

Edge

To arrive at the edge of the world's knowledge, seek out the most complex and sophisticated minds, put them in a room together, and have them ask each other the questions they are asking themselves.

https://www.edge.org/conversation/ed_boyden-how-the-brain-is-computing-the-mind

Printed On Fri April 29th 2016

CONVERSATION : CONVERSATIONS

How the Brain Is Computing the Mind

A CONVERSATION WITH Ed Boyden [2.12.16]



The history of science has shown us that you need the tools first. Then you get the data. Then you can make the theory. Then you can achieve understanding.

ED BOYDEN is a professor of biological engineering and brain and cognitive sciences at the MIT Media Lab and the MIT McGovern Institute. He leads the Synthetic Neurobiology Group. [Ed Boyden's Edge Bio Page](#)

Includes *Edge* Video and Audio

HOW THE BRAIN IS COMPUTING THE MIND

How can we truly understand how the brain is computing the mind? Over the last 100 years, neuroscience has made a lot of progress. We have learned that there are neurons in

the brain, we have learned a lot about psychology, but connecting those two worlds, understanding how these computational circuits in the brain in coordinated fashion are generating decisions and thoughts and feelings and sensations, that link remains very elusive. And so, over the last decade, my group at MIT has been working on technology, ways of seeing the brain, ways of controlling brain circuits, ways of trying to map the molecules of the brain.

At this point, what I'm trying to figure out is what to do next. How do we start to use these maps, use these dynamical observations and perturbations to link the computations that these circuits make, and things like thoughts and feelings and maybe even consciousness?

Ed Boyden

46:12

Edge

There are a couple of things that we can do. One idea is simply to go get the data. A lot of people have the opposite point of view. You want to have an idea about how the brain computes, the concept of how the mind is generating thoughts and feelings and so forth. Marvin Minsky, for example, is very fond of thinking about how intelligence and artificial intelligence can be arrived at through sheer thinking about it.

On the other hand, and it's always dangerous to make analogies and metaphors like this, but if you look at other problems in biology like, what is life? how do species evolve? and so forth, people forget that there are huge amounts, centuries sometimes but at least decades of data that was collected before those theories emerged.

Darwin roamed the Earth looking at species, looking at all sorts of stuff until he wrote the giant tome, *On the Origins of Species*. Before people started to try to hone in on what life is, there was the tool development phase: people invented the microscope.

People started looking at cells and watching them divide and so forth, and without those data, it would be very hard to know that there were cells at all, that there were these tiny building blocks, each of which was a self-compartmentalized, autonomous building block of life.

The approach I would like to take is to go get the data. Let's see how the cells in the brain can communicate with each other. Let's see how these networks take sensation and combine that information with feelings and memories and so forth to generate the outputs, decisions and thoughts and movements. And then, one of two possibilities will emerge.

One will be that patterns can be found, motifs can be mined, you can start to see sense in this morass of data. The second might be that it's incomprehensible, that the brain is this

enormous bag of tricks and while you can simulate it brute force in a computer, it's very hard to extract simpler representations from those datasets.

In some ways, it has to be the former because it's strange that we can predict our behaviors. People walk through a city, they communicate, they see things, there are commonalities in the human experience. So that's a clue; that's a clue that it's not an arbitrary morass of complexity that we're not going to ever make sense of. Of course, being a pessimist, we should still always hold open the possibility that it will be incomprehensible. But the fact that we can talk in language, that we see and design shapes and that people can experience pleasure in common, that suggests that there is some convergence that it's not going to be infinitely complex and that we will be able to make sense of it.

Biology and brain science are not fundamental sciences in the sense that physics is. In physics, there are particles and there are forces, and you could write down a very short list of those things. But if you're thinking about the brain and the brain is going to have these cells called neurons, and the neurons have all these molecules that generate their electrical functions and their chemical exchanges of information, those are encoded for by the genome. In the genome, we have, depending on who you ask, 20,000- to 30,000-odd genes, and those genes produce gene products like proteins, and those proteins generate the electrical potentials of neurons and they specify at least some parts of the wiring. The way that I look at it is we're going to want to understand the brain in terms of these fundamental building blocks, and we can always try to ignore some detail, this concept of the abstraction layer.

