

CortexCast

Season 3, Episode 4

Building Intelligence Beyond Biology: In Conversation with Ed Boyden

Introduction

Ivan: What would it take to connect the life of the mind to the machinery of the brain? For this, we may need to look beyond thought, memory, and behaviour at the surface, and instead trace them down to the molecules, circuits, and physical laws from which they emerge. With a sufficiently precise map of this hidden architecture, neuroscience may one day move from simply observing the brain to repairing it, augmenting it, and perhaps even totally rethinking what intelligence itself can become.

In this episode, Yicheng Su, co-president of Oxford University Cortex Club, talks to Ed Boyden, Professor of Neurotechnology at MIT and one of the most influential figures in modern neuroscience. Over several years Professor Boyden has pioneered transformative tools including optogenetics and expansion microscopy, technologies that have completely reshaped how scientists observe, control, and understand the brain.

The conversation today, however, moves beyond the history of these discoveries and instead explores the deeper philosophy behind them. We discuss Boyden's vision of *ground truth* in neuroscience: the search for the fundamental building blocks of the brain and the rules by which they interact. From molecularly annotated connectomes to AI-driven models of brain dynamics, from therapeutic possibilities to the ethics of augmentation and neural privacy, we ask what kind of science is needed to link molecules to minds.

Along the way, we also consider how young researchers should choose problems, how different disciplines can converge on neuroscience, and whether the future of intelligence will remain biological, become digital, or take forms we cannot yet imagine.

You're listening to CortexCast.

Interview

Yicheng: So first of all, Professor Boyden, thank you for speaking to us today. I think what really excites us is that your work doesn't just sit at the level of biology or technology but it sits at the intersection where ideas start changing what science can ask fundamentally. You're someone who changes what we can even think about for neuroscience. Many of us have been inspired how you've crossed fields from math and chemistry to electrical engineering to the neuroscience and then somehow turned those experiences into new tools that reshape how everyone else works. So I think today, rather than covering the history of optogenetics, or expansion microscopy or some of your other very exciting works, which you've already shared beautifully elsewhere, we're hoping that we can go a little bit deeper into how you think about building ideas, how disciplines talk to one another and what the next generation of neuroscience might need to really link molecules to minds. So if you're happy, I'd love to start there with your own path. How do you think learning across so many disciplines early on shaped the way you invent?

Ed: Well, I guess I think a lot about systems that are built on top of fundamental disciplines. I think of the brain as a computer science system that runs on chemistry according to the laws of physics and that's kind of how my training went. I started college pretty young. There was a program in Texas where you could skip some years of high school. I ended up going to college when I was 14 and my initial thought was to be a chemistry major. So that was my focus for my two years there at the Texas Academy of Math and Science. And I transferred to MIT and then ended up majoring in physics. And then just because I had a lot of friends interested in the topic, I also added computer science. This was the year that the web browser hit it big and everybody was interested in this Internet thing and figured it couldn't hurt to learn it. So as time went on, I got to learn kind of a large number of the fundamental sciences, right? The computer science and math axis, the physics and chemistry. And I started looking for problems to work on. And I tried working on problems like quantum computing or the origins of life. But those are fundamental problems, you could argue. They're really, really hard. And then the third area that I plunged into neuroscience, I realized was more of a systems problem. If we treated the brain as a computer science system running on chemical reactions, maybe we didn't have to have that stroke of luck that leads to discovering a black hole or the Hubble constant or gravitational waves, right? Which are truly multi-decade quest for a single thing, let's say. But rather, we could just try to decompose the brain into those chemical systems. And so that's what I decided to do.

Yicheng: So how did you sort of choose to pivot from the pre neuroscience era, in to say, I want to do a PhD in neuroscience?

Ed: Yeah, yeah. Well, I was working on my master's thesis on quantum computing. And along with some friends, we started MIT's autonomous submarine team. And eight weeks later, we won the world's first championship on that topic. So it seems like, you know, quantum computing is really hard. On the other hand, I mean, submarines have been built before. You know, did I want to spend my life doing that, too? And so, you know, despite the fact that we have a lot of fun, I think the people who contributed to the project did brilliant work. You know, it also felt that the submarine was an established paradigm. So I was really looking for a problem that was the right level of difficulty, right? You know, you don't want to try to land on the moon 500 years ago, but neither do we want to learn how to fold a piece of paper in half, right? You got to find the right middle ground. And so I ended up doing a neuroscience undergrad project with Michael Fee, who was a postdoc at Bell Labs at the time, and Sebastian Seung, who had just started at MIT. And yeah, it was it was great. I enjoyed it. We tried to make a model of how birds sing. I tried to construct a neural circuit model of that, and I realized that we couldn't really test my model because what data could test it, you know? And so that's why I decided that maybe all those forays into physics, chemistry and computer science could help me ultimately to make models of things like consciousness and cognition. But, you know, I could also apply those skills to the ability of tools to get me there.

Yicheng: So I think in some of your past interviews or writings, you've stressed a lot about math and then like the word, well actually probably two words, ground truth. So I think it sort of follows up to the next question I wanted to ask. So I think there's this idea that the brain is optimized for survival rather than for the truth. So if we have all the sensory, auditory, vision sort of inputs and the brain is sort of evolutionarily optimized for us to just to use the minimal amount of energy and to maximize our chances for survival rather than this being the actual

ground truth of what the world is. So if this is correct, what kind of tools or experiments do you think could let us test whether our current perceptual basis is the limiting kind of truth that we can discover? And I think a deeper follow up is that do you think it would even matter?

