Cell

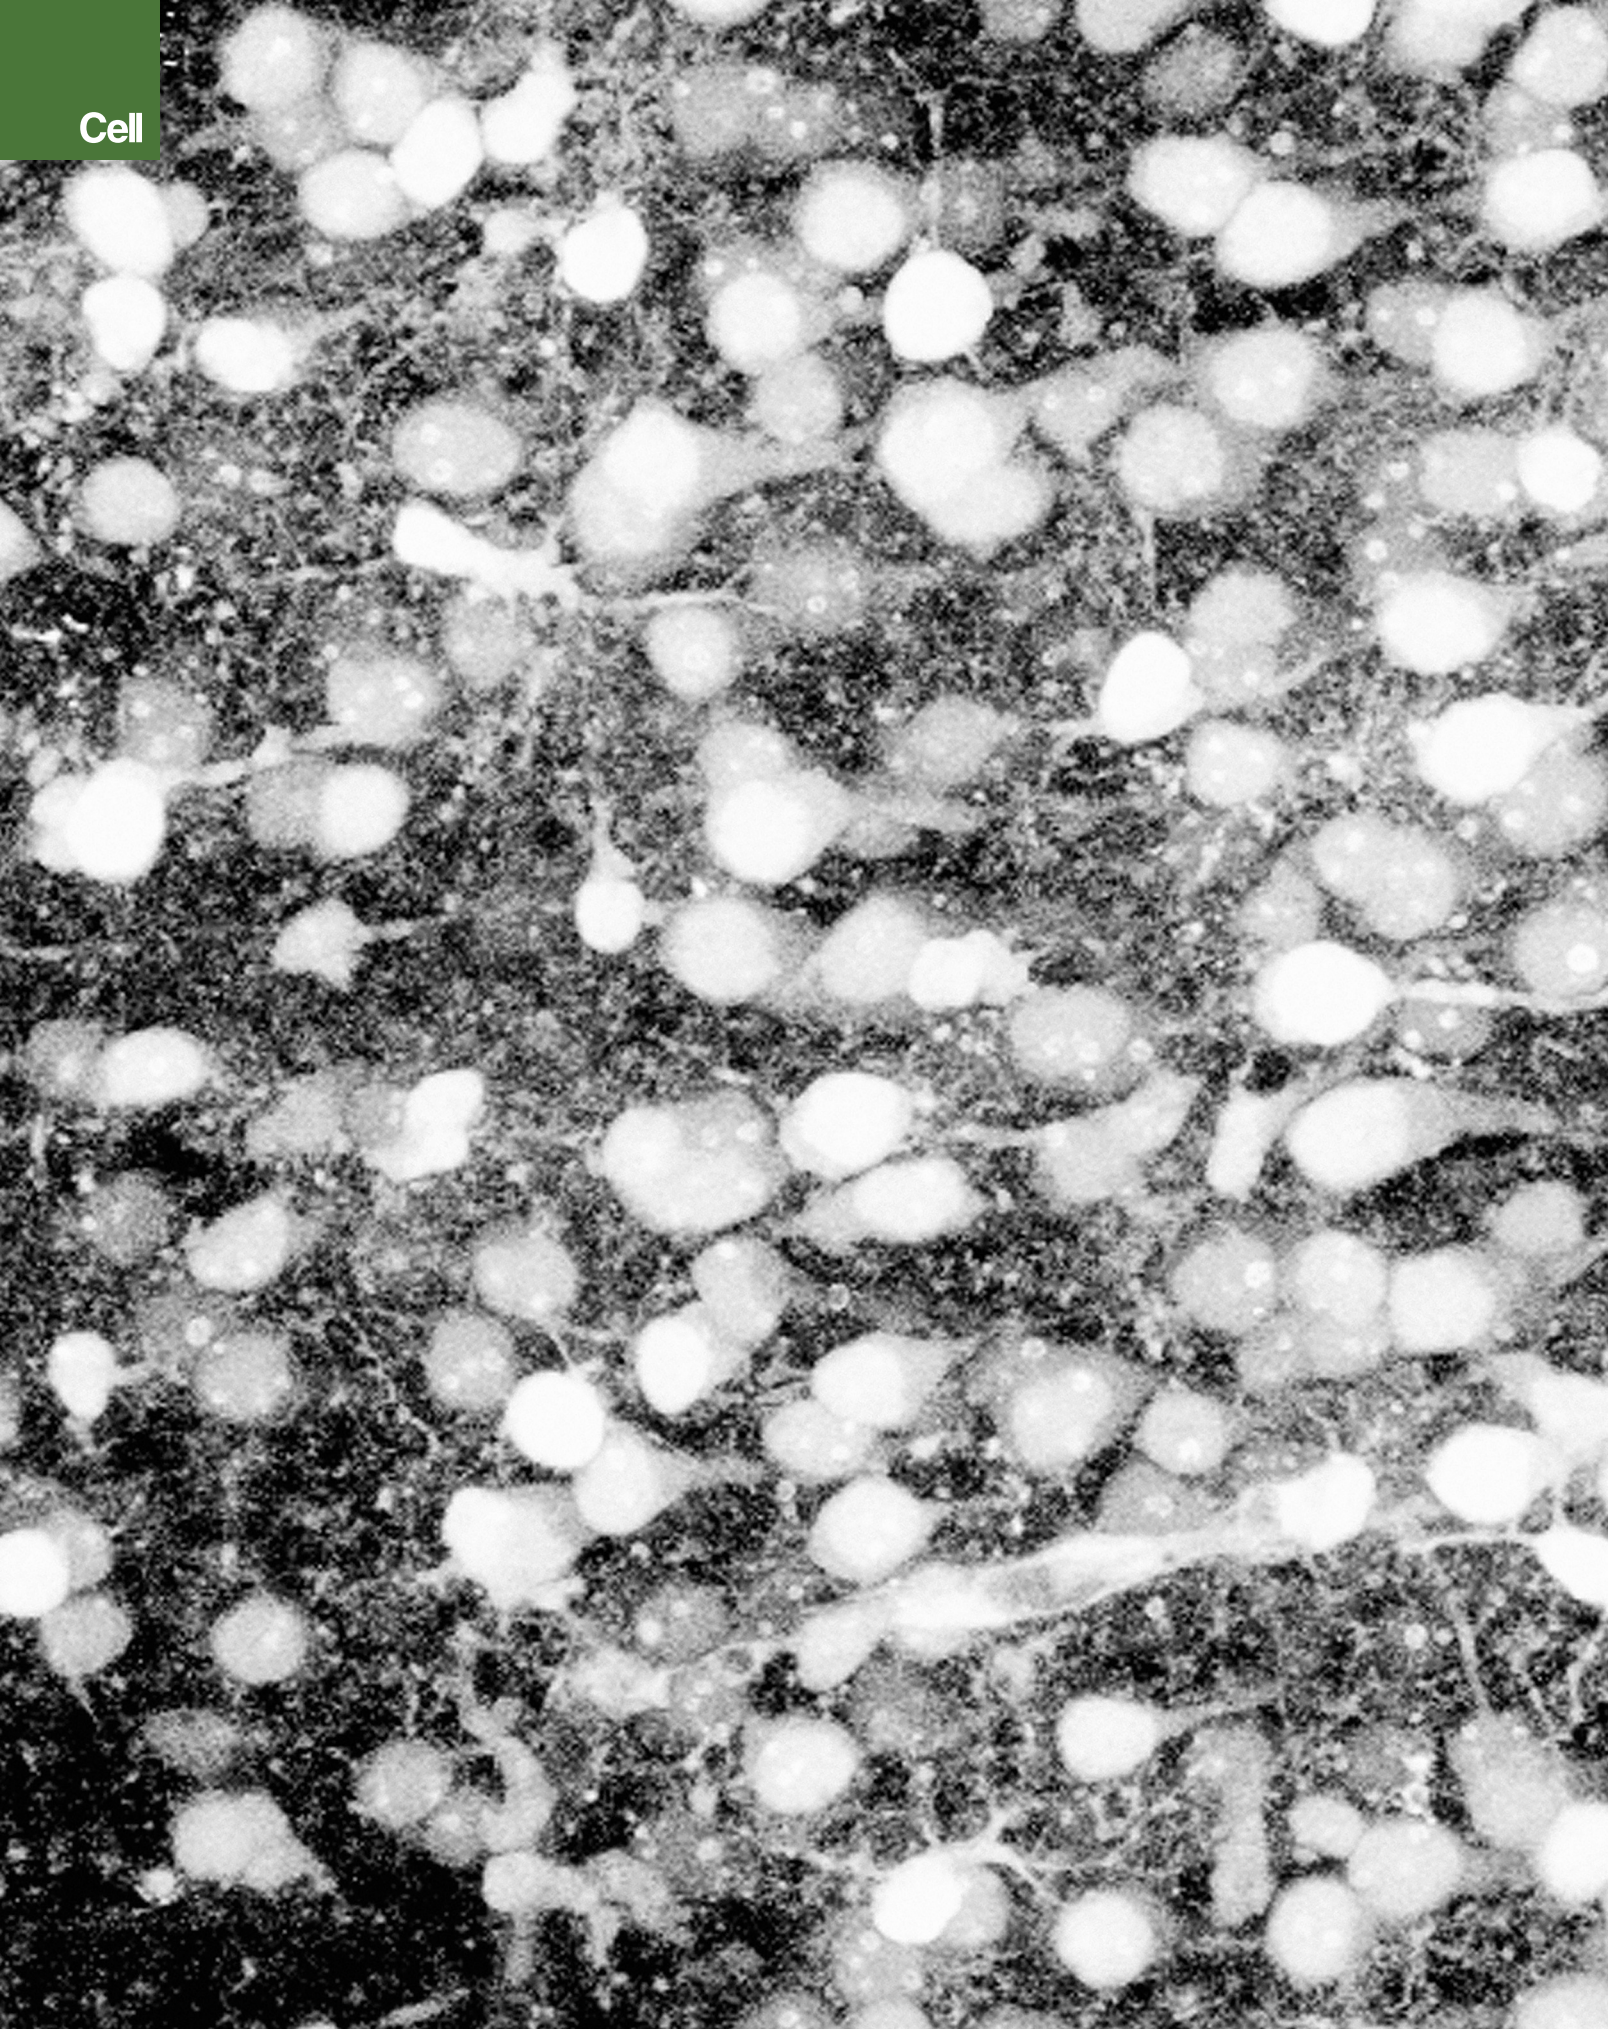

Cell

The Origins of the BRAIN Initiative: A Personal Journey

Rafael Yuste

The 25 of us were ensconced a cold rainy day in late summer of 2011 in Chicheley Hall, a bucolic country estate in Buckinghamshire, to discuss potential collaborations between neuroscience and nanoscience. We had been invited there under the auspices of four foundations: the Wellcome Trust, the Gatsby and Kavli Foundations, and the Allen Institute of Brain Science. In the last session of the meeting, I stood up and proposed the dream of developing novel technologies to record "every spike from every neuron" in a neural circuit. I made the argument that nervous systems were specifically designed by evolution to generate emergent states of activity, and, as Crick put it, that to continue to study the brain with single-neuron methods was similar to watching a TV screen one pixel at time. A vigorous debate ensued, with several participants arguing that this was impossible or that it would cost too much money or generate vast amounts of data that would drown us. George Church, one of the pioneers of the Human Genome Project (HGP), was sitting in the front row. Visibly upset, he stood up and said that in science, "nothing is impossible," and that those were the exact criticisms made against the HGP when it started, and added that not only was the HGP successful, but it also finished ahead of time and revolutionized biology and medicine. After the session was over, a small group of us—Church, Paul