Can we ignore everything below a certain level of description and just focus on the higher level concepts? But modern neuroscience is now almost 130 years old, since the neuron was discovered, and so far, the attempts to ignore below certain levels of description have not yielded universally accepted and explanatory theories of how our brains are computing our thoughts or feelings or movements.

The way that we approach things is pretty radically different from the past. The premise that I launched my research group at MIT on was that we needed new technology. The reason people are shying away from these very, very detailed measurements of brain function, getting the deep data, was because we didn't have the tools. The history of science has shown us that you need the tools first. Then you get the data. Then you can make the theory. Then you can achieve understanding. No theory with no technology. It's very difficult to know that you've solved it.

Before Newton's Laws, there were lots of people like Kepler and Galileo who were watching the planets, and they had decades and decades of data. Why don't we have that for the brain? We need tools for the brain like the telescope and the microscope, and now, we need to collect the data, ground truth data, if you will, where we can see all those cells and molecules in action, and then, we're going to see a renaissance in our ability to think of and learn about the brain at a very detailed level, but to extract true insight from these datasets.

Let's think for a second about the hypothesis that biology is not a fundamental science. If you think about books like *The Structure of Scientific Revolutions*, this and other attempts to explain the path of science, we often have these models: here's my hypothesis, somebody comes along and disproves it, and if it's a big enough disproof, you get a revolution.

But let's think about biology: suppose I want to figure out how a gene in the genome relates to an emergent property like intelligence or behavior or a disease like Alzheimer's. There are so many genes in the genome, most hypotheses are probably wrong just by chance. What are the chances that you got the exact gene that's most important for something? And even if you did, how do you know what other genes modulate it? It's an incredibly complicated network.

If you started thinking of how different genes of the genome, how their products interact to generate functions in cells or in neurons or networks, it's a huge combinatorial explosion. Most hypotheses about what a gene is doing, or especially what a network of genes is doing, much less a network of cells in the brain, they're going to be incorrect. That's why it's so important to get these ground truth descriptions of the brain.

Why can't we map the circuits and see how the molecules are configured, and turn on or off different cells in the brain and see how they interact? Once you have those maps, we can make much better hypotheses. I don't think the maps of the brain equal the understanding of the brain, but the maps of the brain can help us make hypotheses and make them less assumption-prone, make them less likely to be wrong.

One thing that I hope a circuit description of the brain will help us understand about humanity is, as we know from psychology, there are countless unconscious processes that happen. One of the most famous such experiments is you can find regions of the brain or even single cells in the brain that will be active even seconds before people feel like they're making a consciously-willed decision. That leads to what you might maybe slightly jokingly say, we have free will but we're not conscious of it. Our brains are computing what we're going to do, and that we're conscious after the fact is one interpretation of these studies.

What I suggest though is that if we peek under the hood, if we look at what the brain is computing, we might find evidence for the implementation or the mechanisms of feelings and thoughts and decisions that are completely inaccessible if we only look at behavior, or if we only look at the kinds of things that people do, whereas if you find evidence that something you're about to do, something you're about to consciously decide, your brain already has that information in advance. Wouldn't it be interesting to know what's generating that information? Maybe there are free will circuits, quote, unquote, in the brain that are generating these decisions.

We know all sorts of other things that occur, feelings that our brains are generating, and we have no idea about what's causing them. There are very famous examples where somebody who has an injury to a part of their brain that is responsible for conscious vision, but you tell them when you see something, I want you to have a certain feeling, or when you see something, I want you to imagine a certain kind of outcome, and people will have that occur even though they're not consciously aware of what they're seeing.

There is so much processing that we have no access to, and yet, it's so essential to the human condition for feelings and decisions and thoughts, and if we can get access to the circuits that generate them, that might be the fastest route to understanding those aspects of the human condition.

I've been thinking a lot over the last decade primarily about the technology that helped us figure out what we need to understand about the brain in terms of circuits and how they work together. But now that those tools are maturing, I'm thinking a lot about how

we use these tools to understand what we all care about.

Up until now, we mostly have been giving our tools out to other neuroscientists to use. We've been focusing very much on technology invention, and other groups have been discovering profound things about the brain. I'll just give you a couple of examples.