Ed: I guess by ground truth, I mean something very different. By ground truth, I mean for a scientific field, what are the fundamental building blocks and how do they interact? So physics, right? Before quantum mechanics, physics was very difficult. After quantum mechanics, you could build lasers and transistors and microchips and telecommunications and the list goes on and on and on, right? I mean after the microchip even, you could drop out of school and start Apple, Microsoft, Facebook, you know, some of the biggest companies of all time. So you see this huge reduction in what I like to call science risk. The other, you know, chemistry went through a similar transformation, right? With the periodic table and the molecular bond and now you can make all sorts of things, right? Polymers and ceramics and all sorts of things much more easily than if you don't understand it. So I'm a big believer in the idea of ground truth, meaning the building blocks and how they work together. And for the brain, I think that in the end has to be the biomolecules, right? We're made of biomolecules. Biomolecules interact with each other through contact or through fields. And, you know, if they contact, they can cut and or, you know, form bonds or change shape. In response to fields, of course, we get all sorts of interesting things, especially in the brain with neural activity and electric fields. So by ground truth, I mean, can you understand the brain in terms of all the parts and how the parts work together?

Yicheng: So it is more like sort of, for example, some of your work, you build tools and through those tools, there are these little blocks that allow us to build more work on top of that to find more ground truth. Is that sort of what you're saying?

Ed: Well, I think there is just the one ground truth, which is the parts of how they work together. Once you have those, you can try to understand things in terms of it. I don't know if there are additional ground truths. Maybe there are. You just have to be lucky, right? You can't count on those. You know, the genome, right? The ground truth would be the four letters A, T, C and G. Now, the genome forms three dimensional loops. You have all sorts of interactions between chromosomes. Are any of those complex phenomena reliable enough and describable enough and repeatable enough and so forth that you could call a ground truth? Maybe, but it's also possible that we're not there yet, right? We don't know what's fundamental beyond the ATCG. It's quite possible that there are a lot of fundamental things that are there, but we just don't know them yet. One example, I guess, would be the genetic code, right? We have a triplet code where three bases will specify amino acid. So that, you could argue, is a higher level ground truth that rests on top of the lower level ground truth. Of course, it took many decades to show that it was universal and that it does apply in all the cell types and in, you know, all the cell species and so forth.

Yicheng: That's a really interesting idea. So I think you went on to talk about how some ground truth can be more elemental or fundamental than the others. So although people want to search for those sort of ground truths, maybe some are more easily falsified or they already work in certain contexts. And do you think if we have maybe one ground truth over here and then another ground truth somewhere else, do you think it would be a good way to try to work to see if there's like something in the middle? It could be a less, it could be like a shallower level of ground truth, but it could still reveal some useful information.

Ed: I guess I need to know what you mean by ground truth. You seem to be using it like a generic term for fact or truth. But by ground truth, I mean fundamental building blocks. What do you mean by ground truth? Because maybe I don't understand what you mean by it.

Yicheng: Well, I think initially my ground truth, my initial definition of ground truth is how we sort of perceive the world. But then I think after you explain your version of it, I think we're more aligned on the terms that they are just like fundamental dogmas of how things work. And using those, like for example, E equals MC squared or F equals MA , and using those we can build on more truths. Do you think that's similar to what you're saying or is that still a little bit different?

Ed: I mean, by ground truth, I mean just what are the building blocks and how they interact. So for E equals MC squared, you could say we have electrons and electrons interact through electromagnetic fields, right? I mean, and then you said F equals MA , you have masses and they interact through forces, right? I mean, that's how I would translate what you're saying into ground truth. But I think it's very dangerous to say that all truths are ground truth.

Yicheng: Yeah, that's really interesting. Maybe I think if we just sit on that a little bit and then we can come back to it later, if that's all right.

Ed: I mean, if you go back to the physics and chemistry analogies, right, I think you'll see that that is the case, right? You know, in physics, you can have very high-level theories of gas, right, like the ideal gas law, and that can be very convenient, right? PV equals nRT . But it's not really ground truth, right? You don't have to know anything about atoms or molecules. And if you go deeper, which is the field of statistical mechanics, then you can get really, really interesting insights that you cannot get simply by thinking at the high level, right? You can get the Boltzmann distributions and you can look at kinetic energy and you can derive temperature and so forth, and from velocity of particles and distributions and so forth. And all that's very, very difficult to do, if not impossible, if you only confide yourself to the macroscopic observables that can be measured without precision equipment.

Yicheng: How do you think this sort of ground truth influenced how you work?

Ed: Well, I mean, it's the laser focus of our group, and when I switched into neuroscience, I outlined this 50-year-ish quest of four parts. Build the tools, get the data, build the models, and then ideally, can we understand philosophy, right? What is subjective experience? Why do we do what we do and feel what we feel? And so, throughout that 50-year vision, I guess, the idea is that the ground truth should be the guiding vision. Once you have that, then, analogous to physics after quantum mechanics or chemistry after the periodic table of the molecular bond, you can have an explosion of bottom-up thinking, because once you know the building blocks and how they interact, you can build anything or understand anything. And yes, feel free to make approximations if you want, but just make sure that they are consistent with the ground truth, right? What Bohr might have called the correspondence principle back in the day. For physics, in the limit, quantum mechanics in the limit of classical reality should look like classical reality, right? Because we don't see weird things happening as we walk around on the street, right? So yeah, it's fine to have such outcomes that you judge and compare it to everyday reality, but they shouldn't be considered as the fundamental truth, I think.

Yicheng: So, zooming out just a little bit, if we follow the trajectory of your work, as you described, there's the mapping the brain and then there's the perturbing it. Do you think there's a theoretical path to, for example, something like high-fidelity brain-to-brain data transfer, because I think this would touch upon the topic of subjectiveness. Do you think this is something that's possible in the future if we understand the building blocks well enough?

Ed: Well, that's, I think, the dream that would emerge from an understanding of the building blocks. If you really know how the parts work together and you really can understand the whole system in terms of those parts, right, then it becomes possible to repair and maybe someday to augment and to study these higher order processes in terms of the parts, absolutely.

