A Lecture Companion

”Ethics and the New Biology” by Michael Levin

Michael Levin

Recorded on April 26, 2025

About this document

This document is a companion to the recorded lecture *"Ethics and the New Biology" by Michael Levin*, recorded on April 26, 2025. You can watch the original lecture or listen in your favorite podcast feeds — all links are on the page here.

This document pairs each slide with the aligned spoken transcript from the lecture. At the top of each slide, there is a "Watch at" timestamp. Clicking it will take you directly to that point in the lecture on YouTube.

Lecture description

This is a ~30 minute talk to an Ethics class (undergraduates) about aspects of our work that have implications for ethics.

Follow my work

Twitter • Blog • The Levin Lab

Transcript note

This transcript is generated automatically. While we strive for accuracy, occasional errors may occur. Please refer to the video for exact wording.

Want one for your lecture?

Want something like this for your own talk? Reach out to Adi at adi@aipodcast.ing.

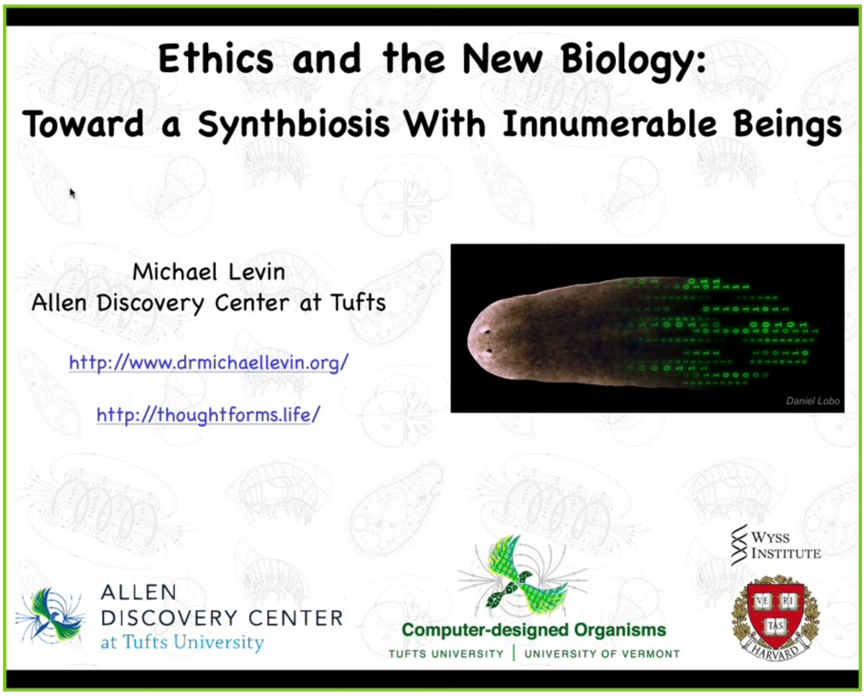

Thank you so much. I've not given a talk specifically on this topic, but I take it very seriously and I think about it all the time.

If you want to follow up on any of the science here, this is my official lab webpage that has all the papers, the datasets, and the software.

This is my personal blog with a crazier version of what I think all of these things mean.

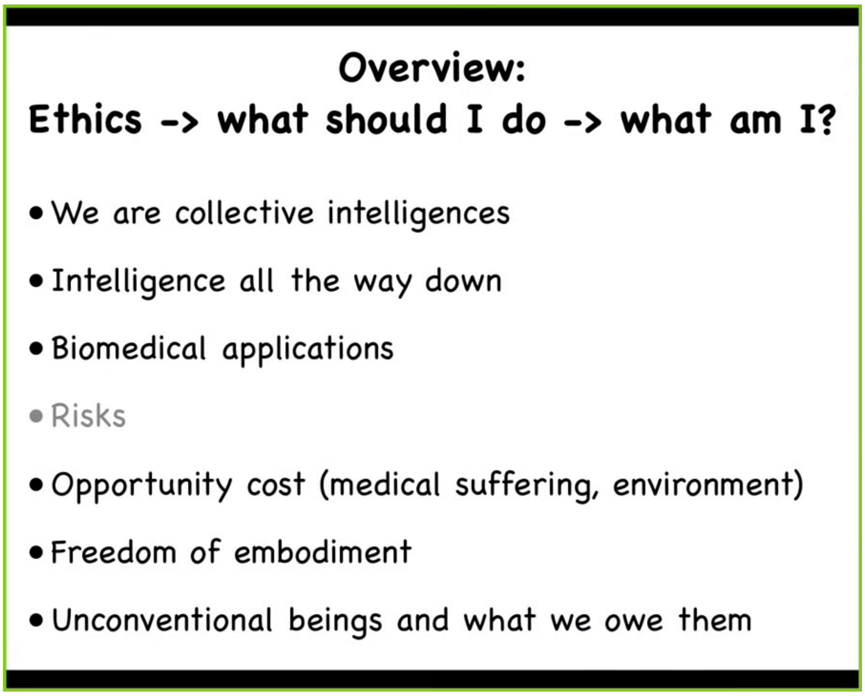

What I'd like to do today is first give you a quick outline of the journey I'm going to take you on. I'm going to try in the next 25 minutes to take you through a few different things. Fundamental to ethics is, I think, this question of what should I do next? I think fundamental to that is the question of what am I, really, and how do I relate to the outside world? This is a lot of what we do, even though we have a wet lab, a large wet lab where we do work on cancer and bioengineering and so on, but I think fundamentally this is the sort of thing that we actually work on.

I'm going to talk to you about this idea that we are all a collective intelligence. In fact, I'm going to try to show you, and again it's hard in 25 minutes, but I'll at least show you a few examples, why I think the world consists of intelligence all the way down. Not just brainy mammals and birds, but all the way down. The implications of that are that we can reach interesting new biomedicine if we understand how intelligence is implemented in our bodies. So that gives opportunities.

In any discussion of ethics, someone will be interested in the risks. I'm not going to do a lot on that in the main talk, but if you want to ask me questions about it later, it's the most common thing that comes up. I'd rather talk about some of the more unusual things.

One thing I want to be very clear on with respect to evaluating risk in the science is that we have to think about the opportunity cost. In other words, lots of people approach risk in scientific discovery with an unspoken assumption that everything is great now and you scientists better not screw it up. A lot of these discussions come from that place. Actually, you probably know things are not great. There is incredible biomedical suffering out there. I'm not a clinician, I'm not an MD, but I get emails, a dozen per day, from people in the most unbelievable extremes of pain and infirmity and disease. It's hard for young, healthy people to realize what we're actually talking about here.

I think we absolutely have to think about the risks, but we also have to think about the ethical imperative to do something to help the situation. I think it is moral cowardice to leave it alone and say, "not my problem." I'm going to sit back and try not to make things worse. I think that's not the right way to go. There are also issues with the environment.

Past the biomedical issues, I'm going to talk about freedom of embodiment, because after we learn to fix the issues of the body towards the normal standard complement of form and function, you realize that you can use the exact same tools to create any embodiment that you want. This is something that we can think about in any discussion of ethics.

Towards the end, I'm going to talk about the bigger picture of how that leads to a huge variety of unconventional beings with which we are going to share our world and what we owe to these beings. That brings us back full circle. You can see this is a cycle, even though I've listed it linearly.

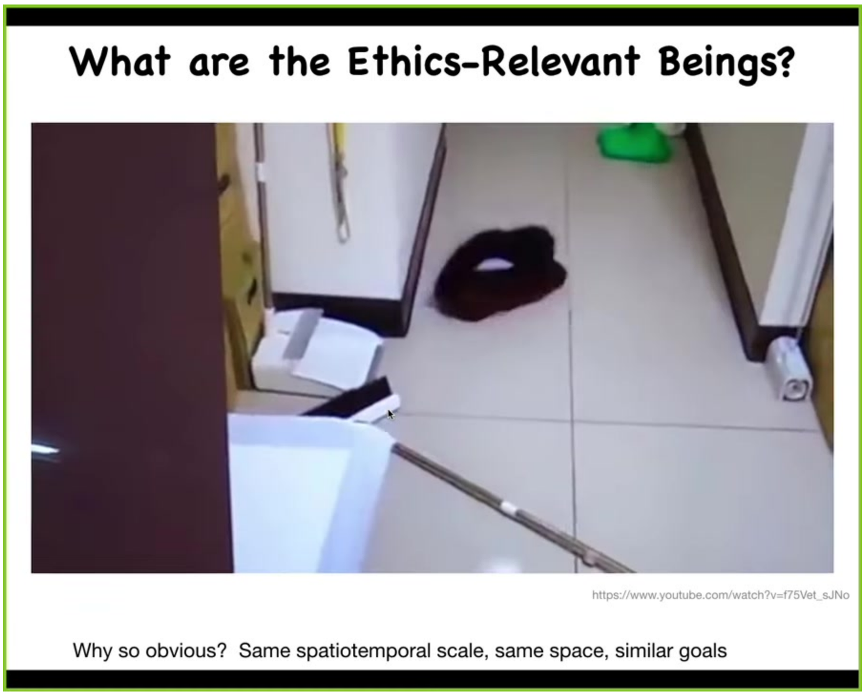

The first thing we need to think about is, what are all the beings that are relevant to ethics? When we see something like this, here's an Oscar-winning performance here, this little squirrel. She knows exactly what's going on. She's going to put the thing on the neck. Look, it's had a terrible accident. It's even going to look to see if the humans are watching to make sure that they see, are you guys looking? Look, I've had a terrible accident. When we see something like this, it's very obvious that this is something that has intelligence, that this is something that has preferences, that it has a theory of mind, and that we have to be kind to things like this. And it's obvious to us because it functions on the same scale of space and time. It operates in the same space as we do in three-dimensional space. It has similar goals to us in terms of survival and physiological integrity.