There's a group at Caltech and they use one of our technologies, a technology that makes neurons activatable by pulses of light. They put these molecules into neurons deep, deep in the brain, and when you shine light, those neurons are electrically active, just like when they're normally being used. They found that there are neurons deep in the brain that trigger aggression or violence in mice, so they would activate these neurons and the mice would attack whatever was next to them, even if it was just a rubber glove.

I find it fascinating to think about something as ethically charged, as essential to the human condition, as involved with our justice system and all sorts of stuff, as violence. You can find a very small cluster of neurons that, when they're activated, are sufficient to trigger an act of aggression or violence. So of course, now, the big question is what neurons connect to those? Are they violence detectors? Oh, here is the set of stimuli that makes us now decide, oh, I should go attack this thing next to me even if it's just a glove.

And then, of course, where do these neurons project? What are they driving? Are they driving an emotion, and downstream of that emotion comes the violent act? Or are they just driving a motor command: go attack the glove next to you? For the first time, you can start to activate very specific sets of cells deep in the brain and have them trigger an observable behavior, but you can also ask, what are these cells getting, what are these cells sending messages to, and looking at the entire flow of information.

I'll give you another example that is fascinating. One of my colleagues at MIT, Susumu Tonegawa, trained mice on a learning task, so that certain neurons in the brain become activatable by light. They used some genetic tricks to do that. Now, what happens is those mice can be doing something else much later, they shine light on the brain, and those neurons, the ones that had been activated earlier when they were learning, they get reactivated and the mice make a memory recall. It's like they were there in the earlier place and time.

That's interesting because for the first time, they can show that you can cause the recall of a specific memory, and now they are doing all sorts of interesting things. For example, you can activate those cells again, and let's say that's a happy memory; let's say it's associated with pleasure or a reward. They have shown that that can have antidepressant effects, that you can have an animal recall, a memory when you shine light on certain neurons, now the memory that is recalled triggers happy emotions; this is how they interpreted it. And that can counteract other stressors or other things that make the animal normally feel not so good.

Literally, hundreds and hundreds of groups are using this technology that we developed for activating neurons by light to trigger things that are of clinical and maybe even sometimes philosophical interest.

~ ~ ~

I studied chemistry and electrical engineering and physics in college, and decided that I cared about understanding the brain. To me, that was the big unknown. This will seem

kind of cheesy, but I started thinking about how our brains understand the universe, and the universe, of course, gives us things like the laws of physics upon which are built chemistry and biology, upon which is built our brain. It's kind of a loop. I was trying to think about what to do in a career; I thought, what's the weak point in the loop? And it seemed like the brain was very unknown.

I was very impressed by people who would go build technology to tackle big problems, sometimes very simple technology. All the chemists in the 1700s and 1800s who built ways of looking at pressure and volume and stoichiometry, without that, it's inconceivable that we would have things like the Periodic Table of the Elements and quantum mechanics and so forth.

What stuck out in my mind was you need to have that technological era, and that then gives you the data that you want, that then yields the most parsimonious and elegant representations of knowledge. And for neuroscience, it seemed like we had never gone through that technological era. There were bits and pieces, don't get me wrong, like electrodes and the MRI scanner, but never a concerted effort to be able to map everything, record all the dynamics, and to control everything. And that's what I wanted to do.

At the time I started graduate school at Stanford, I went around telling everybody I wanted to build technologies for the brain and to bring the physical sciences into neuroscience. A lot of people thought it was a bad idea, frankly, and I think the reason why was at the time, many people who are physicists and inventors were trying to build tools for studying the brain. But they were thinking forwards from what was fun for them to do, and not backwards from the deep mysteries of the brain.

The key insight that I got during graduate school was if you don't think backwards from the big mysteries of the brain, and you only think forwards from what you find fun in physics, the technologies you built might not be that important. They might not solve a big problem. What I learned was we have to take the brain at face value. We have to accept its complexity, work backwards from that, and survey all the areas of science and engineering in order to build those tools.

During the first decade that I've been a Professor at MIT, we have mostly been building tools. We built tools for controlling the brain, tools for mapping the detailed molecular and circuit structure of the brain, and tools for watching the brain in action. Right now, we're at a turning point; we're ready to start deploying these tools systematically and at scale. Don't get me wrong, the tools still need improvements to be equal to the challenge of studying the brain, but for small organisms like worms and flies and fish, or for small parts of mammalian brains, we're ready to start mapping them and trying to understand how they're computing.

