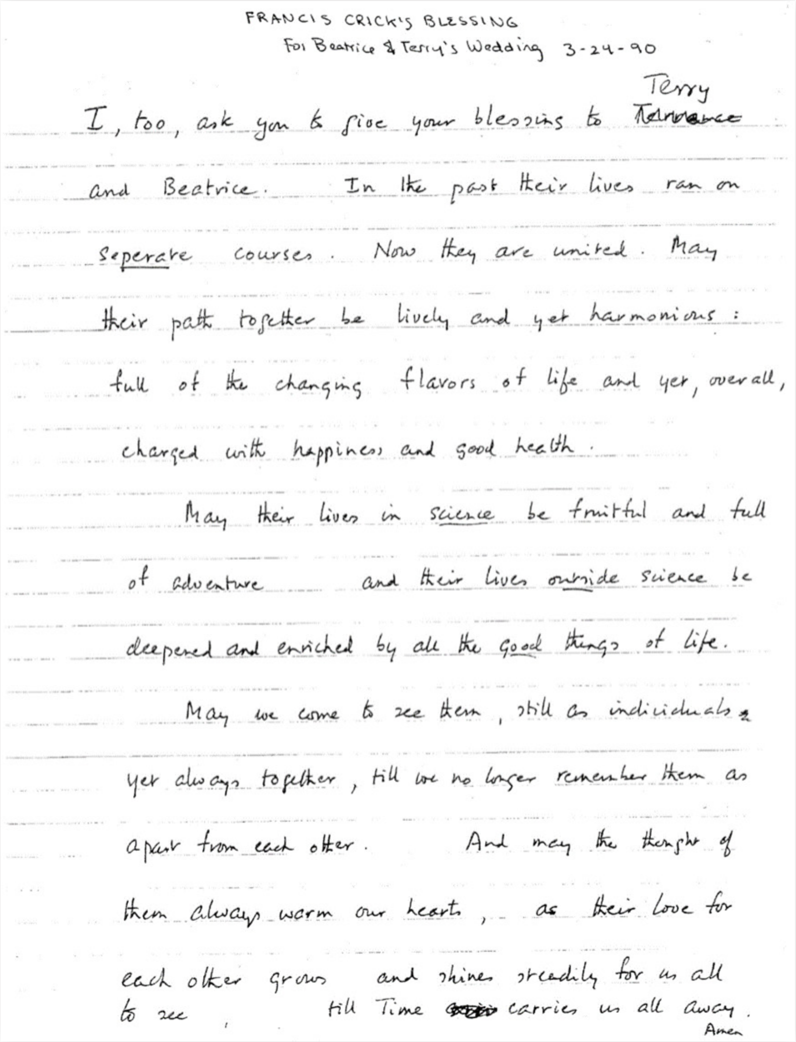

FRANCIS CRICK'S BLESSING
For Beatrice & Terry's Wedding 3-24-90

I, too, ask you to give your blessings to ~~Terrance~~ Terry
and Beatrice. In the past their lives ran on
seperate courses. Now they are united. May
their path together be lively and yet harmonious:
full of the changing flavors of life and yet, overall,
charged with happiness and good health.
May their lives in science be fruitful and full
of adventure and their lives outside science be
deepened and enriched by all the good things of life.
May we come to see them, still as individuals,
yet always together, till we no longer remember them as
apart from each other. And may the thought of
them always warm our hearts, as their love for
each other grows and shines steadily for us all
to see, till Time carries us all away.
Amen

Francis Crick delivered this blessing at the Caltech Athenaeum where Beatrice and I were married in 1990.

In 1992, Patricia Churchland and I published The *Computational Brain* with the MIT Press, which shifted the focus from the traditional emphasis on the response properties of single neurons to distributed representations in large populations of neurons. Francis Crick once wrote in Nature that the small neural networks that we could train were far from real brain models, but they were demonstrations that it was possible to represent and compute with brain-like networks. Our book was aimed at two audiences: Neuroscientists who were curious about insights from neural network models and researchers with no background in neuroscience who wanted to learn more about the brain through the perspective of brain models. It is still in print more than 40 years later, and we have been gratified to hear from many over the years who said it had a seminal influence on their interest in neuroscience.