Yicheng: And I think just to follow on your previous point, I think there's, for example, if we know about quantum mechanics, we can sort of calculate how molecules interact. However, I think due to computational reasons, we wouldn't try to simulate using quantum mechanics how macromolecules will even like cells on a bigger scale, because even though we know that the physics apply to them and it must apply to them, but it's not like it doesn't give us useful information, but rather we cannot use this to perturb or simulate those systems. So do you think, even if you know the ground truth, some of the ground truth must be so grounded that it is not the most convenient way for us to know about the real world?

Ed: Well, even in physics and chemistry, right, you know when you reach the ground truth, right? So in physics, if I'm building a microchip which will operate at room temperature here on my laptop over there, we know that certain things will have to take into account, but do we care about the subatomic particles that are inside a neutron? Do we care about the gluons and all these other... You know, we don't, right, because those things don't apply at the energy scales of a laptop computer, right? They're irrelevant. So similarly, people have done calculations and measurements of things like quantum effects in biology, and for many processes, maybe you can ignore them completely, but not all of them, right? I mean, just in the last year, a couple groups have shown that proteins used by thousands and thousands of groups might be able to have reasonably long-lived quantum-coherent states, right? And so anybody who ignored quantum mechanics in biology and who just said that it can never apply to anything at a body temperature cell, you know, they're starting to look like, you know, that might be incorrect, right? Maybe there are, you know, effects that could last biologically relevant timescales, and, you know, making assumptions that you could ignore things should be ideally backed up by data and not just by guesswork.

Yicheng: Yeah, that's really interesting.

Ed: Just because it's hard to compute doesn't mean it's not important, right? Pretending the problem is simple doesn't make it simple.

Yicheng: Yes, I agree with you here. I think what I am trying to suggest is that there could be a very fundamental ground truth, and then from this ground truth, we can sort of, there's going to be other shallower ground truths that adhere to the more fundamental ground truth, but maybe those are because they are slightly on a bigger scale. It could be easier for us to use those ground truths, but I agree with you completely that just because it's hard to use doesn't mean that it's not important.

Ed: And I would hesitate to call those higher-level ground truths ground truth right away, right? I mean, it can take many decades to know whether they really are. I think it's a little bit dangerous if we call all truths ground truth, right?

Yicheng: Yeah, I think so. Just to change gears a little bit, I remember asking you last term in your talk whether the future bottleneck in biology would be data acquisition, processing, or storage. You replied that it doesn't hurt to have more tools. Would it be fair to say that, from your perspective, the real limiting factor is what parts of reality we actually have access to and are able to measure?

Ed: I mean, it's useful to say how biology is different from physics and chemistry, right? Physics, you have a couple particles. Chemistry, you have several dozen atoms that you have to think about, right? You know, biology, we have countless biomolecules that interact in complex ways. The alphabet from which you build the number of building blocks in biology is much larger than in physics and chemistry. So that's kind of, I think, the issue. If you want to see and control a biological system in terms of its building blocks, you know, system scale down to building block scale, then you need to have sophisticated tools. That's the guiding principle that's governed my graduate school career and my running the lab here for the last 20 years, is to build tools that get you, ideally, ground truth at system scale. That way you don't have to make any guesses of which middle grounds can be reliably rested upon.

Yicheng: Right. So after we may know the system a little bit better from observing it, we're measuring it. If we want to move towards, for example, therapeutics, do you think the bottleneck is sort of the delivery problem or making this system stable enough? What do you think the real bottleneck is for therapeutics?

Ed: Well, it's hard to see beyond a fundamental scientific wall. Sometimes people get lucky, right? People discovered penicillin and many scourges of humankind are almost unheard of because the bubonic plague, for example, right? Wiped out what, a third of Europe or something? I forget the exact number. Now nobody gets the bubonic plague, almost nobody. So sometimes you get lucky and sometimes you don't, right? You know, almost all non-infectious diseases, diabetes and most cancers and brain conditions and so forth, you know, are really challenging, right? None of them, you could argue, can be fully cured. And the treatments have a lot of partiality in their effect and a lot of side effects. So what will we find when we actually hit the ground truth? And of course, it's going to be so complicated that we are going to need AI and computer models to make sense of it. Maybe we find a node of the network that, if you tickle it just right, solves everything. Unlikely, but it could happen. Or maybe you find out that you need to have a multidimensional attack on the problem and sculpt the network of the brain in order to heal it. That could happen, too. But it's very hard to see beyond the end of a science. And if you look at, let's say, science fiction before quantum mechanics, and I searched really hard, zero people, zero, had the idea of the microchip. Lots of people thought about computers, right? You could have a electromechanical switch. You can imagine mechanical gears. Lots of people thought about computing devices in the 1800s. Zero people, as far as I can tell, could think of a microchip because a small piece of glassy-ish material that's made out of sand, right? I mean, how could that process a billion things a second? It just doesn't make any sense. So it's very hard, I think, and that's physics, right, with only a couple particles. Imagine now the

brain, which has so many more building blocks. I think if we understand the ground truth of the brain, it's very hard to predict what will happen next in the sense of therapeutics. Because we might find out that it's way more complicated than we thought. Or sometimes we get lucky, we might find it's much simpler than we thought. But it's very hard to know.

Yicheng: Yeah, that's quite deep.