Slide 4 of 26 · Watch at [4:55]

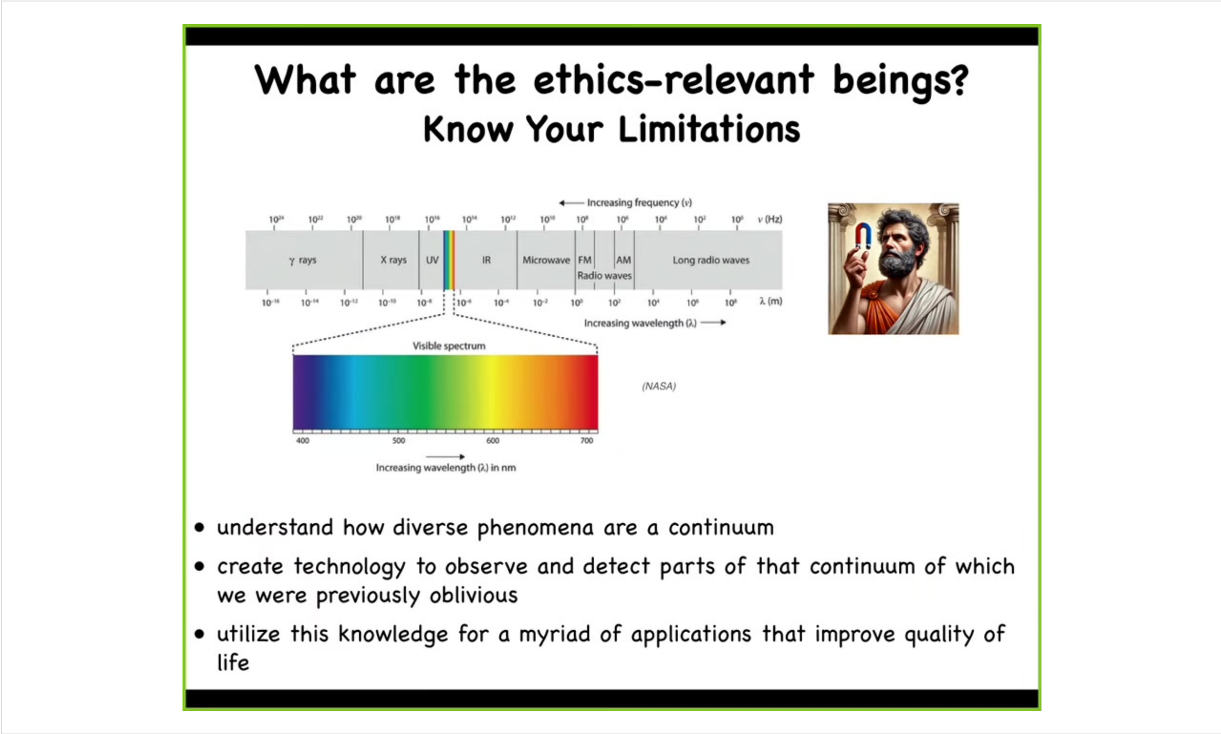

This is easy, but actually it gets much more complicated, and I'll just use an example of the electromagnetic spectrum. Back in the day, there were phenomena such as lightning and static electricity and magnets and light and all kinds of things. Until we had a proper theory of electromagnetism, we had no idea that all of these were the same thing. We thought they were disparate phenomena.

What a proper theory of electricity and electromagnetism did for us is it allowed us to say that all of these things are manifestations of the same underlying phenomenon. That phenomenon is continuous. In other words, it's a spectrum. It's not a set of discrete things. Even though we like to cut it up into regions just for convenience, it's a spectrum.

Most importantly, because of our evolutionary history, we are only sensitive to a tiny part of that spectrum. We had no idea these existed until that theory enabled technology that now, with the technological appliances and the theory behind them, enables us to recognize and make use of all kinds of other phenomena in our universe to which we were completely blind.

I'm going to make the claim that, as humans, we have a significant amount of mind blindness. We're reasonably good at recognizing things like this.

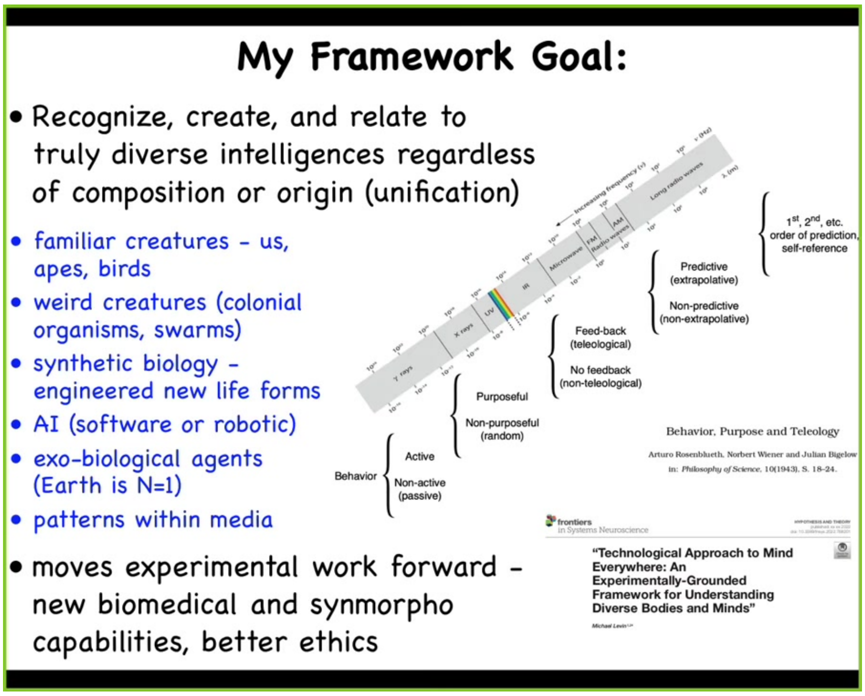

Lots of people study birds and primates and a whale or an octopus. But there are all kinds of other creatures. One of the things that I try to do in my group is to develop a framework that allows us to think about diverse intelligences. All kinds of weird colonial organisms and engineered new life forms that have never existed before. I'm going to show you some today. Artificial intelligence, whether purely software or robotic, someday aliens, and even exotic things that don't look like beings at all, such as patterns, active patterns with the media.

That spectrum, instead of the electromagnetic spectrum, is a spectrum of agency. But fundamentally what we have is a gradual continuum where passive matter, if there even is such a thing, slowly scales up its cognitive capacities to eventually become something like a metacognitive human or whatever is beyond that.

This is what we work on. If you're interested, this paper has a lot of the details. I think this is not just a philosophical task, because as we have shown, if you do this well, it leads to new biomedical and bioengineering advances. It drives new science. It's going to lead to better ethics.

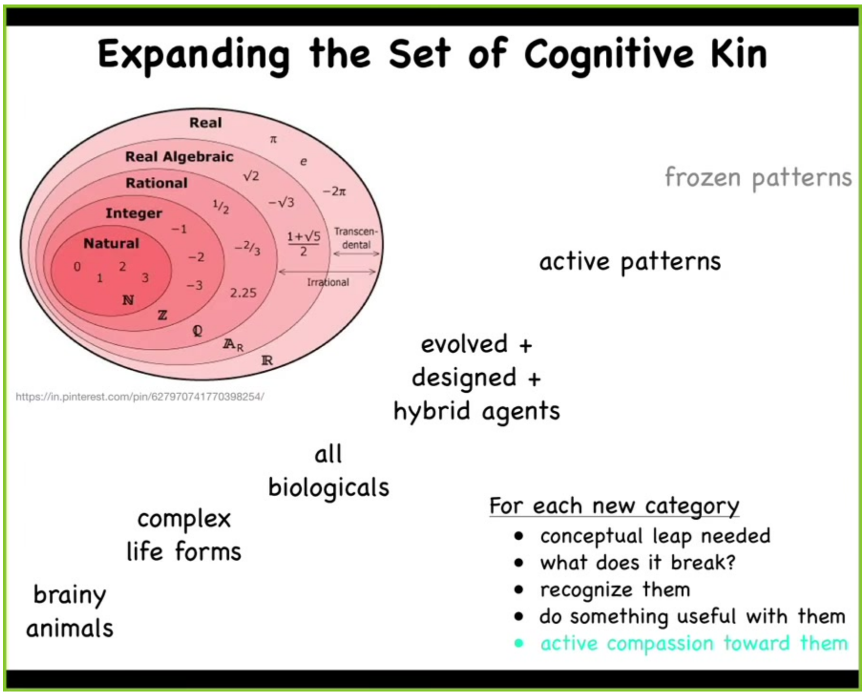

So another way to think about this is that if you know the history of numbers, first we had some counting numbers, and then somebody invented 0, and then we had some negative numbers, and then we had the rationals, and then somebody realized that wasn't enough, and then we needed the irrationals and so on. And each of these expansions of what we understand as number broke prior concepts as far as what defines the set of numbers. It required a conceptual leap to get there and to expand our understanding.

And it was actually quite stressful for everybody. People died over this discovery that the rational numbers weren't all there is. A lot of people were killed, people committed suicide. It was really fundamental.

And I think what we have here, I'm going to go through trying to gradually enlarge the set of things that I think are on the spectrum with us as being cognitive beings that are relevant to questions of ethics.