The work progresses through primarily philanthropic as well as government grant funding. We have been very lucky that there has been a bit of an increase in people interested in funding high risk, high reward things. That's one reason why I'm at the MIT Media Lab, and you might ask why is a neuroscience Professor in the School of Architecture at MIT?

As we were discussing earlier, neuroscientists long had a deep distrust of technology, that technologies often didn't work, the brain was so complicated that the tools could only solve toy problems. When I was looking for a professor job, the search was hit-or-miss.

My collaborator, Karl Deisseroth and I had already published a paper showing we could activate neurons with light, a technology that we've called ever since "optogenetics," "opto" for light and "genetics" because it's a gene that we borrow from a plant to make the neurons light-sensitive. But a lot of people at the time were still deeply skeptical: is this the real deal or is this yet more not-quite working technology that will be a footnote? I went to the Media Lab to complain about how political and complicated academia was, and I was very lucky; they were wrapping up a failed job search and they said, "Why don't you come here?" And so I went, and we've been incubating a lot of neurotechnology there since then.

When I first got to Media Lab, a lot of people were deeply puzzled about what I would do there. Was I going to switch into, "classical publicly-perceived Media Lab technology," like would I have developed ways of having cell phones diagnose mental illness or other things like that? I wanted to get to the ground truth of the brain. In some ways, the Media Lab was a perfect place to start. We could incubate these ideas, these tools out of the cold light of day until they were good enough that neuroscientists could see their value. And that took several years.

It was about a three-year period until this started to get mainstream acceptance, and then, there was another three-year period where people said, wow, how do we get more technology, and that led to initiatives like the Obama BRAIN Initiative, which is an attempt to get widespread technology development throughout neuroscience.

The BRAIN Initiative started at the instigation of the Kavli Foundation. They were hosting a series of brainstorms about what nanoscientists and neuroscientists could do together, and Paul Alivisatos and George Church and Rafael Yuste and many people at that border were at these early sessions. And in late 2012, I was invited to one of these sessions where many inventors were invited and we started talking about maybe brain activity mapping is great and all, but the technologies might be much more broad than that; you might need more than just maps.

You might need ways to control the brain, ways to rewire the brain.

That was an interesting turning point because it went from activity mapping to broadly technology, and four or five months later, Obama announced this BRAIN initiative which, somewhat recursively, stands for Brain Research for Advancing Innovative Neurotechnologies, and they are now devoting tens to hundreds of millions of dollars a year, depending upon which year, to try to get more technology made to help understand the brain.

The BRAIN initiative now is run by different government agencies. They have their own priorities, so, for example, DARPA is very interested in short-term human prosthetics, for example, no surprise there. The National Science Foundation is interested in more basic science, and so forth. The different agencies have their own agendas now.

IARPA is involved. They are trying to do a hard push for short-term mammalian brain circuit mapping based upon existing technology, and sort of a small part of that more on the technology development side. Most of the money is on the application side. But we have some new tools that we think can be very, very helpful.

~ ~ ~

Companies are great if you can work hard and be smart and solve the problem. But if

you're tackling something like the brain, or the biggest challenges in biology in general, a lot of it's serendipity. A lot of it is the chance connections when you bring multiple fields together, when you connect the dots, when you kind of engineer the serendipity and make something truly unpredictable, and that's hard to do if you have closed doors. That's hard to do if you don't allow open, free collaboration.

Our group is very big; I think we're the second biggest research group at all of MIT. But we work with probably about 100 groups, people who are genomics experts and chemistry experts and people making nanodiamonds and all sorts of stuff. The reason is that the brain is such a mess and it's so complicated, we don't know for sure which technologies and which strategies and which ideas are going to be the very best. And so, we need to combinatorially collaborate in order to guarantee, or at least maximize the probability that we're going to solve the problem.

You want to have academia for that serendipitous ability to connect dots and collaborate, and you want companies when it's time to push hard and just get the thing done and scale up and get it out the door. What I would hope to engineer in the coming maybe decade or so are hybrid institutions where we can have people go back and forth because you might need to have an idea that would go back and forth a bit until it matures.