I attracted some of the best and brightest students and postdocs to my lab in the 1990s, a golden decade. We developed a popular learning algorithm for Independent Component Analysis (ICA), which could separate the sources from mixtures of independent signals, such as someone talking from background music (Bell and Sejnowski, 1995).

When we trained a network on patches from natural images, the independent components remarkably resembled the response properties of neurons in the visual cortex discovered by Hubel and Wiesel. This was consistent with his earlier work on shaded surfaces and confirmed Horace Barlow's conjecture that the visual cortex evolved to represent natural images efficiently. ICA has many other uses and is now routinely used to analyze neural recordings and brain imaging data.

Salk Institute retreat at the La Casa del Zorro at Anza Borrego in 2002. There are four Nobel Prize winners in the first row. Where is Terry and Beatrice?

We developed a computational explanation for the responses of dopamine neurons found in the midbrain. While this neuromodulatory system was known to be part of the reward system, what dopamine represented was obscure. We predicted that the output of the dopamine neurons represents not the reward itself but reward prediction error based on a reinforcement learning algorithm called temporal difference learning. This paper was published in 1996 after two years of resistance from reviewers. Our hypothesis was subsequently confirmed by recordings from dopamine cells in monkeys by Wolfram Schultz and in humans with fMRI. This theory revolutionized our understanding of how humans respond to rewards, risks, and temporal discounting in making cognitive decisions.

I am indebted to many neuroscience collaborators, especially Mircea Steriade on modeling sleep spindles, Mary Kennedy on simulating the biochemical dynamics at synapses, Chuck Stevens on synaptic plasticity, and Kristen Harris on reconstructing dense neuropil and finding a way to estimate the precision of synaptic plasticity. The computer simulation skills of Tom Bartol in my lab were essential for the success of these collaborations. These long-term collaborations were especially important for the students and postdoctoral fellows in my lab, who greatly benefitted from the deep knowledge in these labs and access to experimental data. The research community at UC San Diego is also extraordinarily cooperative, and I have benefitted from many other fruitful collaborations.

I was on a committee that established the McDonnell-Pew Centers for Cognitive Neuroscience and served as Director of the Center at UC San Diego and Salk, and a subsequent NIH training grant also helped train a generation of cognitive neuroscientists. I was also involved in the genesis of the BRAIN Initiative, announced in 2013, whose goal was to develop innovative neurotechnology that could accelerate advances in brain research. As a member of the Advisory Committee to the Director of NIH on the BRAIN Initiative, I helped set the goals and milestones for computational, modeling, statistical, and theoretical brain research. The progress over the last ten years far exceeded what we thought could be accomplished.

Some of my heroes

I had a wonderful scientific opportunity to interact with colleagues at Caltech as a Fairchild Distinguished Scholar in 1992-93. In particular, I was adopted by Carver Mead's lab, which was the birthplace of neuromorphic engineering.

Carver is a brilliant engineer and visionary who realized that transistors on VLSI chips near threshold had dynamics similar to neuronal ion channels. I first met Carver at a workshop near Pittsburgh in 1984. I was impressed by his silicon retina, which was based on the same principles as the retina in our eyes, an ultralowenergy and light-weight alternative to traditional frame-based cameras. Christof Koch, Rodney Douglas, and I applied to the NSF to support

Participants in first Telluride Neuromorphic Engineering Workshop in 1995. Terry in the first row is next to Rodney Douglas and Misha Mahowald, two influential neuromorphic pioneers.

Terry with Sydney Brenner in La Jolla, 2008.

a Neuromorphic Engineering Workshop, first held in Telluride, Colorado, in 1995. This workshop has become the longest-running annual workshop sponsored by the NSF.