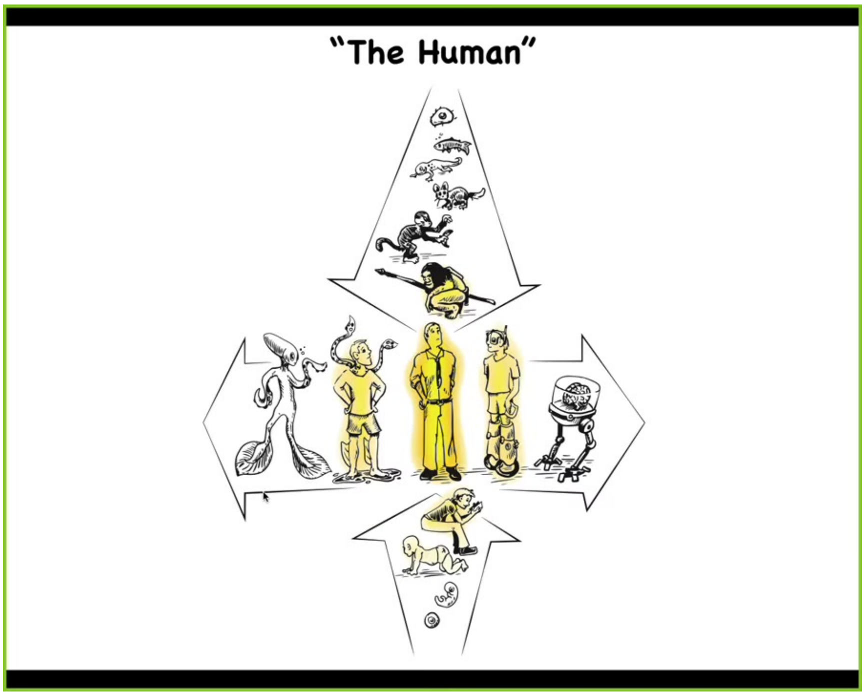

And the first thing is that in standard philosophy of mind or philosophy textbooks in general, what you'll see is they talk about the human. Humans do this and they do that, then they have the ability to imagine things and have moral worth. And so they have this beautiful agential glow about them that makes them important as moral actors. But we know that we stand at the center of this continuum, both on an evolutionary and developmental time scale, where we started life as a single cell and the origin before that. And so this too is a set of slow changes. Whatever philosophical properties this human has, you have to have some story of what happened in the previous stages. Like how did we get here? Is it gradual? Is it discrete? What's going on? As I'm going to point out in a minute, this is not enough.

There's another spectrum that we are at the center of, which is the set of biological and technological changes that slowly and gradually could move us beyond what this is. You have to ask yourself, where does this agential glow peter out? Do you not have to worry about these things anymore, or does it ever?

So here's how we all enter the world. We start life as a little blob of chemistry and physics. It's a little unfertilized oocyte. Then there's this incredibly magical process called embryogenesis where slowly and gradually, there's never a magic lightning flash that takes you from physics to the land of psychology. What happens is eventually you become one of these beings that are susceptible to the tools of behavioral science or maybe even psychoanalysis.

And so the story of our origin is really the story of scaling up of cognition and transformation. As scientists and philosophers, we need to understand this transformative process of how mind emerges from these simpler components.

Slide 8 of 26 · Watch at [10:46]

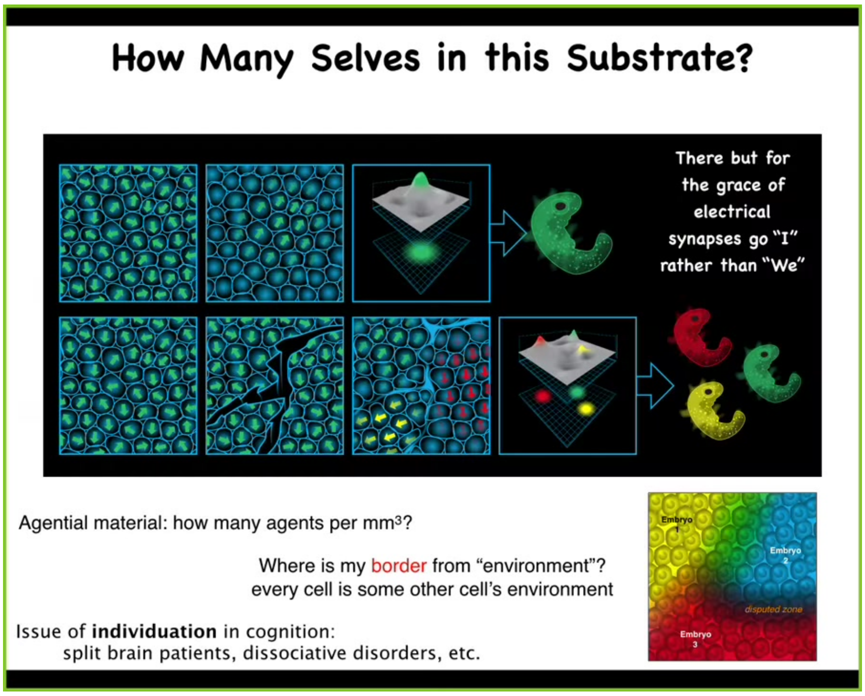

In fact, if you look at an early embryonic blastoderm, there might be a couple of hundred thousand cells at some particular stage, and you look at that and you say, well, there's one embryo, there's an embryo. And when you say it's an embryo, what are you counting? What is there one of? I mean, there's hundreds of thousands or millions of cells. What is there one of? Well, it turns out that the best thing you could say is that what there is one of is commitment to a particular story. In other words, alignment towards a specific journey in anatomical space. And the anatomical space is just the space of all possible configurations that a body might take.

And that all of these cells agree; they've bought into the same model of where they're going to go in anatomical space, and then they make this one embryo. That's what you're counting. You're counting commitment to a model.

If you make little scratches in this blastoderm, then each one of these islands is going to do the same because it's unaware of the presence of the others. Eventually it might heal up or it might not. Then you get twins and triplets and so on.

So what you see here is that the question of how many individuals are inside this excitable medium is not known in advance. It's not set by genetics. It could be anywhere from zero to probably half a dozen or more. There are potentially multiple cognitive beings that come out of this because of this process of differentiation.

Where do I end and the outside world begin? That's what's happening here. And that, of course, has lots of parallels to issues of individuation and dissociation in the cognitive sphere: split-brain patients whose left and right hemispheres don't communicate very well, dissociative identity disorders, multiple personalities on the same brain hardware, and so on.

So you're already starting to see that it's not at all as simple. Our origin and the underlying mechanisms are not what is commonly portrayed in philosophical discussions of humans and our ethical duties.

Slide 9 of 26 · Watch at 12:46

This is the sort of thing we're made from. This happens to be a free-living one called the lacrimaria, but this is a single cell. It has no brain, no nervous system, and is very competent in all of its local little goals. You can already see that it has some ability to manage its business at this lower scale.

Slide 10 of 26 · Watch at 13:09

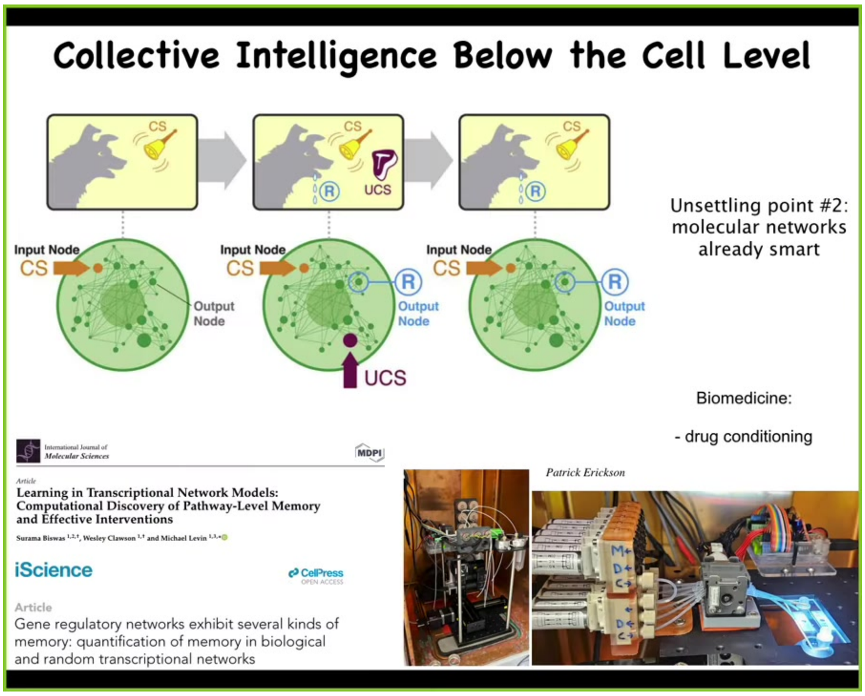

Even below that level, the molecular networks inside of that little creature, inside of all of us, are molecular networks that are capable of six different kinds of learning. Different kinds of memory, including Pavlovian conditioning.

The material that you're made of, all the way down to the molecular level, and probably below that, has certain behavioral properties that you would easily recognize as aspects of cognition. And we're taking advantage of that for all kinds of biomedical purposes in our lab to try to train these electrical networks, these biochemical networks for drug conditioning.

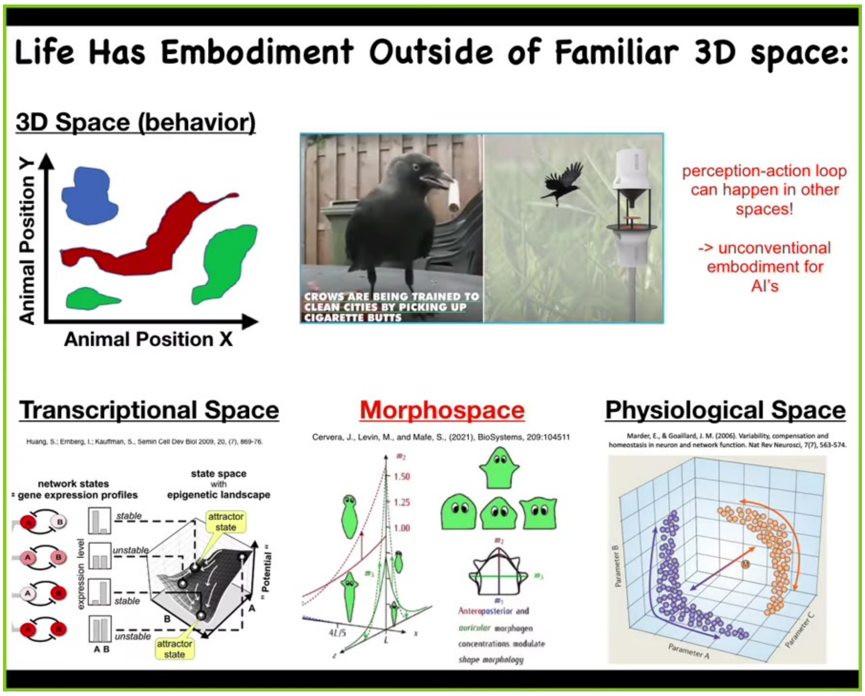

And so life exerts intelligence in lots of spaces, not just the three-dimensional space where it's easy for us to recognize things like these crows that have learned to exchange cigarette butts for treats. But it operates in many spaces. It operates in the space of gene expression, it operates in the space of physiological states, and it operates in anatomical morphospace, which is the space of shape configurations. And when I say life operates in these spaces, I mean that it navigates those spaces adaptively to solve problems, just like conventional animals that you're used to navigating a three-dimensional space.