I'll give you an example. We're building new kinds of microscopes and new kinds of nanotechnologies to record huge amounts of data from the brain. One of our collaborators was estimating that soon some of these devices we're making might need some significant fraction of the bandwidth of the entire internet in order to record all the brain data that we might be getting at some point. Now, we need some electronics, right? We need electronics to store all the data and computers to analyze the data. But that's an industrial thing.

It's much easier to get that done in a company than in academia because people in industry can turn the crank and make incredible computers, so we started a collaboration. A small startup here in Cambridge, Massachusetts, does these computers with us. Now we're working on the nanotechnologies, and that fusion of two different institutional designs allows us to rapidly move faster than companies alone or academics alone. These new hybrid models are going to be essential to balance the need for luck and the need for skill and ability.

The thing that I'm excited about also is how do we get rid of the risk in biology and medicine? Most medicines, most strategies for treating patients, they are found in large part by luck. How do we get rid of the risk? We talked a bit about how there are fundamental sciences like physics, and then, you have higher order sciences like biology. Medicine also might have different scientific methods for different kinds of disease. We have made huge inroads against bacteria and viruses because of antibiotics, because of vaccines. Why have these been so successful? It's because we're trying to help our body fight a foreign invader, right? But if you look at the big diseases, the ones that nobody has anybody clue what to do about, there are brain disorders, a lot of cancers, autoimmune conditions, these are diseases where it's our body fighting ourselves, and that's much harder because you can't just give a drug that wipes out the foreign invader because the foreign invader is you.

How do we understand how to de-risk the tough parts of medicine? We have to think about drug development and therapeutic development from a different point of view. The models that give us new antibiotics and new vaccines and so forth might not be quite

right for subtly shifting the activity levels of certain circuits in the brain, for subtly tuning the immune system to fight off a cancer but not so much that you're going to cause an autoimmune attack, right?

One thought is, well, if it's your body fighting yourself, what you want is very deep knowledge about the building blocks of those cells and how they're configured in the body. The basic premises behind ground truthing the understanding of the brain might be also right what we need in order to de-risk medicine, in order to understand how cells and organs and systems go awry in these intractable disorders. That's something I've been thinking a lot about recently as well: how do we de-risk the goal and methodology and path towards curing diseases?

There was just a study released about how taking a drug from idea to market can cost \$2.5 billion now. And if you look at the really tough diseases like brain diseases, like cancers and so forth, the failure rate to be approved for human use is over 90 percent.

This got me thinking that maybe this is the same kind of intellectual problem as why we don't understand how brain circuits compute thoughts and feelings. We have these large 3D systems, whether it's a brain circuit or a cancer or the immune system, and knowing how to tweak those cells, make them do the right thing, means finding the subtle differences that make those cells different from the normal cells in our body. I've been thinking a lot about how we can try to take these tools that we've been developing for mapping the brain, for controlling the brain, for watching the brain in action and applying it to the rest of medicine.

~ ~ ~

I can tell you about a collaboration that we have with George Church. George's group for about fifteen years now has been trying to work on a technology called in situ sequencing, and what that means is can you sequence the genetic code and also the expressed genes, the recipes of cells, right there inside the cells?

Now, why is that important? It's important because if you just sequence the genome, or you sequence the gene expression patterns after grinding up all the cells, you don't know where the cells are in three-dimensional space. If you're studying that brain circuit and here is how information is flowing from sensation into memory regions towards motor areas, you've lost all the three-dimensionality of the circuit. You just have ground up the brain into a soup, right? Or for a tumor, we know that there are cells that are by the blood vessels, there are stem cells, there are metastasizing cells; if you just grind up the tumor and sequence the nucleic acids, you again have lost the three-dimensional picture. A couple years ago, George's group published a paper where they could take cells in a dish and sequence the expressed genes.

That is, you have DNA in the nucleus, that expresses in terms of RNA, which is the recipe of that cell, and the RNA then drives all the downstream production of proteins and other biomolecules. The RNA is sort of in-between the genome and the mature phenotype of the cell. It's kind of the recipe. George's group was sequencing the RNA. I thought that was amazing: you could read out the recipe of a cell.