Sydney Brenner was a legendary biologist who was present when Francis Crick and Jim Watson first revealed their model of DNA in 1953 and later contributed to unraveling the genetic code. He made a singular gift to biology by developing C. elegans, a roundworm, as a new model system. Sydney was a Senior Fellow in the Crick-Jacobs Center for Theoretical and Computational Biology, which I directed at the Salk Institute. We had many dinners together when he visited La Jolla. I learned from Sydney and Francis how to ask key biological questions. When Sydney's health failed and he could no longer travel, I would visit him in Singapore. In 2017, I was the interlocutor with Sydney on a series of dialogs that ranged over his long career. The audience was predominantly younger scientists: one of the themes that ran through our discussion was that young scientists should be given more independence earlier in their careers. Our dialog was transcribed and published *In the Spirit of Science: Lectures by Sydney Brenner on DNA, Worms and Brains.*

I have had a long-standing interest in education and co-directed an NSF-sponsored 10-year Science of Learning Center at UC San Diego. I learned that education is not as much a scientific problem as a sociological one, with gatekeepers at every door.

One personal success was "Learning How to Learn," a free Massive Open Online Course (MOOC) that Barbara Oakley and I launched in 2013. Barbara is an engineer with a massive talent for teaching. This Coursera MOOC gives students practical advice on improving their learning skills based on our know ledge of how brains learn. Over the last ten years, over 4 million learners ages 10 to 90 have taken the course in 200 countries. We get a constant stream of grateful feedback, some telling us it has changed their lives. In over 40 years of classroom teaching, I have directly influenced only a few thousand students and received three grateful letters from students whose lives were changed, once by a podcast interviewer.

Tea time at the Computational Neurobiology Laboratory at the Salk Institute in 2008. Francis Crick frequently attended tea.

Looking back and moving forward

As computational neuroscience matured, neural models helped illuminate many brain functions, such as attractor states in the hippocampus and working memory in the cortex. New conferences were established based on computational brain models, including the Computational

Neuroscience Meeting (CNS) and the Computational and Systems Neuroscience Meeting (COSyNE). The pathway from physics, mathematics, and engineering to neuroscience became a highway. A new generation of highly motivated young researchers is analyzing a cornucopia of experimental data to uncover new insights into brain function.

The recent convergence of systems neuroscience with artificial intelligence is accelerating progress. This only became possible in the 21st century once these two areas of science and engineering converged on the same computational architecture: Massively large numbers of neural processing units highly interconnected by synaptic weights learned from data. The BRAIN Initiative and dramatic advances in AI are fueling this convergence. I was privileged to live through these remarkable events and contribute to both.

Terry and Beatrice visiting Iceland in 2004.

Terry with President Obama at the White House shortly before the BRAIN Initiative was announced in 2013.

At the top of a mountain in Lofoten, Norway in 2018.

Haim Sompolinsky

Professor, Harvard University and
Hebrew University

I was born in 1949 to David and Ilona Sompolinsky. David, a native of Denmark, was a medical student during World War II who became a Holocaust hero. Alongside the Danish Resistance and everyday citizens, he played a pivotal role in saving 700 Jews from extermination by the Nazis, orchestrating their escape to Sweden in October 1943.

In Sweden, David met my mother, Ilona, a young woman from Romania who had lost nearly her entire family in the Holocaust. She was sent to Sweden for rehabilitation at the war's end. There, my parents married and later moved back to Copenhagen, where I was born, the third in a family of ten children.

My parents.

In 1951, our family immigrated to Israel, and I was raised in Rishon Le'Zion, near the Asaf Harofeh Hospital, where my father, a microbiologist, headed the bacteriology laboratory in addition to establishing the Microbiology Department at Bar-Ilan University. Over a span of more than seventy years, he dedicated himself to both basic and clinical research in microbiology. Our household adhered to Jewish laws and customs, yet it was unusual in many respects. Through my father's life, I learned the balance between religious observance and scholarship, a deep love for Israel—the ancient and modern homeland of the Jewish people—a secular academic career, and a steadfast dedication to universal human values.

After completing high school, I devoted three years to intensive study at the Talmudic Academy the Ponevezh Yeshiva. During this time, I cultivated the ability to focus on challenging problems for extended periods and honed my skills in time management. Talmudic deliberations and debates taught me to question consensus and to continually approach old problems from new angles. However, I found the scope of study and worldview presented there to be monochromatic, narrow, and overly reliant on uncritical acceptance of authority.