Ed: Yeah, I started writing a couple articles for the public to try to talk about this more. Konrad Kording and I wrote an article for The Transmitter, the Simons Foundation's newsletter, about bottom-up neuroscience. Denisa Lepădatu and I have been writing a blog that we playfully call Universe Conquering, which is the idea that most human activities focus on trying to conquer a field or the world or whatever. But science is about the universe itself, right? If cancer or Alzheimer's disease is causing suffering, that's an invisible enemy that we have to fight. So we started to write these articles about the ground truth and why you would get this bottom-up explosion of creativity and problem-solving if it was achieved. The latest article we just posted on arXiv, again Konrad Kording and I, along with many wonderful colleagues, is this idea of what we're calling the compiler. You can make maps of molecules thanks to expansion microscopy and other technologies that build upon it. And again, if you really take the idea of ground truth seriously, molecules and how they interact should give you enough information to predict what a biological system does. But how do you get there, right? And so one thought is, let's build a machine learning engine that will learn to predict the physiology from those ground truth maps. And then making ground truth maps should scale up very, very well, right? That essay has this little plot showing how fast expansion microscopy data acquisition is going to go in the coming years. It might be one of the fastest-growing techniques of all time in the sense of just how fast you can go, right? It's only limited by things like camera speed. Whereas physiological imaging, recording live neurons in an awake-behaving mouse, that's much harder to scale up, right? There's very limited real estate. Whereas if you are doing molecular mapping, you could take a biopsy from a person, chop it into many pieces, and then have robots process them all in parallel, right? Anyway, this essay tries to talk about how the compiler, as we call it, would learn to predict physiology, including disease physiology, from these molecular maps. And that, I think, is a very powerful addition. So anyway, this idea of ground truth is starting to take off, I think. The reason, I think, that it wasn't popular before is that the tools wouldn't give it to you. But if, indeed, expansion lets you map down to nearly atomic scale, if, indeed, cameras that are being designed for self-driving cars and drones are going to allow imaging at faster speeds than the previous generation of cameras, which were governed by the cell phone industry and other market drivers, that could really raise all boats and help everybody solve problems in this regard.

Yicheng: I just want to talk about this compiler idea that you mentioned. So I think you can either feed it a list of ground truth you think they are, or, at the same time, you can maybe just show it molecular data. I'm assuming you mean like the latter, where you want the machine learning model to actually learn what are ground truths themselves, rather than telling them a few prerequisites. Is that correct?

Ed: Well, I mean, it's maybe more straightforward than that. We just want to make a map of the molecules in a brain, and then maybe we'll have physiological recordings from that very same brain, and so can we learn to predict the physiology from the molecular recipe? Over time, as modeling and computation and efforts like large language models for protein

prediction and AlphaFold, new generations, and so forth, as the molecular underpinnings of machine learning get better and better, maybe we won't even need the physiological measurements to calibrate. Maybe you could do a pure, bottom-up, unsupervised model, and maybe someday that's how we'll do it. But I think the fastest way to get going is not to ignore the physiology, right? Because why ignore that, right? It may not be the most fundamental ground truth, but it is a truth, and you can use it to guide the interpretation of what those building blocks really mean and what they do.

Yicheng: Yeah, that's a really cool idea. I'm just sort of trying to absorb that a little bit. So I think what my early question about delivery and homeostasis, I think what the idea I had in mind was an area of research that I'm quite interested in, is that haploinsufficiency disorders, which means that instead of a toggle switch on and off, we need like a tuner to sort of make it in the right, we need protein in the right degree amount. And if we have too much of it, we get disease. If we have too little of it, we get disease. And I think a lot of genetic delivery or genetic editing, it works in mice, but to treat humans, there will be, at least in the present, there's ethics issues. So if we use like other methods, for example, cell-type-based targeting of RNA or something, I think that is a small bottleneck we have on treating those diseases. Another small bottleneck we have is how do we maintain those cells in sort of this range? And then how do we make sure that treating one cell type doesn't impact the whole network? I think that's sort of what I had in mind, a more concrete example. And I think, as you said previously, a lot of those big breakthroughs require some degree of luck and a lot of trial and error. But I'm thinking, is there a way that we can speed up this process or at least have more constructive fallbacks?

Ed: Well, I mean, in terms of gene regulation, I'm not an expert in the areas that you just talked about. But as an outsider, my hypothesis is that gene regulation and delivery are fields that are being tackled with many existing disciplines. A lot of great chemical engineers are working on ways to refine delivery and to improve it. Several of my alumni, for example, have co-founded companies to pursue delivery into the brain and all sorts of stuff like that. So I feel like there's a lot of these problems being tackled through prior ground truth fields like chemistry. It's more of a chemistry problem, perhaps, than a biology problem. As far as homeostasis and so forth, if one is lucky and one can truly regard it as a homeostatic problem, that's amazing. Of course, there's lots of other things going on. So you mentioned haploinsufficiency. One could imagine that over developmental time, that there could be changes that accumulate from the change in gene dosage. And just correct to the gene dosage in somebody who is, let's say, 32 years old, might not compensate for all the changes that happened before. So there still are ground truth questions that could arise, unless you have proven it not to be the case.

Yicheng: And I think just following on from what you said, a lot of modern neuroscience breakthroughs requires people from engineering, chemistry, physics to come together to work on these bigger problems. And so my question is that with so many fascinating questions in neuroscience today, like delivery, like therapeutics, like more fundamental neuroscience, etc. How should young researchers, maybe at an undergrad or maybe at a grad level, choose a problem for themselves to devote into?