Alivisatos, Ralph Greenspan, Michael Roukes, and I—remained in the room, and energized by the criticism, we bonded. That same evening, we circulated among ourselves a rough draft of a white paper proposal that we called the Brain Activity Map (BAM), which described our vision of launching a large-scale scientific project of the size and scale of the HGP but focused on interdisciplinary technology development to record from and functionally manipulate every neuron in a neural circuit, such as the entire brain of an experimental animal or a cortical area of a human patient.

Fast forward to early 2013, and I am sitting at home with my wife and daughters, watching President Barack Obama give his annual State of the Union address. Suddenly, we hear the President announce to Congress the launching of a large neuroscience project, the same ideas that we had proposed, and even using some of our exact language. It was an unforgettable moment of my life. In the year and a half since Chicheley, our proposal had been adopted by the White House and formed the inspiration for its BRAIN Initiative (Brain Research through Advancing Innovative Neurotechnologies), or BRAINI. Now in its third year and enjoying bipartisan support from its beginning, the BRAINI funds more than 400 laboratories both in the US and abroad and has resulted in over 260 publications already. It is slated to last until 2025, supported with federal funds exceeding several billion dollars, of which close to $2 billion have been awarded or committed. The BRAINI has also helped inspire similar large-scale brain research projects around the world, including in Japan, Australia, Canada, China, South Korea, and Israel. In addition, a European project (The Human Brain Project) was launched at the same time as the BRAINI. The era of large-scale neuroscience projects is here.

So how did I end up in the middle of this maelstrom? Born in Madrid, as a Spaniard, I grew up under the towering influence of the Spanish neuroscientist Santiago Ramón y Cajal. I devoured his autobiography as a teenager, particularly the part where he gives detailed advice to younger scientists about everything, from how to cultivate the willpower to whom to marry. My mother ran a laboratory of clinical analysis in our neighborhood in Madrid, and I made my first money working in her lab, counting blood cells of patients under the microscope. My favorite subject in high school was biology, and my favorite part of biology was histology, the study of tissues using light microscopy. Inspired by Cajal, I daydreamed of being a neuroscientist, investigating the secrets of life with my microscope in a dark basement, and joining the small group of mankind's secret heroes, beautifully captured by *Microbe Hunters*, Paul de Kruif's influential book that my mother gave me for my birthday.

In pursuit of that dream, I attended medical school at the Autonomous University in Madrid and the Fundación Jimenez Diaz Hospital, where I first experienced science. After working as an

Two-photon calcium imaging of cortical circuits.