After three years, I felt ready to chart my own course. I began undergraduate studies of physics and mathematics at Bar-Ilan University. Although I was inspired by the beauty and elegance of pure mathematics, the possibility of using mathematics to unravel the mysteries of the natural world exerted a stronger pull. Theoretical physics provided the perfect blend of mathematical theory and empirical application, leading me to pursue a PhD in theoretical condensed matter physics under the guidance of Professors Marshall Luban and Shlomo Havlin. My research on structural phase transitions in ferroelectric materials taught me how to develop and solve approximate theories in the statistical mechanics of condensed matter and how to validate these theories against experimental data.

Me with teacher.

During this period, I was fortunate to marry Elisheva, a remarkable woman from the Kirschenbaum family, born and bred in the US. Together, we established our home, creating a nurturing environment for our family. Elisheva has been a cornerstone of my life; marrying her feels like winning the lottery. Her intelligence, love, and unwavering support have been crucial to my ability to pursue a demanding career. Together, we have shared the joy of raising five children and the blessing of watching our family grow to include twenty-two grandchildren. Observing these incredible individuals develop and thrive has filled us with immense joy and wonder.

Upon completing my PhD and fulfilling my military service obligations, we moved to Boston for me to embark on postdoctoral research in theoretical condensed matter physics, under the mentorship of Prof. Bert Halperin at Harvard University. This transition marked a pivotal moment in my life, profoundly influencing my scientific identity. Bert proved to be an exceptional mentor, combining rigorous standards with deep knowledge and an extraordinary intuition for physics. His ability to navigate complex theories and interpret often bewildering experimental data profoundly affected my approach to scientific inquiry.


With my wife.


Prof. Bert Halperin, my postdoc mentor.

Choosing a research focus was challenging, as I sought a topic where the fundamental issues were still elusive. I was drawn to the topic of 'spin glasses,' a class of magnetic systems characterized by structural disorder like that found in ordinary glasses, which had thus far eluded a consistent theoretical explanation. I developed a novel approach to understanding spin glasses, shifting the focus from their static characteristics to their dynamics, uncovering the multi-scale dynamics and complex energy landscapes characterized by a hierarchy of "valleys within valleys." Collaborating with another postdoc, Annette Zippelius, we worked on a sophisticated statistical mechanics framework, Dynamic Mean Field Theory. Both the Spin Glass theory and Dynamic Mean Field Theory would later become invaluable in advancing the theory of neural circuits.

Upon concluding my postdoctoral research, I returned to Israel to take up a position as an Associate Professor of Physics at Bar-Ilan University. Concurrently, I expanded my international collaborations, particularly with Bell Laboratories, where I spent a significant portion of the following years.

Three years later, I was appointed as a Professor at the Racah Institute of Physics at the Hebrew University of Jerusalem, a move that was orchestrated by my Hebrew University colleagues, physicists Daniel Amit and Hanoch Gutfreund. Around 1983, we began collaborating on the application of spin glass theory to neural networks, marking our foray into neuroscience. Our interest was piqued by John Hopfield's 1982 seminal paper, on the emergent capabilities of neural networks and physical systems. Hopfield drew a tantalizing analogy between the storage of memory in neural networks and the behavior of magnetic spin systems and outlined a style of computation implemented through the attractor states of the neural dynamics, akin to energy minima in physical systems. We decided to try to apply statistical mechanics to study the Hopfield model. We soon realized that spin glass theory, my area of expertise, provided the perfect framework for this endeavor. Our research demonstrated that these networks constitute a unique class of systems. Unlike spin glasses, memory-induced plasticity generates significant structure, yet the complexity of the encoded data gives rise to more complex energy landscape than standard uniform systems.

In addition to charting the salient energy landscape of associative memory systems, this work provided a rigorous evaluation of memory capacity, convincing many physicists of the value in studying neural circuits and computations from a theoretical physics standpoint, especially through the lens of spin glass theory. These initial efforts were instrumental in shaping the nascent field of computational neuroscience and firmly established physics as one of its fundamental conceptual and computational pillars.