Ed: So I'm really bad at giving general advice, because everybody's different. Maybe one person takes too little risk, and they should be more adventurous. But if somebody's being too careless, even foolhardy, it would be bad for me to tell them to take more risk, right? So

I think the best thing I can do, and actually we have a lab blog, engineeringx.substack.com, and we're going to post in the next day or two an essay that we're calling anti-advice. The idea being that nobody can really give you advice, right? And so you need to learn how to create a space of advice for yourself. So I'll have a 4,000-word essay on this, posting any day now. And the problem with asking somebody for advice is that they have different lived experiences, and they maybe grew up in a different time. And the field was different back then. And what was a brilliant, earth-shaking, positive impact in 2005 might be kind of boring and humdrum, and everybody's doing it in 2025, right? So I think the best advice I can give is better advice, which is, you know, it's great to get advice, but then you can build what I like to call a tiling tree, a way of organizing all of your thoughts to be systematic. And then maybe you end up discovering new advice that nobody told you by pursuing this exercise of the tiling tree. And we wrote two articles about this strategy, which are also on that blog I mentioned. And it's a very powerful strategy, and many of our group's technologies and discoveries emerged because of this strategy. And we show examples on the blog. But the point I'm trying to make is you can take, like, a couple samples of advice you get, build a tree structure to organize them, and then fill in all the gaps yourself, and then discover the optimal advice for you. But again, since everybody's different, it's really, really hard for me, anyway, to give general advice. The short version of this is, whenever you hear advice, consider and visualize what would happen if you did the opposite of it. And this can become a playful exercise that you can do every day. And I've heard from many people who have told this over the years that, you know, one person, her professors told her not to start a company. She heard me say about this anti-advice recommendation. She thought, what if I do the opposite of the professors? Now she's working on her second startup company. The first one's doing well, and the second one is a really cool neurotech innovation. So, yeah, I think doing the opposite of advice is a very powerful strategy. I mean, don't do it indiscriminately, right? The opposite of advice could be even worse advice, right? But it's worth visualizing what would happen if you did the opposite.

Yicheng: Yeah, I've read about how you used, like, tiling trees to think about problems, but I don't think I've heard this version of using it to sort of consolidate advice.

Ed: Yeah, we were going to post this any day now. We just really came up with this framing, because I was trying to figure out, like, for 20 years, I would go to a university, and then the students at the lunch or whatever would be like, do you have any advice for us? I would never know what to say. So after many hundreds of lectures and many hundreds of corresponding interactions, I was like, I need to find a way to formalize this 20 years of meta-advice, because it's actually kind of useful. But I'd never written it down before. So yeah, with two undergrads, Claire Wang and Nina Khera, who write this lab blog with me, and others are starting to write articles too, but the two of them have been working with me on the blog since the beginning. Our hope is that we could spread these strategies to more people to help them.

Yicheng: Yeah, absolutely. And I think I've read another blog that is somewhat related to you. And so the bigger picture is that because connectomics give us such rich structural information, and it is something that your group has been working on for quite a little bit recently. I think a blog by one of your previous students, Adam Marblestone, sort of beckoned for a comparative brain connectome initiative. And my question is that when you think about connectomics' explanatory role, what do you see it giving us that we can't get any other way? And then to get it from wiring to things like emotions or even bigger things

like consciousness, what layers of information do you think we need to measure or at least perturb alongside connectomics?

Ed: Yep. Yep. Well, connectome, of course, that word emerged from electron microscopy studies, right? *C. elegans* was the first connectome. The fruit fly was done a couple of years ago. And so electron microscopy will give you the shapes of the cells and how they contact each other. But if you're going to know whether a synapse, a neural connection, is strong or weak or whether it's fast or slow, you've got to know the molecules, right? And so we've been, I think our lab is the only group that's actually used this term. I just looked up in Google Scholar a couple of days ago trying to see how far it's spread. But this idea of the MAC, the molecularly annotated connectome, or MAC for short. And yeah, so if you had a map of the brain and the molecules along the wires of the map, then you could potentially, maybe using the compiler concept that we talked about earlier, make simulations of the brain, right? And also use that map to answer all sorts of interesting hypotheses about the brain. So yeah, I'm very excited about the MAC, the molecularly annotated connectome. And I think tools like expansion and LICONN, which builds on expansion, and E11's barcoding, which also builds on expansion and so forth. Many people are using expansion as a platform for their own innovation, right? Which is really cool. I would love to democratize not just science, but democratize science innovation as well, right? That's sort of a side topic. But I do think that the molecularly annotated connectome might be within grasp maybe for small brains now, and if we could scale it up to larger brains too.

Yicheng: And because of maybe if we just have a static snapshot of a brain may not be enough, maybe we need dynamic, at least like a video recording of it. Do you think because the information gets so complex, this is when you think AI is going to come in and help us?

Ed: Well, I think, I mean, AI is out there already, right? It's like air or water. Everybody's using it. So yeah, I think we need that. But for the brain, just the sheer amount of data makes that important. As far as the video versus the static snapshot, that's really where the compiler concept is key. From static snapshots, if you learn what those molecules do, you could do, in the fashion of generative AI, models of the dynamics that emerge, right? The same way that you can take a structure and, using generative AI, simulate derivatives or conformational changes and so forth. So I think inferring dynamics from structure, of course, is a very exciting arena of all machine learning right now, you could argue. But I think for the brain, it's fairly well posed, although in the short term and to medium term, I do think having physiological observations will be very, very helpful. And so we have been working, for example, on scaling up voltage imaging. Can we image every cell in a brain, the voltage of every neuron? That's kind of a goal that we're pursuing right now at high speed.

Yicheng: Right. And I think, I can't remember whether I read it on another blog or it was an EngineeringX blog, but you mentioned something about how, if hypothetically, we get to know our consciousness, our emotions a little bit better, it will sort of reframe how we think about it. Could you sort of over here expand on that a little bit for us?