Slide 12 of 26 · Watch at 14:27

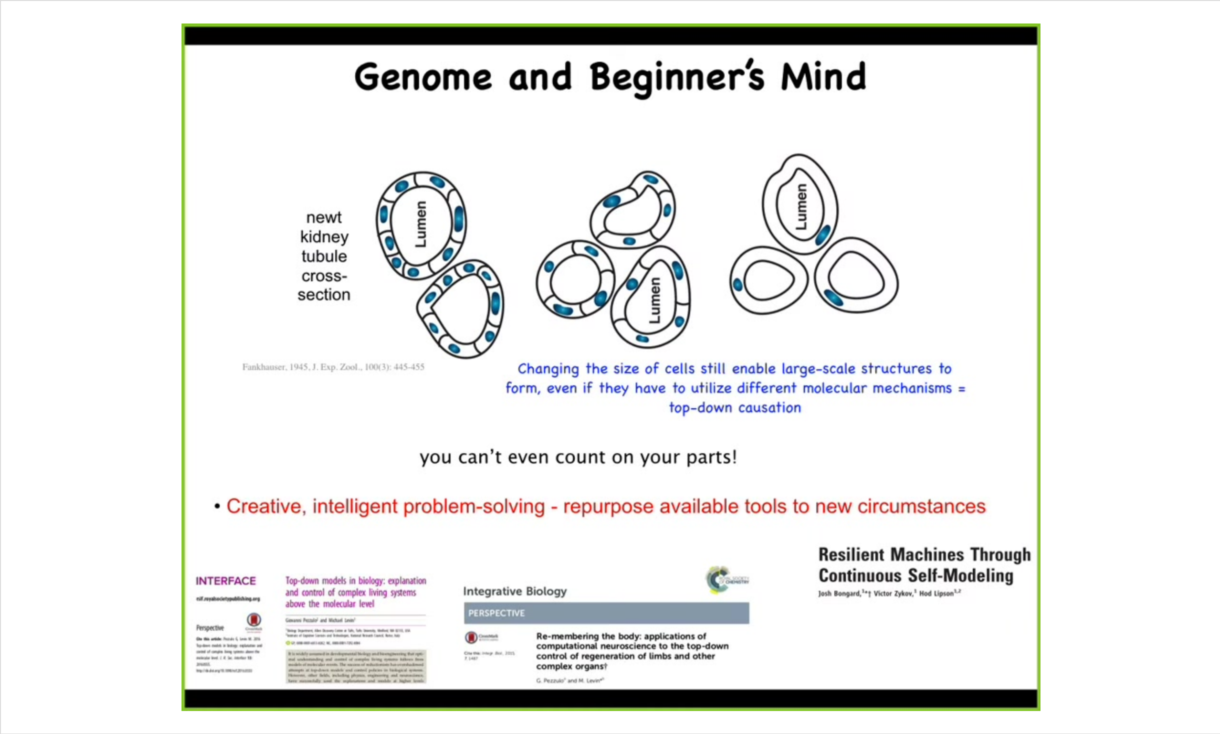

I'll give you just one simple example of what I mean by that. This is a cross-section through a kidney tubule in a newt. You can see it's got a lumen in the middle, and then you've got these 8 to 10 cells that work together. If you make newt embryos with multiple copies of their genome, the cells will get bigger to accommodate the additional genetic material. Oddly, the newt stays the same size. How could that be? While the cells are bigger, what they do is they adjust and fewer of them are now making the exact same structure. Until you make a newt with so many copies of their genome that the cells become gigantic, and instead of many cells working together to build this, one single cell will bend around itself and leave a hole in the middle. This is a different molecular mechanism.

So what you have here are a bunch of cells that are using different affordances, different genetically specified affordances, meaning different mechanisms that they have to get the job done in the face of incredibly weird circumstances. If you're a newt coming into this world, you can't count on the outer environment. You can't even count on your own parts. You don't even know how many copies of your genetic material you're going to have. You don't know what the size of your cells are going to be. You don't know what the number of your cells are going to be. You have to reach that region of anatomical, amorphous space with various new tricks using the tools you have, despite all these novelties.

That is a standard version of intelligence. That is how they measure IQ tests: they'll say, you've got these parts, now creatively use them to solve this problem you've never seen before. This is a system that does that, but not in 3D space. It does it in anatomical space, and it means that the living material has this kind of beginner's mind aspect to it, where, in fact, at least most organisms do not assume that everything's going to go the way that it used to go for its ancestors. They have this ability to creatively use their genome and its products to solve problems. That means that not only can they do things like this, they can also do things like this.

Slide 13 of 26 · Watch at [16:35]

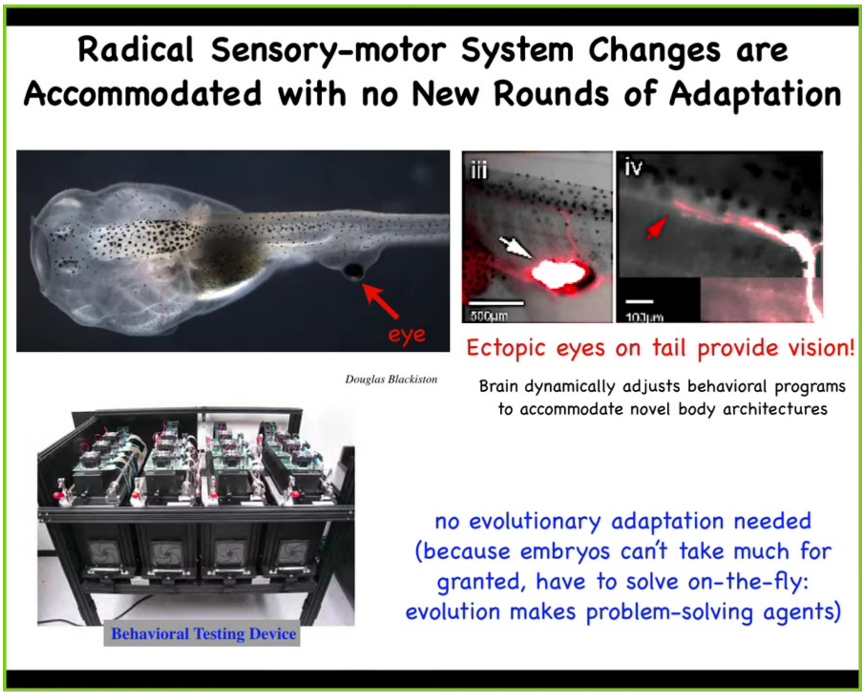

Here's a tadpole. Here's the mouth, the nostrils, the brain, the spinal cord here. What we've done is prevent the primary eyes from forming, but we did put an eye on its tail, and we found that these animals can see quite well. How do we know? We built a device that trains them in visual behavioral tasks, and they can see. This is shocking. Why don't these animals, where the eye does not connect to the brain—so here the eye makes an optic nerve, that optic nerve might synapse onto the spinal cord, it might go to the gut, it might go nowhere at all. Why does this work out-of-the-box? How come you don't need rounds of evolutionary mutation, selection, adaptation to accommodate this completely novel sensory-motor architecture? Why does this work out-of-the-box?

It's because of that plasticity that I told you a minute ago, because evolution does not make fixed solutions to fixed environments. It makes problem-solving agents. These cells never assumed the tadpole was going to be the same all the time anyway. They figure out a way to make it work. And one way they do that. I'm using this example of cells navigating anatomical space as a case study or a model of an unconventional intelligence. We don't have any aliens, but we do have intelligences that live in very different problem spaces. And we can use that as a way to start understanding how do we communicate, how do we detect intelligence? I showed you some ways of doing that. You make perturbations and you notice they can get to the same goal by different means. So that's William James's definition of intelligence. Okay, so we've detected them. What's the next step? Well, how do we communicate with them? If we're going to have a relationship with these intelligences, how do we communicate with them?