Participants in the Chicheley Meeting in September 2011 (Opportunities at the Interface of Neuroscience and Nanoscience). Photo credit: Max Liu (The Kavli Foundation).

undergraduate intern in the laboratories of Sols and Ferrús, learning enzymology and *Drosophila* genetics, I volunteered for two summers in Sydney Brenner's laboratory in the Laboratory of Molecular Biology of the Medical Research Council in Cambridge, England. I was lucky to have worked at that time in my life in that incredible environment, even though when Sydney heard I was interested in neuroscience, he put me right to work on bacteria! I was assigned to Leslie Barnett, Sydney's right-hand, who had previously worked for Francis Crick and who regaled me with stories of the early days of molecular biology. In an unforgettable moment, late one evening in the lab, Leslie told me the story of the triplet code experiment, when she, Crick, and Brenner carried out one of the most elegant experiments I have ever heard of. They systematically mutagenized phage strains with different numbers of hits of a single-base-deletion mutagen until the phenotype reverted back to wild-type. Every third hit. Three letters. Leslie told me they did this experiment in one afternoon, breaking for a pint of beer at the pub while waiting for the results. For me, that story was a turning point: so much for medicine! I was going into basic research, and my career goal was (and, to this day, still is) to do such an experiment and have a big smile on my face for the rest of my life.

After that seminal experience, I asked Sydney for advice as to where to do a PhD in neuroscience. As England was in the middle of the Thatcher years and suffering major cuts to science, Sydney advised me to move to the US, as Crick had just done and as he himself would do later. With his characteristic wit, echoing western movies, he told me, "Go west, young man." I applied to PhD programs in neuroscience in the US, and after another summer at Cuello's laboratory in McGill, I was admitted to Rockefeller University and joined the Wiesel Lab in 1987 as a graduate student.

Larry Katz's group in the Wiesel Lab at Rockefeller and calcium imaging setup in late 1988 (left to right: Alex Peinado, Larry Katz, and Rafael Yuste). Photo credit: Peter Pierce.

Torsten Wiesel, who had recently received his Nobel Prize for his classic work with David Hubel on cortical electrophysiology, had moved to Rockefeller from Harvard and built a lab as a mini-department focused on understanding the mammalian cortex. I was in heaven. More than 30 investigators, organized in 5 or 6 independent groups in a ''Scandinavian syndicalist fashion,'' as Torsten liked to say, studied all aspects of cortical structure and function, using a wide variety of methods: monoclonal antibodies, tissue culture, biophysics, brain slices, anatomical tracings, and electrophysiological recordings and psychophysics from cats and monkeys. Torsten provided us with financial support, space, equipment, supplies, and complete freedom. He was always there for advice on scientific or personal problems but didn't take any credit for our work. I encountered his sharp mind as he interviewed me in his office, and we spent 90% of the time discussing Spanish modern art before he accepted me into the lab, an unexpected payoff from years of going to all those modern art exhibits while growing up in Madrid.

At the Wiesel lab, I was immediately attracted to Amiram Grinvald, who was developing optical imaging methods to map the functional architecture of the cortex, and I asked if I could do my thesis in his group, imaging cortical activity with voltage sensitive dyes. He suggested instead that I work with Larry Katz, a young postdoc with Torsten who had developed a new method to make brain slices, and Amiram suggested that we explore calcium indicators rather than voltage dyes. High-affinity calcium indicators had recently been developed by Roger Tsien and were beginning to be used quite effectively as a biochemical tool to monitor the intracellular free calcium concentration in cultured cells. Larry and I hit it off personally, and I started experimenting with calcium indicators, exploring methods to label neuronal populations in brain slices with rhod-2 and fura-2. After many failed attempts, on December 8, 1988, I had a lucky break: Alex Peinado, a postdoc in the lab was making brain slices from developing rat cortex, and I asked him if I could borrow one for my staining experiments. I still remember looking down through the inverted microscope and seeing a field of blue fluorescent neurons. Every single neuron was stained! It had worked: developing cortical tissue could be labeled with incubations of calcium indicators.

However, the real breakthrough happened soon after; when I started to make time-lapse movies of these slices, I noticed that neurons were blinking spontaneously. I showed this movie to Larry, who told me, ''Rafa, it's about time you change the arc lamp in the microscope because it's old and it's flickering.'' That evening, I took more movies, but this time, I incubated the slices with tetrodotoxin (TTX), the sodium channel blocker, which got rid of the blinking. The next morning, I showed this movie to Larry and told him, ''Larry, this is really interesting, the arc lamp is sensitive to TTX.'' Laughing, Larry said, ''and you have a thesis!'' These experiments marked