Ed: Yeah, yeah. I think there's a couple of reasons why consciousness is interesting. First of all, I mean, society is constructed around our assumptions about consciousness. If a person commits a crime, then they will be punished, right? If a vacuum cleaner robot injures somebody, nobody thinks about putting it in jail or, you know, maybe you throw it in the

trash. But, you know, so both from an accountability standpoint, but also a suffering standpoint, right? If you see somebody who is experiencing pain, I think we would want to help them, right? But if a vacuum cleaning robot is missing a part, most people don't run over to it and try to help. So whether we like it or not, consciousness is kind of, you could argue, in our judgments about consciousness, the basis of almost all of our social constructs. Empathy, justice, accountability, you name it, right? So for something so important, it's kind of worrisome that we can't measure it. So I do think that there are these sort of large humanistic questions that consciousness is really the core of. But also there's lots of practical things too that are extremely urgent. There was an article a couple days ago, I forget where it was published, but pointing out that many people who are in like a coma, many of them have consciousness. And yet they're sitting there, people think they're unconscious for decades sometimes. You know, disorders of consciousness, right? Those are incredibly incapacitating. And yet you could argue that we're just beginning to try to think about those things from a therapeutic standpoint, right? You know, there was this wonderful paper, I think from Nicholas Schiff and colleagues, where they did deep brain stimulation of people in these coma-like states and could restore some degree of responsiveness, right? By stimulating the thalamus. I mean, that's really, really intriguing, not just from a philosophical standpoint, but also from a clinical standpoint. And of course, there's the issue of our own happiness as well, right? What is the meaning of life? And maybe this is maybe the farthest thing that I think about, but I find it fun to talk about. And I think a lot of us in our lives aspire to happiness or meaning, for some people might aspire to what they call enlightenment. But what is that really? And if we understood why our brains want what it wants, could we come up with something even better for that, right? Like, oh, here's why I want these things. But you know what? Now that I've dug a little level deeper and I know really why I do what I do, maybe I want something else. Maybe the everyday analogy is, if I'm a kid, I may not know much about nutrition, so I'll go and eat a lot of candy and cake. But as I learn more about why I eat what I eat and do what I do, almost zero adults are just only eating candy and cake, right? Everybody has brought in their wants thanks to our understanding of reality, right? And so this might apply, I think, beyond things like food and nutrition. Maybe we find that there are new kinds of things that we might aspire to as a species.

Yicheng: Yeah, I think that's a really cool idea. And it's really cool how neuroscience sort of looping back towards ourselves more and more now. I think the other day I read something that relates the free energy principle to meditation. And it sort of connects how we see ourselves to Buddhism and to neuroscience. So I think it is really cool how things are sort of coming together to a loop. And just building on that a little bit, I think I want to talk about ethics just a little bit. I think you have stressed more than once on different occasions that ethics is at the core of balancing innovation with reality. And the possibility of an international workshop of people coming together to talk about this. So my question here is that if this workshop were to take place, could you share with us a little bit what kind of principles of self-governance would you propose? And what would be the thinkings behind to back up those?

Ed: Yeah, I've been talking a lot about this idea of an Asilomar for the brain. In the 1970s, Paul Berg, the Nobel Laureate at Stanford, proposed this workshop at Asilomar at the dawn of the genetic age. What should we do about cloning genes and putting genes in organisms and so forth? And actually a couple of weeks ago, this group called BrainMind

did hold, at Asilomar in California, such a workshop. Unfortunately, I came down with a bit of a bug, a viral bug, and couldn't go. But I think that there are such efforts that are beginning and I think that there are a lot of things to think about. Right. So first of all, there are lots of I didn't attend, so I'm not speaking on behalf of that group, but this is just my own thinking and some chats that we had in our group here. Obviously, they're all the medical ethics concerns that go with any medical field. We have institutional review boards. We have government regulatory bodies. And you could argue that because of the efforts of people like Paul Berg and others and the medical establishment as a whole, there is a lot of activity already in that space. And there's always more transparency about company affiliations, all sorts of ethical levels where the medical establishment pursues excellence. But there are domains that don't necessarily cleanly fit into that bucket. For the brain in particular, I think there are a couple of areas that are not well addressed by any self-governing or regulatory body. One, of course, is augmentation. If I augment my brain, what are the ethics of that or if somebody else wants to augment their brain? Another area is privacy. So most legal systems and so forth, our thoughts are our own. What you write can be seen by others, but if it's your thoughts, it's your own. But what if you could scan somebody's brain? I think privacy is a huge issue. So there's a bunch of topics that I think relate to these brain-specific functions. Maybe amongst the most intriguing is identity. If I stimulate my brain, does it change me? Am I now a different person? Are my actions different or am I different? And how does this relate to questions of ethics, accountability, and all sorts of other things? There are a lot of these questions that have arisen, and my hope is that these dialogues are continuing and that there are maybe both self-governance principles that emerge, but also when appropriate, regulatory governance approaches as well. Because some of these things could, if they go wrong, it could kill the whole field or incapacitate part of it or hurt individuals or all of the above. It's really something we should take seriously.

Yicheng: Which parties do you think we should pay more attention to? Is it the scientists? Is it maybe the philosophers or maybe it's the government officials or maybe the economists? How do you think we should balance those different people's different perspectives and opinions?

Ed: Well, I think a lot of scientists and philosophers and government officials and so forth have already begun to dive into neuroethics, although we could all do more. Maybe the most intriguing arena of neuroethics, though, where I think a dedicated attempt to build consortia and stimulate conversation and arrive at self-governance principles would be in entrepreneurship. There are a lot of companies popping up that want to do brain stimulation or that are doing brain stimulation. I've seen over the years companies proposing everything from using brain scans to attempt lie detection to companies that are offering brain stimulators that might augment your brain performance. So it's out there. These things are beginning, whether we like it or not. And so I think that's kind of the area that I hope will get the most self-reflection. In Silicon Valley, there's this saying that drives many startups, this idea of move fast and break things. And that might be OK if you're making an app for a phone, but you might not want that if it's your brain.