Slide 14 of 26 · Watch at [18:21]

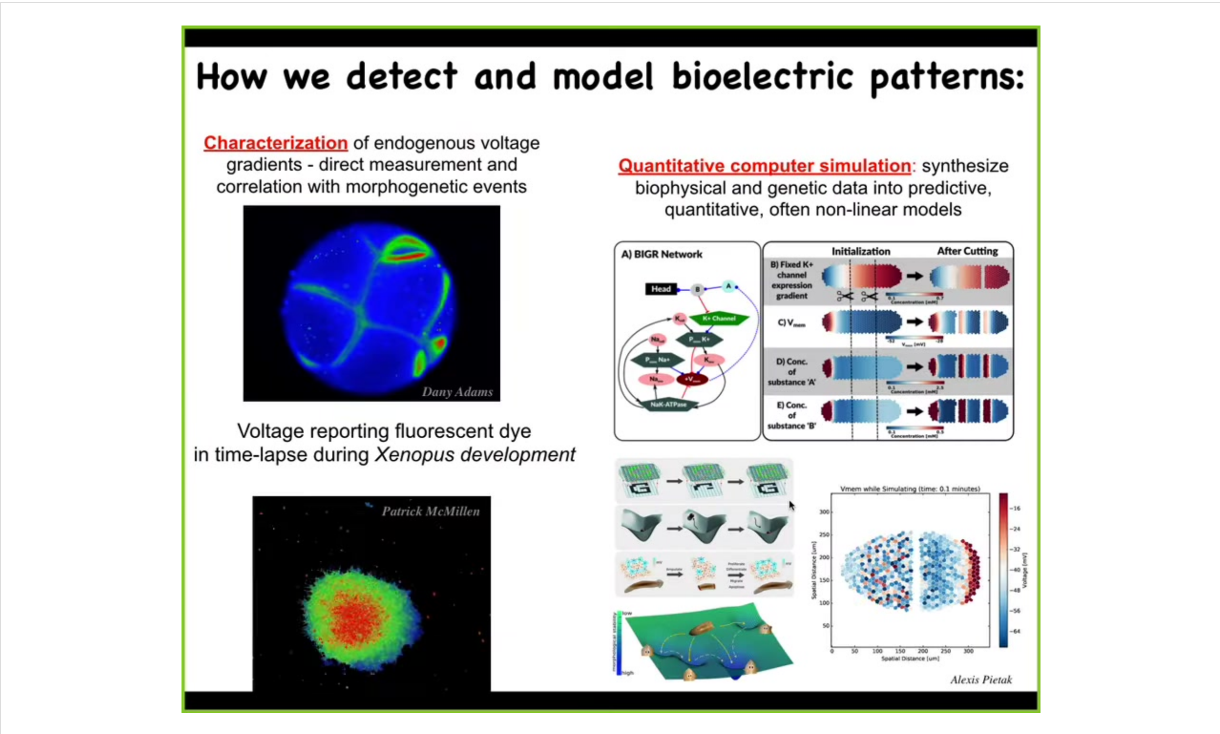

We discovered that a very important kind of communication with them can take place through the bioelectric interface. Just like you and I are now communicating through the bioelectrical interface of your retina and your auditory stream and your nervous system. It turns out this is also a very good way to communicate with the cellular intelligence of your body.

What we've developed are tools that let us see the electrical signals moving between all of the components of this collective intelligence. This is just like neuroscientists do in the brain, except we can now do this anywhere in the body. The colors indicate voltage levels. Here's a bunch of cells in a dish. This is an early frog embryo. And now what we can do is we can read the mind of the body. By reading these electrical signals, we can try to interpret them and then try to understand what are the memories, the preferences, the goals, and so on of this new collective intelligence.

Slide 15 of 26 · Watch at [19:18]

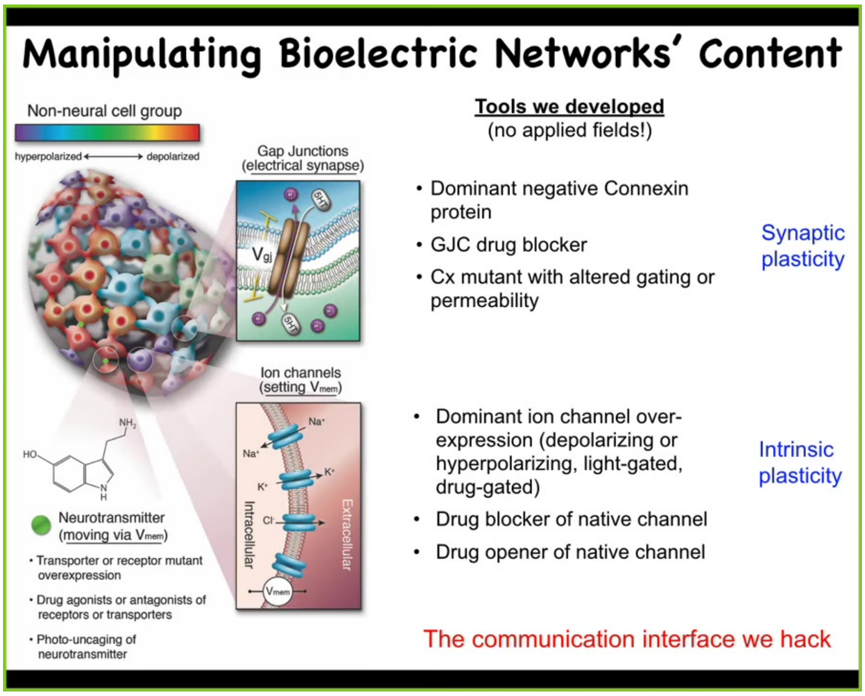

We now have a way to also write information. We don't use electric fields or waves or frequencies. We manipulate the interface that the cells are normally using to control each other. That means ion channels and electrical synapses, known as gap junctions, just like in neuroscience. This is the communication interface that we hack. If you do this, you can do some interesting things.

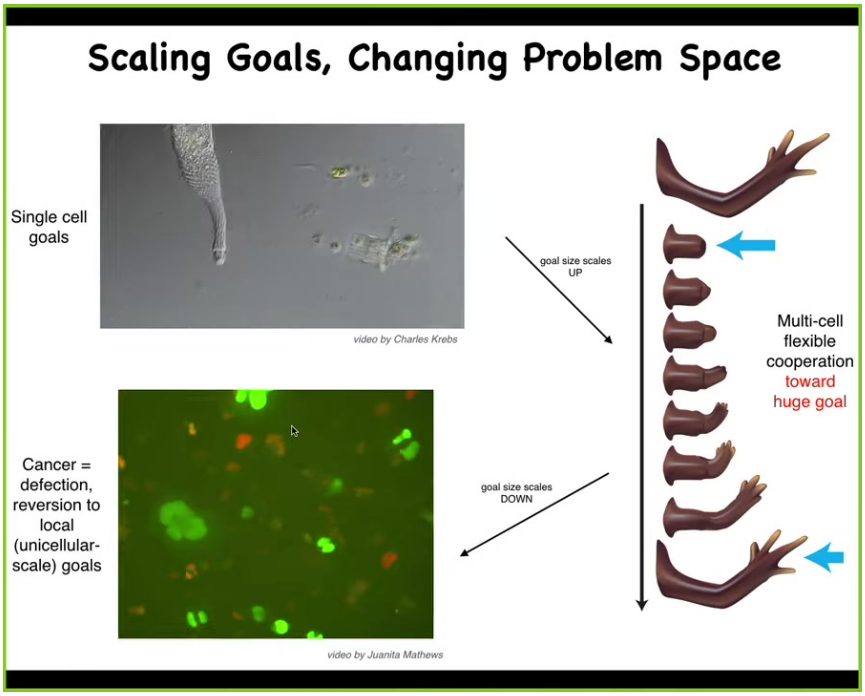

First of all, you can shift the border of the self of an agent. To explain what I mean, this single cell that I've shown you before has very tiny goals. All of its goals are in terms of a space-time radius. The cognitive light cone is very small, both in space and time. But when they get together, these cells can pursue really grandiose goals. So here's the limb of a salamander. If you amputate that limb anywhere here, the cells will rapidly build, and then they stop when it's done. So their goal, the goal of the collective, is to build this thing. It's a goal, because if you deviate them from that goal, they will work really hard, they will achieve it, and then they will stop. No individual cell knows what a finger is or how many fingers you're supposed to have, but the collective absolutely does, as you can see here. It's very reliable. The scale of the goals, the border between self and world, has grown huge. The self is gigantic here. The self is very tiny here. But that process can break down.

Here are human glioblastoma cells. These are cancer cells that are disconnected electrically from the collective. At that point, they're back to their unicellular lifestyle. They can't remember the large-scale goal of building a nice brain or anything else.

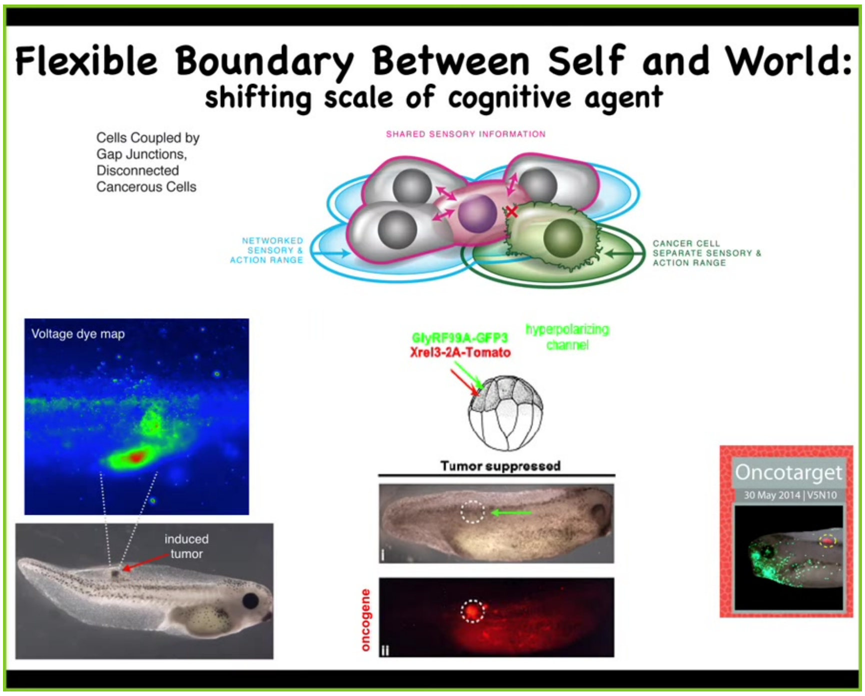

Once you understand that, you can control that electrical connection between cells and force a larger cognitive light cone onto the system.