Amit, Gutfreund, and I began discussions on neural networks with Hebrew University neurobiologists. These early interactions highlighted the vast divide between the disciplinary cultures, evident in the differences in notation, terminology, and metaphors. What was deemed essential by one group was often seen as a mere detail by the other, and vice versa. Nevertheless, we gradually found common ground, with both sides stepping out of their comfort zones. The collaboration eventually grew to include computer scientists and psychologists, culminating in the establishment in 1992 of the Interdisciplinary Center for Neural Computation (ICNC), a new center for research and graduate training in computational neuroscience. at the Hebrew University.

Bell Labs another bastion of multidisciplinary collaboration. As an industrial research laboratory, it was unconstrained by the rigid disciplinary boundaries typical of academic institutions. During my visits to Bell Labs, I interacted with Nobel Laureate physicist Phil Anderson, who had developed a keen interest in biology, and neural network pioneer John Hopfield, among others. I also interacted with the exceptionally talented young biophysicists running experimental neurobiological labs, such as David Tank, Winfred Denk, and David Kleinfeld. Thus, what began as an exercise in theoretical physics transitioned over a few years into theoretical and computational neuroscience research.

With former junior colleagues.

During my time at Bell Labs, I met Sebastian Seung, a bright physics PhD student from Harvard, who later joined my lab at the Hebrew University as my first postdoctoral fellow. Sebastian has since emerged as a leading figure in computational neuroscience and one of the pioneers of the Connectomics field. I was incredibly lucky to have worked with a group of highly talented students and postdoctoral researchers, whose innovative ideas, skills, and dedicated work were indispensable to the success of my research. I take pride in the fact that many of them have become distinguished scientists, significantly contributing to the advancement of theoretical and computational neuroscience.

Throughout my career, I have had the privilege of collaborating with and learning from exceptional experimentalists. Bob Shapley from NYU opened the world of vision science to me, guiding me through the complex anatomy and physiology of the primate visual cortex. Markus Meister, then at Harvard, and I collaborated to decipher some intricacies of the retinal neural code. Florian Engert and Jeff Lichtman introduced me to the field of whole-brain neuroscience, specifically in zebrafish larvae. Eli Nelken and Adi Mizrahi at the Hebrew University provided comprehensive insights into auditory processing.

Computational neuroscience has gradually moved to the forefront of neuroscience research. This shift has been facilitated by the revolutionary development of new electrophysiological, optical, and magnetic imaging technologies that enabled neuroscientists to map out the structure of large neuronal circuits and monitor their activity, sometimes with cellular or even subcellular resolution. Consequently, neuroscience has entered the realm of Big Data. These expansive, high-resolution datasets have not only unveiled the vast complexity inherent in neuronal circuits but have also underscored the realization that in large brains, computation is a collective endeavor: it is distributed across vast populations of neurons, influenced by non-linear and contextual factors. Interpreting such systems necessitates the application of tools originally developed in physics, which emphasize emergent computation intimately linked to the collective dynamical patterns of large circuits.

Given my background in spin glasses, it was natural for me early on to turn my attention to the dynamics of neural circuits. My initial work aimed to address the well-documented, ubiquitous irregular patterns of neuronal activity, whose origins were then a mystery. Collaborating with Italian physics PhD student Andrea Crisanti and German physicist Hans-Jurgen Sommers, we found out that neural circuits with strong randomly connected neurons enter a chaotic state, generating intrinsic irregular spatiotemporal activity patterns. Chaos in random neural networks became a canonical model of the collective dynamics of unstructured neural circuits, with numerous applications in neuroscience as well as in artificial neural circuits studied in Machine Learning.

Later it became evident that the random circuit model failed to account for a crucial aspect of cortical and other brain circuits: the presence of distinct cell types, excitatory and inhibitory neurons, along with their corresponding distinct synaptic pathways. It took an additional eight years to formulate a comprehensive theory for such networks. Together with postdoctoral fellow Carl van Vreeswijk, we found that circuits composed of strong excitatory and inhibitory synaptic currents could dynamically counterbalance each other, achieving a state of equilibrium characterized by intrinsic chaotic spatiotemporal variability. This and other properties mirrored the collective behaviors observed in numerous neuronal circuits. Balance between excitation and inhibition emerged as a key principle underlying the development and function of the cortex in healthy states. Excitation-inhibition imbalance characterizes circuit dysfunctions in various neurological and psychiatric conditions.