Yicheng: I think I have two slightly personal questions. The first is that you've mentioned entrepreneurship, and I think you have a couple of companies that you've spun out yourself. So I was wondering, when do you think you decide that this technology should become a company?

Ed: I think every problem has a natural home. If you have a problem that needs scalable capital, you need lots of employees, you need to deal with regulation and other topics that have structure, a company makes sense. If you need open collaboration, if you need serendipity, if you need timescales that are not so short, then maybe it ends up in academia. Now, something interesting happened. A couple of my students started writing thesis chapters about how they would improve the scientific enterprise if they could, and some of them started making it happen. So a couple of my former group members invented this idea of the FRO, Focused Research Organization, and then three of them started organizations, and now more are in the pipeline to go after it. In fact, E11, the first focused research organization you could consider to be a lab spinout run by my former student, Andrew Payne, is trying to scale up expansion brain mapping out in California. Sam Rodriques started FutureHouse, an AI scientist FRO. So these can tackle problems that are too big scale for an academic group, but not quite ready to make money, and so you wouldn't want to take dilutive capital. So, yeah, I think if we had this perspective that every problem has its natural home, then sometimes you find out that that home doesn't exist, but then you can invent it. And that's something that I think many of our alumni have found very empowering in their own careers. Some of them have even quit tenure-track faculty jobs to go and start one of these new organizations because the problem that they really wanted to solve, that they thought would benefit humanity, didn't fit into a clean existing...

Yicheng: Yeah, this is really sound advice. And I think actually building up quite nicely from that, you being a neuroscientist, an entrepreneur leading your own research group, and then also like a father, husband, you have so many different roles. How do you sort of balance your time, your energy, do all of these?

Ed: Well, I think in science, a lot of people like to talk about work-life balance, but I think one can also, and it's not always an easy road, but if one structures one's scientific life, it offers a lot of freedom. So, for example, I've always had this inclination that if other people could work on something, I didn't want to work on it. Now, that does have stressful moments, like I struggled to get a faculty job. I struggled to get funding for many projects like expansion microscopy. But once you get going, you have very little competition. And so I think CRISPR started coming out around the same time that we were working on expansion. And I think we all sat down at one point and were like, OK, we're not going to work on CRISPR. Everybody's working on that. We're going to double down on microscopy. And we ended up kind of reinventing a lot of our group as a microscopy group. You know, expansion, lots of people in the group worked on that, microscopes for voltage imaging. There are lots of new directions that we can go into. Anyway, back to your question. By choosing these kinds of paths, it gives a lot of flexibility. And then I would also seek alignment in the early days. In the early days, my co-parent and I would take our kids to conferences. Now they're old enough, they're teenagers, that isn't interesting to them anymore. But they've been to many Society for Neuroscience meetings, or at least the zoo and the museums nearby over the years. And so, yeah, is there alignment where I can try to get everything kind of in the same direction? And it's not always easy, you know. In addition to struggling to get a faculty job and struggling to get funding, getting on-campus daycare used to be really hard. I think it took two or three years for my son to get into MIT's daycare. So the first couple of years were a bit hit or miss in terms of finding help. We were very lucky that two of our parents were able to help out a bit. Anyways, yeah, so seeking alignment, I think, is an interesting design principle to explore. But maybe there are strategies, like being a bit contrarian, that

give a bit more freedom. Now, of course, this goes back to advice not being always helpful. You know, the current scientific establishment, where at least in the US, right, big funding cuts, large slashes in scientific support and so forth. You know, maybe things get even harder for the contrarian projects. And so it's a little bit unclear what the best thing to do is right now. But these perspectives, of course, could be filtered through the advice tree methods like we talked about earlier.

Yicheng: Yeah, absolutely. I just want to move back to a slightly big and more obscure question. We've mentioned AI a couple of times. So with AI and neurotechnology moving so quickly, do you think the future form of intelligence will be something that's still biological, or maybe something in between, like a hybrid or something fully digital, or perhaps another form? What do you think that means for us to be, you know, humans to be humans?

Ed: Yeah, that's a good question. Again, it's a bit hard to see beyond the end of a science, right? So I would argue that when neuroscience gets to its ground truth, we will know so much about how the brain works, that bottom-up design will become possible. Maybe people learn how to grow brains that, you know, could augment our brains. Maybe people learn how to do certain things in silicon chips that seems daunting to do in silicon, but once you know the map of the brain, aren't so hard to do. So it's really hard to predict, I think. But I do think once we see the answer, which in my book would be understanding the structure and function of the brain in terms of the fundamental building blocks and interactions, it could be that there'll be, like Silicon Valley after the microchip, an explosion of productivity in all directions. It's a little bit like before the microchip, will computers be large, wearable, implanted? Well, hard to say, right? But once the microchip came along, the answer was all of the above. There are supercomputers the size of a city block, and there's microchips that are implanted in people for cardiac pacemakers or deep brain stimulators. And so it's quite possible that intelligence might take many forms once we understand it. I think what we don't have now is a good understanding. And that's why, for example, large language models, despite billions and billions and billions being poured into them, and despite the fact that they do pretty well at language, but they still struggle with decisions and thinking and so forth. And every day there are examples of a large language model making an error that an eight-year-old child would know not to make. So I think it's very clear that large language models do make huge contributions to language, and people try to build upon it and iterate to get towards thought and decisions. And progress is being made, but is it the ultimate architecture for intelligence? My bet is that the brain's got some interesting tricks up its sleeve.

Yicheng: Which layer of science do you think on that? Obviously this is just us hypothesizing, but which layer of science do you think is going to be key to tackle or identify that trick?

Ed: Well, that's kind of what I want to do by decomposing entire brain circuits into the building blocks and simulating them. That's kind of where I'm focusing on. Anything else, of course, people could get lucky, but it requires luck. You don't know you're done until you're done. And so if somebody stumbles across the secret by looking at brain region 423, that's amazing, right? But that hasn't happened yet.