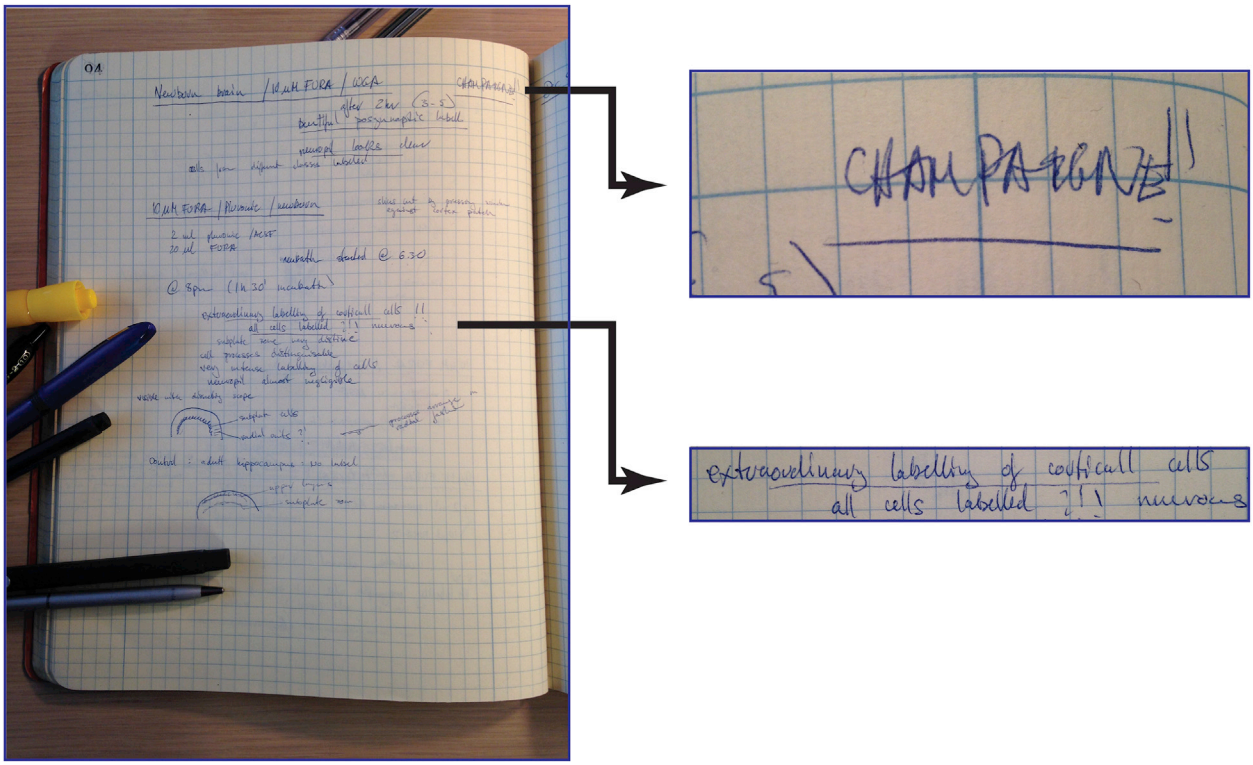

Lab notebook entry of first successful calcium indicator staining of entire neural circuits in December 1988. Note the "champagne" remark on top right and "extraordinary labeling…all cells labeled" comments in middle.

the beginning of the use of calcium indicators to monitor the activity of neuronal populations. I remember the vertigo I felt realizing that I could spend the rest of my life watching how neural circuits get activated and deciphering their function. Indeed, I have essentially spent my career doing that, as "functional histology," using calcium indicators to monitor cortical function. A functional version of Cajal's histological drawings, but with the Brennerian angle of using that data to "break the neural code."

Using calcium imaging, we discovered that brain slices from developing cortex (rats, mice, ferrets, and cats) showed spontaneous activity, with groups of neurons firing together. Our 1989 Society for Neuroscience abstract describing patterned calcium changes in cortical populations caught the attention of David Tank, at AT&T Bell Labs in New Jersey, who was using calcium imaging to study dendritic biophysics. David invited to me to Bell Labs to present my work and offered me a postdoc on the spot, which I took, also on the spot. Love at first sight! I was very impressed by Bell Labs and by the small but vibrant group of physicists that David had assembled in his division, developing novel methods in neuroscience. Though it was odd for me to go to a postdoc with New Jersey physicists rather than to one of the neuroscience power houses in glamorous California, as everyone had recommended, it seemed completely natural to me. I loved Tank's style and rigor, and this was the ideal place for me to train.