Now, there was a tricky part: it didn't work well in large 3D structures like brain circuits or tumors. Our group had been developing a way of taking brain circuits and tumors and other complex tissues and physically expanding them to make them bigger. What we do to make the brain or a tumor bigger is we take a piece of brain tissue and we chemically

synthesize throughout the cells, in-between the molecules, around the molecules, in that piece of brain, a web of a polymer that's very similar to the stuff in baby diapers. And then, when we add water, the polymer swells and pushes all the molecules apart, so it becomes big enough that you can see it even using cheap optics.

One of my dreams is you could take a bacterium or a virus and expand it until you can take a picture on a cell phone. Imagine how that could help with diagnostics, right? You could find out what infection somebody has just by making it bigger, take a picture and you're done.

We started talking with George: what if we can take our sample and expand it and then run their in situ sequencing method—because sequencing, of course, is really complicated. You need room around the molecules to sequence them. This is very exciting to me, if we can take stuff and expand it and then use George's technology to read out the recipes of the cells, we could map the structure of life in a way.

We can see how all the cells look in a complex brain circuit, or in a tumor, or in an organ that's undergoing autoimmune attack like in type 1 diabetes. That's one of the things that excites me most is this in situ sequencing concept. If we can apply it to large 3D structures and tissues, we might be able to map the fundamental building blocks of life.

Our current collaboration with George's group has been focused very much on small pieces of tissue that we have: mouse brains probably, other model organisms in use in neuroscience. But we know that if they work in those systems, they'll probably work in human tissues as well. Imagine we get a cancer biopsy from somebody, we use our group's technology to expand it physically, making everything big enough to see, and then, we can go in and use George's in situ sequencing technology to read out the molecular composition.

When we first published the idea of expanding something, a lot of people were very skeptical about it. It's a very unconventional way of doing things. To convince people that it works, we went down [the following] line of reasoning: a design method. When we synthesized the baby diaper-like polymers inside the cells, we would anchor through molecular bonds specific molecules to the polymer, and then we would wipe up all the rest. We can use enzymes and so forth to chop up the rest. That way, when we expand the polymer, our molecules that we care about are anchored and move apart, but the rest of the structure has been destroyed or chopped up so that it does not impede the expansion. That's a key design element.

One way to think of this is—chemistry is a way of doing fabrication massively in parallel. So suppose that I want to see two things that are close together, like my two hands here. But of course, lenses cannot see very, very small things, right, thanks to diffraction. So what if we took my two hands and anchored them to these expandable polymers and then destroyed everything else? There might be a lot of junk here we don't care about. We add water and the polymer swells, moving my hands along with it until they're far apart enough that we can see the gap between them. That's the core idea of what we call expansion microscopy where we take the molecules in a cell or the molecules in a tissue, a brain circuit or a tumor, and we anchor those molecules to a swellable polymer. When we add water, the molecules we care about, the ones we've anchored—that we've nailed to the polymer, as it were, have moved apart until they're

far apart enough that we can see them using cheap, scalable, and easily deployed optics like you could find on an inexpensive microscope or even a webcam.

After we published our paper on expanding tissues, a lot of people started to apply them. For example, suppose you wanted to figure out how the cells are configured in a cancer biopsy. You can take the sample and if you look at it under a microscope, you can't see the fine structures, but if you blow it up and make it bigger, maybe you could see the shape of the genome; maybe you could see that one cell is extending a tiny tendril, too tiny to see through other means, and maybe that's the beginning of metastasis.

A lot of people are trying to use our technology now for seeing things that you just can't see any other way, and we're finding a lot of interest not just from brain scientists because now you have a way of mapping brain circuits with nanoscale precision in 3D, but also from other brain-like problems: tumors and organs and development and so forth where you want to look at a 3D structure but with nanoscale precision.

We've spun out a small company to try to make kits and maybe provide this as a service so that people can use this widely. Of course, we've also put all the recipes on the Internet so people can download them, and hundreds and hundreds of groups have already started to play with these kinds of tools.

We want to make the invisible visible, and it's hard to see a 3D structure like a circuit that might store a memory or a circuit in the brain that might be processing an emotion, with the nanoscale resolution that you need to see neural connections and the molecules that make neurons do what they do.