During that period, mounting evidence pointed to another basic cortical mystery: neurons within local cortical circuits predominantly interact with each other, so that inputs from sensory organs play a surprisingly minor role. This apparently paradoxical finding prompted me to develop a simple neural circuit model, known as the ring model, wherein the interactions between pairs of neurons are modulated based on their functional distance from one another. By applying the concept of ‘symmetry breaking’ from physics, I demonstrated that the network can exhibit a continuous manifold of attractor states. Each stable state, or ‘bump state,’ is characterized by neuronal activity concentrated around a specific location on the ring, with different initial conditions leading to the stabilization of the bump at different locations. A weak external signal is sufficient to move the bump to a position aligned with it. The ring attractor model has since been generalized to higher dimensions, including 2D toroidal manifolds. Direct experimental validation of the manifold of bump states in a ring-like circuit came only recently with the observation of such networks in the fly navigation system, revealed through an exquisite series of functional and structural measurements. Additional support for this concept has emerged from recent discoveries of latent toroidal manifolds in the grid cell system of the rodent medial entorhinal cortex, as well as ring manifolds in its head direction system. Working out the dynamics of neural circuits suggests a view of the brain as a generator of a rich repertoire of spatiotemporal activity patterns serving myriad functions, from dreams and imagery to formulating hypotheses about the state of the world, as well as facilitating spontaneous thought and creativity.

Theoretical neuroscience has branched in multiple directions from the last decade of the 20th century through the first decade of the 21st century. One of them was research of learning in neural networks, spurred by the burgeoning interest from engineering and computer science. Pioneering the statistical physics approach was the seminal work of Elizabeth Gardner who demonstrated the applicability of statistical mechanics to the characterization of the collective states of the synaptic connection matrix in a neural network undergoing learning. Gardner theory sparked a wave of physics-inspired research into neural network learning, and it influenced my own recent work on neural manifolds. However, biological considerations often necessitated a different focus. One such effort was the development of learning rules for real neurons—which, unlike their artificial counterparts, communicate through discrete spikes. This included the investigation of unsupervised spike-timing dependent plasticity (STDP) by Larry Abbott and colleagues, and the supervised Tempotron learning rule developed by myself and Robert Gütig. Other unsupervised learning rules such as the ICA algorithm, developed by Terry Sejnowski and Tony Bell, introduced new methods for neural data analysis and offered fresh insights into the neural code. Reinforcement Learning, researched by Peter Dayan and collaborators, emerged as a powerful framework for examining the principles of goal-directed behavior, decision-making, and planning, illuminating the function of the brain's reward system.

This remarkable progress in the field laid a solid foundation for understanding the dynamics and functions of neuronal circuits and how they are shaped by experience. Computational and theoretical neuroscience has secured its role within the