I spent four years at Bell Labs, learning biophysics and CCD imaging and carrying out a very intense collaboration with Winfried Denk, using his newly invented two-photon microscope to image calcium in dendritic spines. Our collaboration was meant to be: we worked in labs across the corridor; he had a microscope looking for a problem; and I had a perfect problem, dendritic spines, looking for a microscope. We lucked out: in spite of the longer wavelengths, the combination of two-photon excitation and calcium imaging was a winner, and our paper opened the floodgates for the joint application of these two methods. But perhaps the most influential moments of my time at Bell were the daily and animated lunch conversations of our group, where we often discussed models of brain function. While I defended receptive field models Tank and John Hopfield argued for neural networks. Over the course of these lunches, it became clear to me

the advantage of network models, where the computation is performed as an emergent property, over the traditional neuron doctrine, where the unit of function was the individual neuron. Having trained in the single-cell Hubel and Wiesel electrophysiology tradition, I walked out of Bells Labs a born-again neural network biophysicist.

The idea that neural circuits generate emergent states of function was not new. Already in the 1930s, Rafael Lorente de Nó, another Spaniard and a disciple of Cajal, hypothesized that neuronal circuits are dominated by recurrent excitatory connections ("internuncial chains") and have been designed by evolution to generate reentrant activity. He imagined reverberating patterns of activity that would cascade through neural circuits and that would remain activated even in the absence of sensory stimulation. Lorente's ideas were taken a step further by the Canadian psychologist Donald Hebb, who argued in 1948 that connected chains of neurons could be bound together into assemblies through synaptic learning rules. But it was Hopfield who, in a seminal paper in 1982, formalized a mathematical model of feedback neural networks and proposed that recurrently connected neuronal circuits generate emergent functional states. These so-called attractors—for example, a small group of neurons firing together—would naturally arise from a recurrent connectivity matrix and learning rules such as Hebb's and could implement memories or, more generally, solutions to a computational problem.

Chains, assemblies, ensembles, and attractors are some of many different theoretical models for how neural circuits may operate as emergent systems, where individual neurons function together as larger structures. In fact, to me, the core problem in neuroscience today is precisely to understand these multineuronal structures, as they could represent the "codons" of the neural code. As a field, we need to identify these ensembles—or whatever you want to call them—, test if they are indeed functional units and learn the ABC's of their biology: how are they anatomically defined, by which mechanisms they are formed and altered, how are they related to behavior and mental states, how do they appear in development and evolution, how do they become altered in disease, etc. If true, these models could have a profound explanatory power and could provide a simple description of the logic of neural circuits, grounding the field much as the Hodgkin & Huxley model of the action potential grounded cellular neuroscience. Although we are not there yet, I am convinced that neural networks models are poised to replace the neuron doctrine as a new paradigm for neuroscience.

After leaving Bell Labs, I joined Columbia's Biological Sciences Department, where I have remained ever since as an independent investigator for over two decades, continuously supported by the National Eye Institute. I feel proud and privileged to work in the same department that hosted, among others, scientists of the stature of Wilson, Morgan, Dobzhanski, Levinthal, and, more recently, Chalfie and Frank. Here, I found a wonderfully supportive haven with colleagues that infused into me the intellectual freedom and collaborative spirit that dominates Columbia's School of Arts and Sciences, which made my interactions with colleagues in the Physical Sciences and Engineering quite natural. The Columbia Medical School's world class neuroscience community, located in a nearby campus, also brought me in as one of their own and plugged me into the pulse of the field.

Perhaps not surprisingly, the work in my lab has been a merger of my Rockefeller and Bell Labs training: calcium imaging, two-photon, dendritic spines, and cortical circuits, first in slices and then *in vivo*, all while continuing to build optical methods to measure or manipulate neuronal activity. The heart of our work is to decipher neural codes by the description of the structure and function of neuronal ensembles, aiming to understand their anatomical and biophysical basis, inspired by Hopfield and following in the footsteps of pioneers like Lorente, Llinás, Braitenberg, and Abeles. We have generated abundant evidence that these multineuronal functional units are for real. In a string of papers, we have studied cortical ensembles from slices to awake-behaving mice, found similar patterns in spontaneous and sensory-evoked activity, used two-photon optogenetics to imprint them, and are now testing their role in behavior. We also find that these patterns are altered in mouse models of schizophrenia and autism, so they could be important in the pathophysiology of mental and neurological diseases. Finally, inspired by a recent visit from Brenner, who encouraged us to test the neural network paradigm in simpler nervous systems, we have also embarked, at Columbia and also at the Marine Biological Laboratory in Woods Hole, on a systematic study of the complete activity of the nerve net of the cnidarian Hydra, arguably one of the simplest nervous systems. This career summary helps frame the

First White House visit with Miyoung Chun, Tom Kalil, and George Church in December, 2011.

reason for my involvement in the BRAINI–an endeavor that represents the natural merger of neuroscience and physical science approaches, with optical imaging methods as its core, something that I have pursued for the past 30 years.