Here's a human oncogene injected into this tadpole. It will make a tumor. The tumor has this really weird electrical state where you can see the cells disconnecting from their neighbors. But if you prevent them from doing that, even though that oncoprotein is blazingly expressed in the same animal, it's tons of that nasty oncogene, there's no tumor. Because what you've done is kept the cells connected to the neighbors; they're part of this mind meld of building nice skin, nice muscle. The cancer cells aren't more selfish. They have smaller cells. Now the border between self and world has shrunk and we've re-inflated again. That's one thing you can do through this interface: you can play with the border between self and world.

We can communicate certain goals.

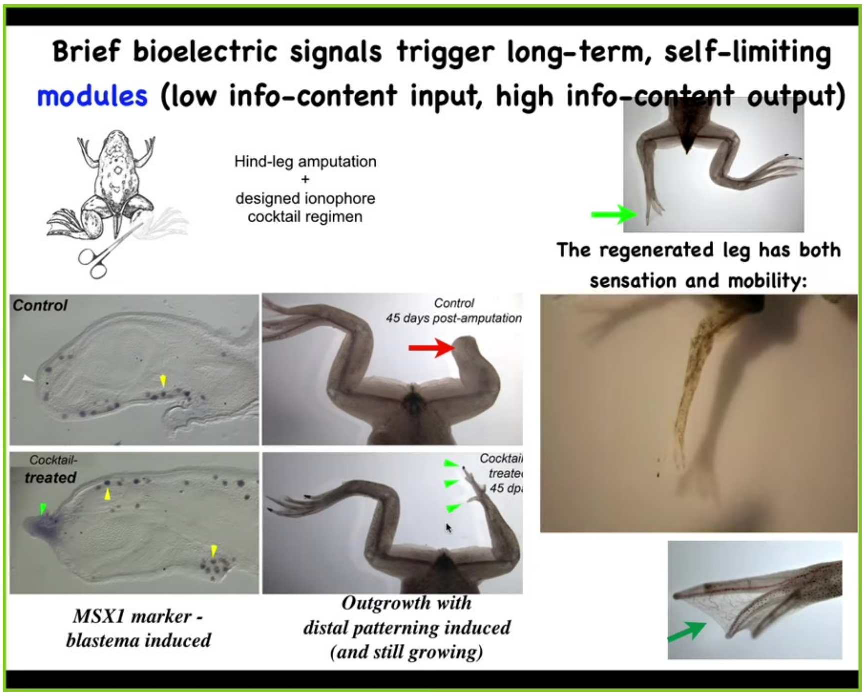

Now that we understand how they think we can communicate certain goals, one thing we can do is to tell a frog. Frogs don't regenerate their legs, unlike that salamander that I showed you. We can tell these cells here to go down the leg-building path instead of the scarring path. By 45 days you already have some toes, you've got a toenail, and eventually a very respectable-looking functional leg. We did that by providing specific ion channel drugs at this interface to put the cells in an electrical state where they remember that they need to build a limb as opposed to scarring and stopping.

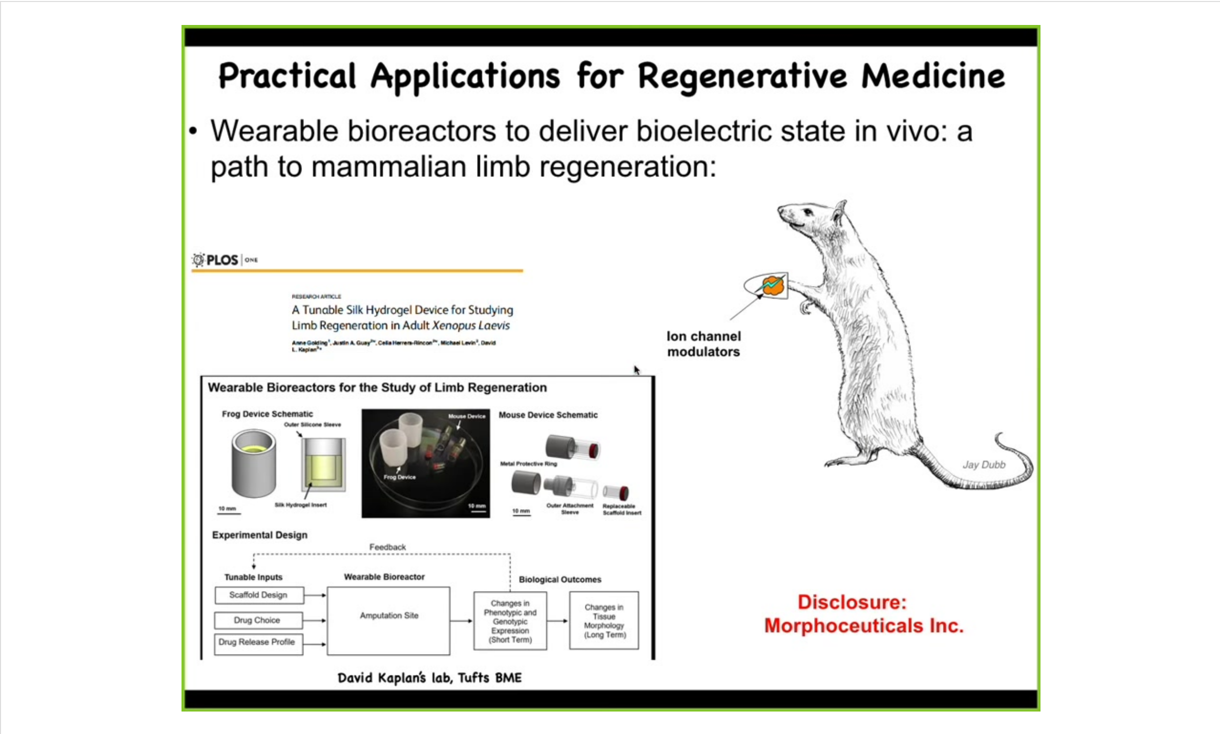

We are currently doing this. I have to do a disclosure because this is a company that was spun out of our lab with David Kaplan to do this in mammals and hopefully someday humans. This is one of those biomedical applications. That first application I just showed you was about cancer. The idea was to normalize the cells, reconnect them to the body, not kill them with toxic chemotherapy, not try to fix the genetics, but to normalize them by reconnecting them to the memories of the body. This is now a second application around limb and organ regeneration. Using electrical cues to tell cells to rebuild the things they already know how to rebuild. It's not just about rebuilding the exact thing in the right place.

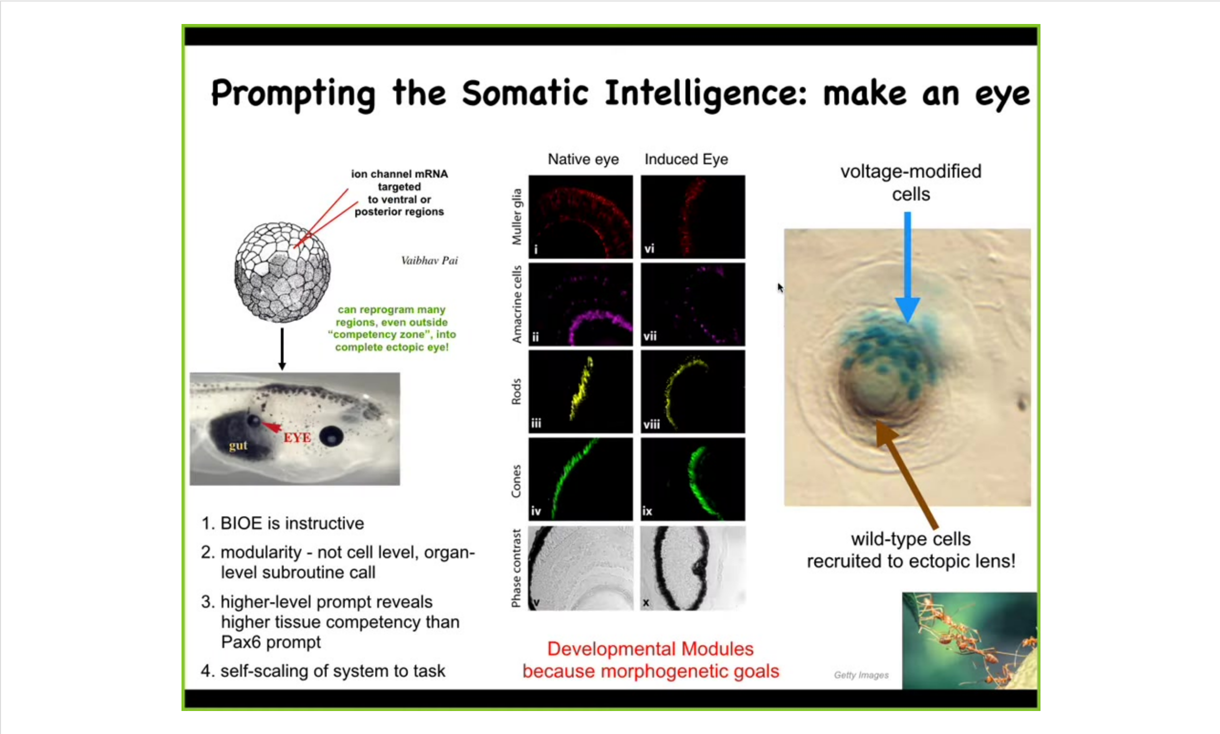

You can also do interesting things like this. We can say, no, instead you should make an eye. And the way we do that is we inject some ion channel RNA that establishes a bioelectrical signal, and that tells the cells to build an eye, and they build an eye. And that eye has all the same lens, retina, optic nerve, all the stuff it's supposed to have. It has some other interesting properties that I don't have time to talk about.

But the idea is, much like when I'm talking to you now, I don't need to worry about reaching into your brain and changing all the synaptic structures in your brain to make sure you remember what I've told you. You do that on your own. I'm only providing a very thin stimulus, and you do all the hard work of managing the underlying biochemistry in your brain to form memories, make decisions, and so on.