The fundamental limit on how fine we can see things is related to a technical parameter called the mesh size; that is basically the spacing between the polymer chains. We think that the spacing between the polymer chains is about a couple nanometers; that is, around the same size as a biomolecule. If we can push all the molecules away from each other very evenly, it's like drawing a picture on a balloon and blowing it up: you might be able to see all the individual particles and building blocks of life, but you know what? We have to validate the technology down to that level of resolution. So far, we have validated it down to about a factor of ten bigger than that, in order of magnitude. But if we can get down to single molecule resolution, you could try to map the building blocks of living systems. We haven't gotten there yet.

I've been amazed at how fast neurotechnology has started to move. Ten years ago, we had relatively few tools for looking at and controlling the brain, and now, ten years later, we have our optogenetic tools for controlling brain circuits, this expansion method for mapping the fine circuitry, and also, we have developed 3D imagining methods that basically work the way that our eyes work to reconstruct 3D images of brain high speed electrical dynamics.

In the coming fifteen years, two things are going to happen and a third thing, might happen. One thing that will happen is that our ability to map the fine details of neural circuits and see high speed dynamics and control it will probably be perfected; that might happen as soon as five years from now but definitely within fifteen years, I would predict that.

The second thing is that we're going to have some detailed-enough maps of small neural circuits that maybe we could even make computational models of their operation. For example, there is a small worm called *C. elegans* that has 302 neurons; maybe we can

map all of them and their molecules and their dynamics and perhaps we can make a computational model of that worm. Or maybe a slightly larger brain: the larval zebrafish has 100,000 neurons, mice have 100 million—ballpark—and humans have 100 billion. You can see there are some multistage logarithmic jumps there that we have to make.

The speculative thing is that we might have some tools that might let us look at human brain functions much, much more accurately. Right now, we have so few tools for looking at the human brain, there is functional MRI which lets you look at blood flow that is downstream of brain activity, but it's very indirect and it's very crude. The time resolution is thousands of times slower than in brain activity, and the spatial resolution, each little block that you see in these brain scans contains tens to hundreds of thousands of neurons, and we know that even nearby neurons can be doing completely different things.

What we most need right now, I would say, is a method for imaging and controlling human brain circuits with single cell, single electrical pulse precision, and the jury is out on how that could happen. There's lots of brainstorming. I haven't seen any technology generated so far that can provably do it although there's lots of interesting speculation. That's something I would love to see happen and we have started to work on some ideas that might allow you to do it.

There's a lot of speculation about whether there are quantum effects that are necessary for brain computations. At body temperature, it's very likely that quantum effects, if any, are going to be very, very short-lived, maybe much shorter than the kinds of computations that are happening in the brain. It's quite possible that if such effects are important, we would need far more powerful tools to see them, or perhaps you can explain all of the biophysics of neurons known to date, for the most part, with completely classical models.

The thing that I loved about working on the quantum computation project, this was with Neil Gershenfeld back in the day, was this greater philosophy of how information and physics are linked. There are many theories of fundamental physical principles of computation; there is even the phrase, "it from bit," where people talk about the fundamental thermodynamic limits of how information processing occurs in physical systems. For example, there are so many bits associated with a black hole, there is, based upon temperature, a fundamental amount of information that might be encoded in a specific transition. The brain for the most part is operating, because it's at body temperature and all that, far above those physical fundamental limits in terms of information processing.

On one level, the most parsimonious models of the brain are analogue because we know that there are different amounts of transmitters being released at synapses, we know that the electrical pulses that neurons compute can vary in their height and in their duration. Of course, if you dig deep enough, you could say, well, you could just count the neurotransmitters, you could count the ions, and it becomes digital again, but that's a much more detailed level of description that might not be the most parsimonious level because you had to count and localize every single sodium ion and potassium ion and chloride ion. Hopefully, we don't have to go that far. But if we need to, we would probably have to build new technologies to do that.

My co-inventor, Karl Deisseroth, and I both won Breakthrough Prizes in Life Sciences for our work together on optogenetics, this technology where we put molecules that are light

sensitive into neurons and then we can make them activatable or silence-able with pulses of light.