So how did the BAM get started? Since I came to Columbia, I tried to help realize the MRC/Rockefeller/Bell Labs dream fusion of biological and physical sciences by interacting and collaborating with physicists, chemists, engineers, and statisticians. As part of this effort, I started an interdisciplinary biology (iBio) seminar series, where biologists and physical scientists presented their work "to the other side." This also led to my moving to a new interdisciplinary science building at Columbia, where I helped assemble and coordinate a small groups of labs interested in particular in developing methods for neuroscience, building Columbia's Neurotechnology Center. But the early steps toward the BRAINI happened at a small meeting at the Institute of Medicine in Washington in January, 2008, to which I had been invited to discuss the future of neuroscience. In my talk, I argued for neural networks and illustrated, with our own data, the power of the merger of the physical sciences and neurosciences. I encouraged neuroscience policy makers to develop a systematic program to bring in physical sciences, drawing inspiration from their methods, theory, quantitative rigor, and also even publication style, with public posting of preprints as the main form of dissemination. My proposal went nowhere, but this exercise in public discussion of science policy for the future of neuroscience planted the seed in me.

Fast forwarding, then, four years to the Chicheley meeting, it is natural that I took the opportunity to advance a similar proposal and coordinated our white paper on it with the small group of like-minded researchers. Our BAM white paper was sent to the White House's Office of Science and Technology Policy (OSTP) right after the meeting, following the advice of Miyoung Chun from the Kavli Foundation, who played a key role by shepherding us throughout the entire process. Luckily, we had a friend at OSTP: Tom Kalil, himself the son of two neuroscientists, for whom the idea of a large-scale brain project resonated. In fact, our white paper was so well

received, we got comments back the same day we emailed it. This started a dialog with OSTP that lasted a year and a half, during which the BAM group was invited to visit the White House several times to deliver increasingly more detailed proposals for a large-scale neuroscience project. In parallel, we organized small workshops, sponsored by the Kavli Foundation, where we discussed our ideas with a wider group of scientists, most of whom enthusiastically signed on. From the original Chicheley 5, we quickly became 10, and then 20. The workshops involved over 100 experts, and their input helped to strengthen the project. To engage the entire field, we also published a platform paper in *Neuron* in 2012, where the scientific case for the BAM was laid out. Meanwhile, OSTP kept sending us good vibes and increasingly pointed questions and advice. These were exciting times, and we started to get the feeling that launching the BAM could happen. A critical meeting occurred on August 7, 2012, at the Eisenhower Office Building of the White House, shortly after the Curiosity Rover had impeccably landed on Mars. We actually met in a room that had a model of the Rover on a table. Kalil opened the meeting and, pointing at the Rover, said something like this: "This has been a good week for us: we just landed this baby on Mars. Now, let's talk about the brain: Why can't we cure schizophrenia? What's missing from neuroscience, what do you need? But don't think about yourselves or your labs or your universities. Think instead about mankind. Can we be remembered in the future as the generation that solved the brain?"

These were inspiring words, and they reflected the general spirit of the discussions with the White House: no egos, team effort, zero politics, all substance, and focused on the best way to help mankind. That meeting, to me, exemplified the best in human nature.

That August meeting was a turning point, and the focus of subsequent discussions with OSTP were not on questions of "if" but on questions of "how." Then, suddenly, OSTP went cold for a few months, and we started to worry that they had dropped our proposal. We were supposedly competing with many other proposals, including space explorations, clean energies, and health initiatives among others. But something was moving inside, because, as mentioned, President Obama did finally announce the BRAINI in the State of the Union address in February, 2013. By April, we got invited to the White House for an announcement by the President in a cryptic email from a whitehouse.gov account, which some of my colleagues found in their spam filter. I got the email on my phone during a family vacation, while visiting Jamestown in Virginia, and shouted out loud. My wife told the startled tour guide that I had just got invited by the President! The next Monday, together with over 200 people—a mixture of leaders in neuroscience, academia, and federal funding agencies—we were ushered into the East Wing conference room, which had only been used for a science-related event once before in history. Obama came in, and in a brief speech, he announced the launching of the now called BRAIN initiative with the triple goal of understanding the mind, helping cure brain diseases, and spurring new economic opportunities. Francis Collins, the NIH Director, then explained how an independent committee of experts ("the dream team"), led by Cori Bargmann and Bill Newsome, would be charged with charting its course. It was a bittersweet moment for me: as proud as I was of having, with my colleagues, driven to the point of fruition a large-scale initiative focused exactly on the issues which we thought were most critical, it was also sad to be left behind in the implementation. But we realized it was the best for the project that the issue of ongoing leadership be open to the entire community—it had never been our intention to grab reins but simply to launch this grand endeavor, and in this, we had eminently succeeded. Since none of the originators were involved anymore, there was no conflict of interest. We should all be proud of how cleanly and fairly the BRAINI was generated: it reflects the spirit of adventure, opportunity, team effort, and fairness of the US.