Yicheng: Yeah, I think a lot of people are alive with the rapid expansion of AI and are very scared that AI is going to overtake us whatsoever. So I think, obviously, neuroscience and AI are two disciplines that I can feed into each other. Do you think that progression in neuroscience is going to help us design better AI systems and AI systems are going to in turn help us to progress neuroscience research? But do you think, for example, if one of the branch progresses a lot faster than the other, I don't think it's going to be too much a problem if neuroscience progresses faster than AI. But do you think if AI progresses to a way that it is still exponential whereas neuroscience remains more linear, this is going to in turn cause problems?

Ed: Well, two thoughts. One is maybe we should accelerate neuroscience. I've been talking to lots of people about whether we should start a new institute or maybe even a new university that could accelerate neuroscience, right? We have these new tools. Many groups adopt one or more of them. Are any groups integrating all of them into a new workflow? It's difficult to do that, right? But maybe there could be a new intellectual and organizational structure that could allow that to happen. The second thought is I really wonder if for ultimate AI you will need neuroscience. And again, lots of my friends in AI don't seem to think that way. They're very excited to train a model that then will train a new model and so forth. And maybe that'll work. But I think it's also worth pointing out that despite the billions and billions and billions and billions being poured in, there still are lots of examples on a daily basis where an AI will make a very elementary mistake and then people try to patch it up. But that doesn't seem to be how the brain does things. The brain seems to have some interesting, not yet described mechanisms for intelligence that I think are very tantalizing.

Yicheng: Maybe more money should be poured into neuroscience.

Ed: I think it's going to be, right? You know, Merge Labs has raised a quarter billion to build brain-machine interfaces. I've heard rumblings of other startups raising hundreds of millions of dollars to build commercial neurotech of different kinds. Yeah, I think the turning point is starting.

Yicheng: This sounds really exciting. Can we just revisit my initial question on ground truth a bit more? If the brain is optimized for survival rather than for the actual truth, and when I say truth here, I mean how the world actually works. If this is the case, what kind of experiments or tools do you think could let us test whether our perceptual basis is limited to the kind of truth that we can discover?

Ed: I mean, to one level, right? There are lots of optical illusions and other examples that are almost trivial of how reality differs from our perception, right? The waterfall effect, all sorts of countless optical illusions, right? Yeah, so I think it's not too surprising to say that the external world is not perfectly represented in our brains. As far as understanding, two thoughts. I think it's interesting to dissociate the feeling of understanding from actual understanding. A famous physicist once said, if you understand quantum mechanics, then you don't, right? Because it doesn't make any sense, right? If you feel like you understand quantum mechanics, that's an indication that you don't understand it. So I just wanted to point out that the feeling of understanding, and now some neuroscientists are studying this, by the way, the feeling of understanding and actual understanding can be different. I remember a fascinating study where they asked people how confident they were in their solving of puzzles, and the people who doubted themselves actually did better. I forget the

exact details, it was so long ago, and I probably misquoted the study a little bit. So anyway, back to your original question. I guess we just have to see what is there, right? A brain didn't evolve to be understood, but the history of science does suggest that if you try a ground truth play, like chemistry with its periodic-table/post-molecular-bond renaissance, or physics with its post quantum mechanics renaissance, if you do try to go for the ground truth, maybe patterns emerge that are understandable. There's no guarantee of that, right? Maybe physicists would have found that there's so many particles that it's just incomprehensible.

Yicheng: Yeah, I think this is really, really cool work, and this is going to make us change how we think ourselves as a whole, if we do get the next breakthrough. It's like, for example, if we're able to simulate from the bottom up approach, and we're able to sort of reconstruct us in a digital way, then, for example, memory, right now, time goes for us unidirectionally, but then if we're able to sort of simulate every single point, then sort of reverting back to a previous point no longer becomes fighting against all the entropy in the world. And I think that would fundamentally sort of change how we see ourselves, how we see the meaning of life or humanity or a lot of stuff. And I think this is why neuroscience is very powerful for myself.

Ed: Yeah, that's so exciting.

Yicheng: Yeah, like I know before, I think I had the idea that if we have the feeling of understanding, and then if we understand what this feeling is, where it comes from, whether it's from a molecular or a system circuit level, and then if we're able to sort of perturb it very accurately, and then, like even if, for example, you say something, I don't understand it at all, and then you perturb my brain to have this feeling again. Does that mean I understand what you said? Or is that like a pseudo-understanding of sorts? I think those are some very interesting questions that probably none of us have answers for right now.

Ed: Yeah, there's a whole subfield of consciousness studies where they talk about access consciousness versus phenomenal consciousness, and sometimes people will say, oh, I'm very confident, I saw the letters. And they're like, great, what were the letters? Oh, I don't know. And so there's the feeling of understanding that crops up. And one could argue, well, that is the conscious experience. If I felt that I saw the letters but I didn't actually see them, that is the feeling. You shouldn't call it not-consciousness. But anyway, this is where we get into philosophy. But it's a really fun topic.

Yicheng: Well, all subjective questions are philosophical in the very end. This has been a great discussion, Ed. Thank you very much for your time.

Ed: Well, it was great to chat with you today, and yeah, let's stay in touch.

Conclusion

Ivan: As we bring the episode to a close, we are reminded that Professor Boyden's work asks us to imagine neuroscience as a project of reconstruction: not only observing the brain, but learning how its parts fit together well enough to explain, repair, and perhaps one day redesign it. It is an ambitious vision — one that raises as many ethical and philosophical questions as scientific ones. But it is also a reminder that progress often begins with a tool, and with the willingness to look at an old mystery from a new scale.

Thank you for listening to CortexCast. Until next time — all the best, cheerio.