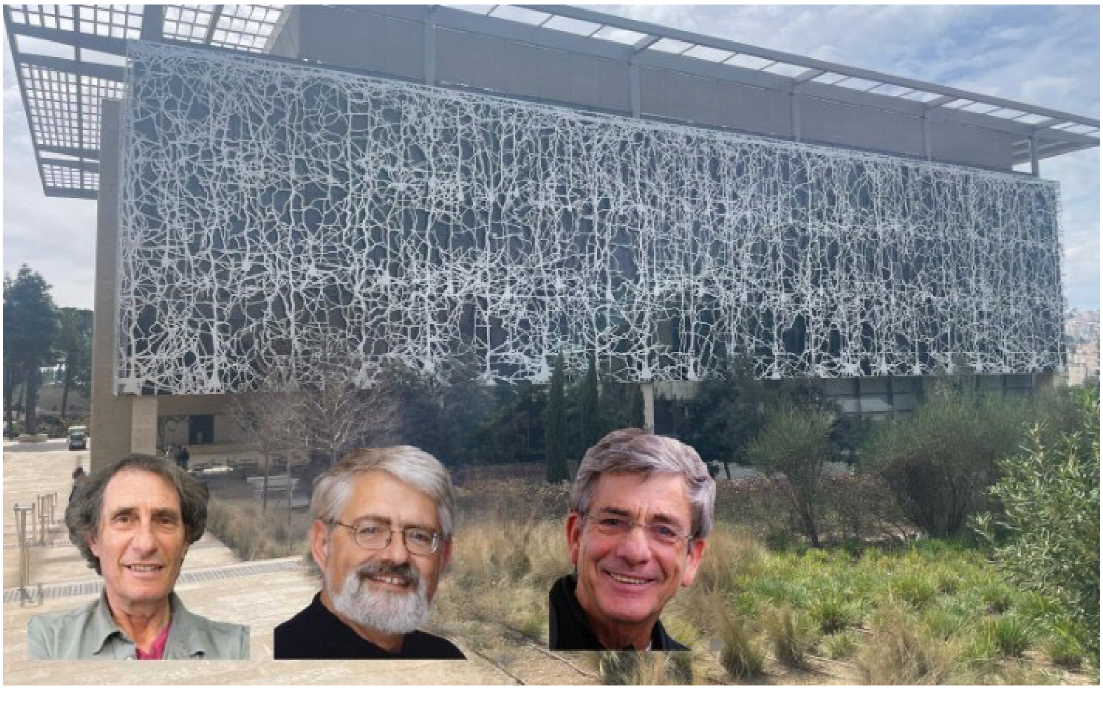

Building housing ELSC with its founders (Eilon Vaadia, right; Idan Segev, left).

With our children.

mainstream neuroscience, as evidenced by the establishment of interdisciplinary brain science centers worldwide. At the Hebrew University, together with outstanding scientists and visionaries, neurophysiologist Eilon Vaadia and theoretical neurobiologist Idan Segev, and supported by the university leadership and generous donations, we founded in 2009 the Edmond and Lily Safra Center for Brain Sciences (ELSC). This state-of-the-art facility is equipped with world-class research labs, core facilities, and a PhD program. ELSC champions multidisciplinary brain research with a strong emphasis on quantitative and theoretical approaches, featuring a uniquely designed PhD program aimed at cultivating the next generation of leaders in computational and theoretical neuroscience.

In the early years of the new millennium, I initiated longstanding collaborations with the neuroscience community at Harvard. I contributed to the neural circuit, behavior, and computation agenda of the new Center for Brain Science, led by Josh Sanes, drawing upon my experiences at the Hebrew University. I have served as a full-time professor there since my retirement from the Hebrew University in 2022. The recent establishment at Harvard of the Kempner Institute for the Study of Natural and Artificial Intelligence marked an exciting new phase of multidisciplinary study of cognition in brains and machines.

Recent AI research has shown that artificial neural networks with a prominent hierarchical (‘deep’) structure are capable of learning complex cognitive functions, sometime reaching human level performance. This breakthrough offers us tentative powerful neural network models for studying complex cognitive processing in large-scale brain networks, heralding a revolutionary era in theoretical and computational neuroscience. In the past decade, I focused my research towards developing new theories of neural processing that leverage deep networks and generative AI. Emerging from this and other studies is the utility of geometric concepts, such as neural manifolds and latent low-dimensional embeddings, in uncovering the regularities underlying neuronal representations in complex circuits processing complex signals.

This line of research is still in its early stages. As we progress, we anticipate identifying both similarities and differences between real and artificial networks, each discovery propelling forward our comprehension of brain function. Currently, significant disparities exist between the capabilities of AI systems and those of animals and humans. A deeper understanding of the brain is poised to yield even more advanced AI systems. Much has been discussed regarding the fragility of AI systems to specific perturbations, yet brains exhibit their own forms of vulnerability. It is

With our children, their spouses, and grandchildren at a family vacation in Holland.

my hope that AI systems will provide deeper insights into the complex mechanisms that link molecular and cellular pathologies to system dysfunctions observed in cognitive disorders.

I consider it a privilege to live in an era marked by remarkable advancements in our understanding of the human brain—how it underpins cognition, behavior and agency—and by the rise of machine intelligence, with its extraordinary promises and significant risks. These developments compel us to re-examine and refine our notions of individual and societal identity, and to redraw the lines between the sacred and the secular, the mundane and the sublime in our lives.

THE
BRAIN
PRIZE

Founded by the
Lundbeck Foundation