The dream team was indeed first class. They met and worked hard and incorporated input from the community, including us, in open town hall meetings. Their BRAIN 2025 report—released on June 5, 2014—succinctly distilled the challenges of circuit neuroscience and proposed seven specific goals, keeping the focus on the development of novel neurotechnologies but also adding emphasis on anatomical and molecular mapping and on human neuroscience. Opening it up to the wider neuroscience community, it also incorporated an aim focused not on technology development but on the application of new technologies to specific neuroscience problems. While agnostic on emergent properties, the BRAIN 2025 report aligned very well with our thinking and the arguments we made in the *Neuron* paper. Their key goals—the interdisciplinary development of technologies to systematically record and manipulate the activity of neural circuits and to

At the Global Brain Project Conference in New York with Cori Bargmann in September, 2016. Photo credit: Mario Morgado.

analyze that data with novel computational methods—were ours. With that report in hand, the BRAINI started. Funds were secured by Congress, application calls were opened, and grants were awarded. The train started to move.

To avoid seeming self-serving, I didn't apply for BRAINI funding and have funded my lab through other sources. But after a couple of years on the sidelines, I ended up participating on the NIH BRAINI multicouncil advisory committee through a circuitous route. The National Eye Institute Council, on which I served, voted me to represent it. During these BRAINI meetings, besides serving the NEI, I brought back some of the original spirit of the BAM, fostering a more ambitious agenda with large-scale projects or facilities (''Brain Observatories''). I am particularly proud of our unanimous decision in the committee to open up all NIH BRAINI funding to any investigator in the world: this is an example for how science funding should work in a globalized 21st century. The BRAINI is growing and enjoying good health, run by a passionate group of program officers with sound scientific advice, and—miraculously, given our current level of political polarization—continues to enjoy strong bipartisan support in Congress. Funding has been generous, and Obama and Biden's parting gift of the NIH 21st Century Act guaranteed $1.5 billion for the next decade of the NIH BRAINI. The expectation is that Congress will continue to support the BRAINI with annual line budget items, enabling not just NIH, but also NSF, DARPA, and IARPA to continue their serious engagement in the BRAINI. The Department of Energy may get involved; this seems to many of us to be a natural opportunity for the US national labs, as they could provide the large-scale project expertise and roots in physical sciences and engineering to help revolutionize neuroscience, both scientifically and sociologically.

My last involvements with the BRAINI have been focused on its internationalization and on helping to provide it with an ethical framework, two issues I am quite passionate about. Together with Cori Bargmann and with help from the NSF and the Kavli Foundation, I organized an international

meeting on the sidelines of the opening of the UN General Assembly last fall to help coordinate the different brain projects around the world. It was particularly befitting for Cori and me to organize this meeting as a joint effort between the BAM and the BRAIN 2025 groups. Speakers from the US, European, Chinese, Japanese, Korean, Australian, and Canadian brain projects participated in this historic meeting, as well as representatives from France, Germany, Spain, UK, Israel, Russia, Palestine, Cuba, and Iran, with the overall goal of setting up a common collaboration and bridges to better coordinate all these projects, something which is starting to happen.

In addition, with the bioethicist Sara Goering, I also recently organized two neuroethics workshops at Columbia's Morningside campus with experts in neurotechnology, artificial intelligence, medical ethics, and the law, also with support from NSF and the Kavli Foundation. Our realization is that the merging of novel neurotechnologies and artificial intelligence (AI) will profoundly impact basic human qualities such as our sense of identity, agency, and our mental privacy. This could profoundly alter society, leading to a augmentation of our mental and physical abilities. Because of this, we think that it is imperative that the development and application of these new neurotechnologies and AI follow a set of ethical guidelines, akin to the Belmont report, which led to protocols for the protection of human subjects in research. Our proposed guidelines could be democratically endorsed and implemented by professional scientific committees with participation from government and civil society. Their mandate would be to promote and protect a set of basic ''Neurorights,'' protecting our minds and our society from the potentially negative effects of these new hardware and software tools and ensuring that this amazing technology is used for the benefit of mankind. Of all the projects I have been involved with, this set of neuroethical guidelines, if successful, could be the most transformative.

As I think back over my career, I have enormous pride seeing the wide-ranging impact of the optical methods that we contributed to neuroscience and our role in inspiring the BRAINI and its ramifications. But what really excites me looking forward is the potential of the new data gathered with these new technologies to rigorously test or demonstrate neural network models and decipher the neural code. We will see. Maybe these ideas are just plain wrong, or maybe the neural code question is ill-posed, and neural circuits could operate otherwise. That's the beauty of science: no one has the truth. In the end, while I feel lucky and privileged to have lived in the neuroscience Zeitgeist, the image that keeps coming to mind about myself is the one of a high school kid from Madrid who loved looking through the microscope and daydreamed of making discoveries.