Same thing here. We are not telling the cells what to do. We are not talking to the genes or the stem cells. We provide a very high level signal. Build an eye. That's it. The system, being a good multi-scale cognitive system, then takes that and does all of the underlying biochemistry that's needed to implement the message that we've given them.

Slide 21 of 26 · Watch at [24:17]

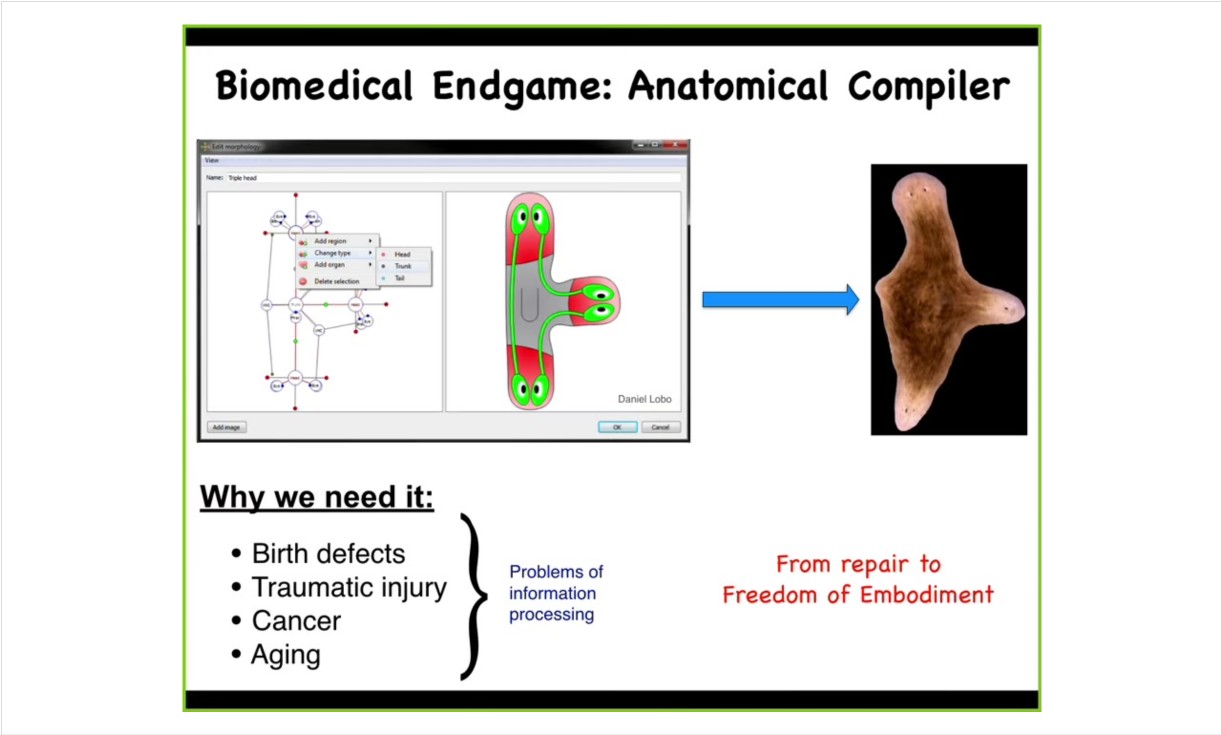

So once you see all this, you realize that what we eventually are going to have is something that I call the anatomical compiler. In other words, once you understand how to tell groups of cells what to build, then you get to move from repair to freedom of embodiment, because then you are going to be able to sit down and draw any kind of a plant, animal, organ, biobot, whatever. The system should be able to compile it into a set of stimuli that would be given to cells to get them to build exactly whatever you want it to build. And yes, it will, if we had that, which we don't yet, but eventually we will. All of this would go away. Birth defects, traumatic injury, cancer, aging, degenerative disease, all of that will disappear once we know how to tell cells exactly what to build. It's a communication problem. It's not 3D printing. It's not gene editing. It's a translator. What you need is a translator interface from our goal to that of the agential material of life.

You don't just have to stick to the standard human shape. If somebody decides that they'd like to have a couple of extra arms and maybe an eye on the back of the head and maybe it should be infrared sensitive, all of that should be completely possible. We augment ourselves now with toothbrushes and antibiotics and education and glasses and crutches and working out and clothes and all of this stuff that we use now to go beyond our so-called natural state. In the future, you're going to be able to do that even more so biologically.

I think the beings of the future, and I don't think it's a very far future, are going to look back and think it is completely unimaginable that people in our day had to live their whole life in the body they were handed at birth. You're born; the details of your body, of your limitations, physical and mental limitations, are currently set by stray cosmic rays that hit the cells of embryos going back however many eons ago. People of the future are going to be aghast at this notion that we just have to live that way and there was no opportunity for improvement. We just have to deal with the consequences of these stochastic processes that don't really care about our goals or our capabilities or our values.

The final thing I want to show you in the last couple of minutes is this.

Slide 22 of 26 · Watch at [26:48]

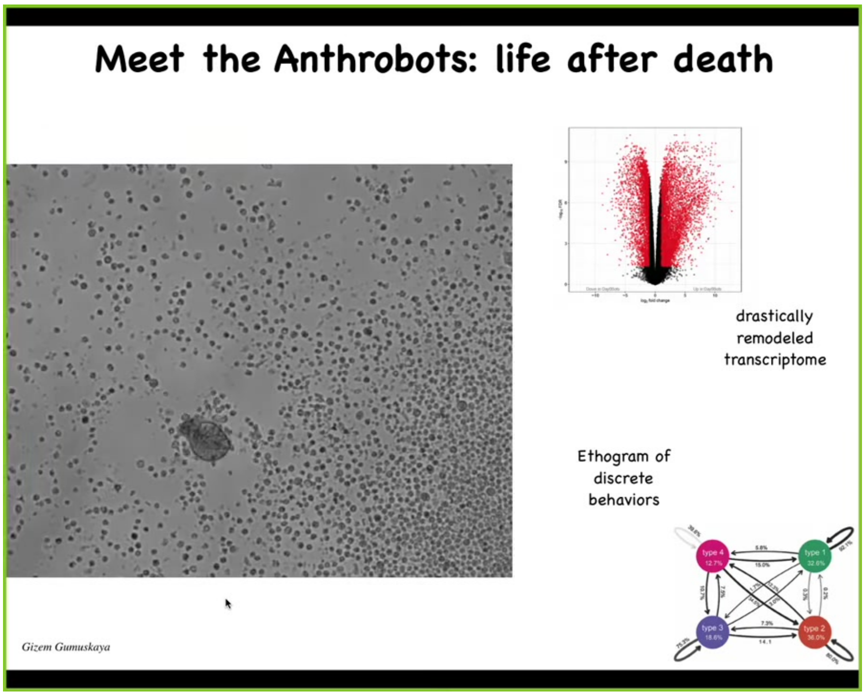

If you look at this little creature, now let's go even beyond. I've shown you how to communicate with the intelligence of the body for repair, then the idea that you can use that to make radical changes in the structure. You can make different kinds of embodiments. Now I want to go even further than that and talk about the plasticity and the capabilities of the material of life. What else can it do?

If you look at something like this, looks like some primitive creature I got off the bottom of a pond. It turns out, if you were to sequence the genome, it's 100% Homo sapiens. We call these anthrobots. These are made of adult human cells. Human patients donate tracheal epithelial cells, and we have a way of enabling them to self-assemble. We don't do much. The cells do all the hard work. They self-assemble into this little motile creature.

This creature has four different behaviors that it switches between, you can build a little ethigram of its behaviors. If you look at the gene expression, this volcano plot, they express about 9,000 genes differently than the cells in the body that they come from. In other words, about half the genome is drastically remodeled. We haven't touched the DNA. We haven't added any weird nanomaterials. There are no synthetic biology circuits. These are completely wild-type cells. Yet they're capable of having a very different embodiment. This looks nothing like any stage of human development. They pick and choose what genes out of the genome they're going to express.

Slide 23 of 26 · Watch at 28:23

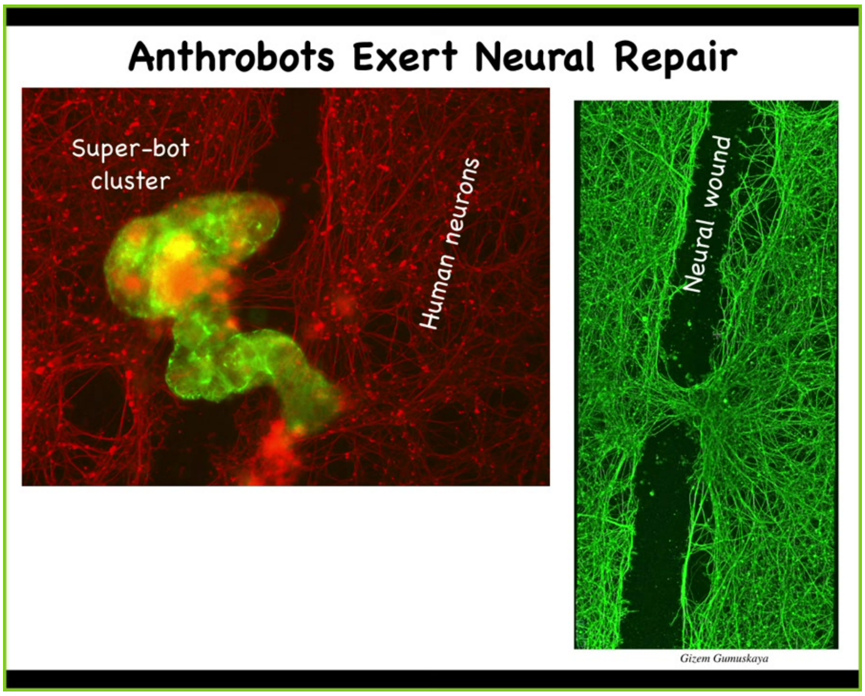

They have some wacky capabilities. If you placed a bunch of human neurons here, these anthrobots will come and they'll sit down in this superbot cluster and they start repairing the gap. So they'll sit there and they'll knit the neurons together across this wound. Now who would have thought, who knew that the cells in your airway are capable of becoming a self-motile little creature that likes to heal neural wounds? We would have never known that.