Our groups have sent these molecules out to literally thousands of basic as well as clinically interested neuroscientists, and people are studying very basic science questions like how is a smell represented in the brain? But they're also trying to answer clinically relevant questions like where should you deactivate brain cells to shut down an epileptic seizure? I'll give you an example of the latter since there is a lot of disease interest.

People have been trying to shut down the over excitable cells during seizures for literally decades, but it's so difficult because which part of the brain and which cells and which projections? It's such a big mess, right, the brain? So a group at UC Irvine has been using our technologies to try to turn off different brain cells or even to turn on different brain cells, and what they're finding is that some cells, if you activate them, can shut down a seizure in a mouse model. But still, who would have thought that activating a certain kind of cell would be enough to terminate a seizure? There is no other way to test that, right, because how do you turn on just one kind of cell?

What they did was there are certain classes of cell called interneurons, and they tend to shut down other cell types in the brain. What this group did is they took a molecule that we had first put into neurons about a decade ago, a molecule that, kind of like a solar panel, when you shine light on it, will drive electricity into the neuron. They delivered the gene for this molecule so that it would only be on in those interneurons, none of the other cells nearby, just the interneurons. And then, when they shine light, these interneurons will shut down their neighboring cells, and they showed you could terminate a seizure in a mouse model of epilepsy.

That's interesting because now, if you could build a drug that would drive those cells, maybe that would be a new way of treating seizures, or you could try to directly use light to activate those cells and build a sort of prosthetic that would be implanted in the brain and activate those cells near a seizure focus, for example.

People are exploring both ideas. Could you use our optogenetic tools to turn on and off different cell types in the brain to find better targets, but then, treat those targets with drugs? Or could you use light to activate cells and directly sculpt their activity in real-time in a human patient? The latter, of course, is much higher risk, but it's fun to think about for sure. And there are a couple companies that are trying to do that now.

When we were talking about the Breakthrough Prize, I thought about the little speech I gave—they give you thirty seconds, but I thought about it for several weeks because I feel like there is such a push to cure things, a push to find treatments, but in some ways, by forcing it to go too fast, we might miss the serendipitous insights that are much more powerful.

I'll give you an example: in 1927, the Nobel Prize in Medicine was given to this guy who came up with a treatment for dementia. What this person did is, he would take people with dementia and he would deliberately give them malaria. Remember this is the greatest idea of its time, right?

Now, why did it work? Well, malaria causes a very high fever. At that time, dementia was often caused by syphilis, and so, the high fever of malaria would kill the parasite that causes syphilis. Now, in 1928, one year later, antibiotics started to come online, and of course, antibiotics have been a huge hit and syphilis-related dementia is almost

unheard of nowadays.


The rush to get a short-term treatment, I worry, can sometimes cause people to misdirect their attention from getting down to the ground truth mechanisms of knowing what's going on. It's almost like people often talk about we're doing all this incremental stuff, we should do more moon shots, right? I worry that medicine does too many moon shots. Almost everything we do in medicine is a moon shot because we don't know for sure if it's going to work.

People forget. When they landed on the moon, they already had several hundred years of calculus so they have the math; physics, so they know Newton's Laws; aerodynamics, you know how to fly; rocketry, people were launching rockets for many decades before the moon landing. When Kennedy gave the moon landing speech, he wasn't saying, let's do this impossible task; he was saying, look, we can do it. We've launched rockets; if we don't do this, somebody else will get there first.

Moon shot has gone almost into the opposite parlance; rather than saying here is something big we can do and we know how to do it, it's here is some crazy thing, let's throw a lot of resources at it and let's hope for the best. I worry that that's not how "moon shot" should be used. I think we should do anti-moon shots!

What's Related

People



Ed Boyden

Neuroscientist, MIT Media Lab; Leader, Synthetic...

Beyond Edge

[MIT Media Lab](#)

[Synthetic Neurobiology Group](#)

Topics

[Conversations](#)

John Brockman, Editor and Publisher

Russell Weinberger, Associate Publisher

Nina Stegeman, Associate Editor

Contact Info: editor@edge.org

Edge.org is a nonprofit private operating foundation under Section 501(c)(3) of the Internal Revenue Code.

Copyright © 2016 By Edge Foundation, Inc All Rights Reserved.