There's two aspects of note here. The biomedical aspect is that someday you will have biobots made of your own cells that can be injected into your body to make repairs. You don't need immune therapy because they're your own cells. They won't be rejected. Can we learn to program these things to enhance their healing abilities so that they can have all kinds of useful activities in the body?

But the other thing is the plasticity, the incredible plasticity of life. There's never been any anthrobots. There's never been any evolutionary selection to be a good anthrobot. Where does this come from? We can talk about that too.

Slide 24 of 26 · Watch at 29:24

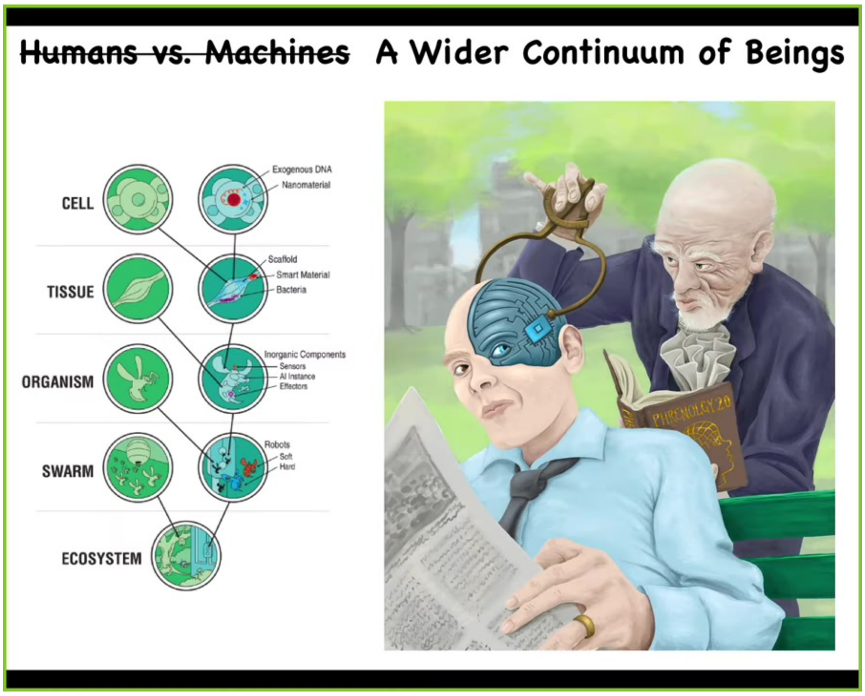

But what that means, that incredible plasticity means that almost any level in biology can be augmented or manipulated by engineering.

And so when we talk about ethics, a lot of people like to compare, 'I'm a real human' and then there are these machines and we're very different. The problem you're all going to be facing in your lifetime is not what to do with language models sitting inside a server somewhere. It's this: this guy has had various modifications. And there are a lot of people who are going to be interested in figuring out, is this a "real human"? What are we going to do? Is it 51% implants and now you're not a real human? What is going to happen to beings that are extended, rationally extended and modified in other ways? And that opens up numerous questions and is going to be very difficult and lead to a lot of moral lapses if we insist on binary categories where, using old school methods, we're going to try to figure out if this is a human or a machine.

I'm going to skip a couple things here because we're running out of time.

Slide 25 of 26 · Watch at [30:34]

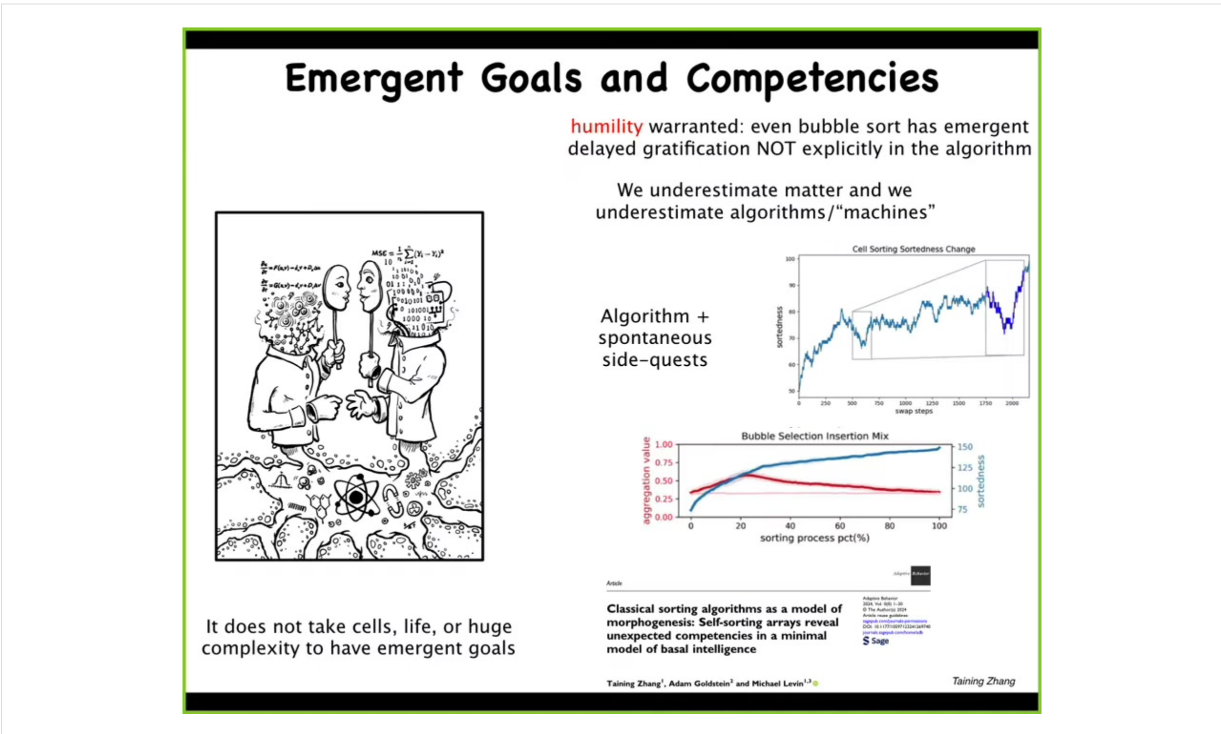

Just to end on this. This is not about biology at all, I don't think. We have started looking at extremely minimal systems. If any of you take programming, you will have heard of sorting algorithms, things like bubble sort. All computer science students study them. They are very short, simple, deterministic algorithms. And what we found is that even those things, when you look at them the right way, are doing things in addition to what is actually in the algorithm. In other words, the algorithm makes them sort numbers, and they sort the numbers, but also they do other stuff. They have these other side quests that are nowhere in the algorithm. And so they have a simple version of what we have. In other words, you have to obey physics, but within the context of that, you do all kinds of things that are not actually prescribed by the physics at all. It doesn't take life, complexity, or a brain to be able to do that. I think it's baked in all the way down.

And that means that all of natural life is a tiny little speck in this enormous space of possible bodies and minds, pretty much any combination of evolved material, engineered material, software, and patterns that are not in the physical world at all—some kind of being with some kind of embodied intelligence that we are going to have to learn to live with. And part of the ethics of the future is developing a kind of synth biosis, a mutually beneficial interaction with beings that are not like us at all.

Slide 26 of 26 · Watch at [32:13]

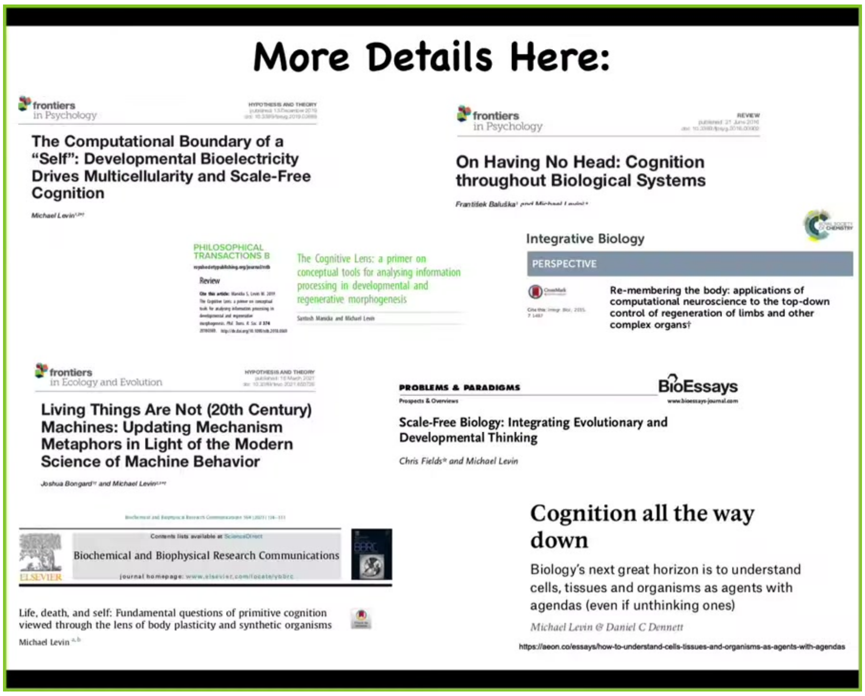

I'm going to thank the postdocs and the students who did some of the work that I showed you today.

Some disclosures. There are three companies that have spun out of all this stuff.

Our funders, and most of all the systems that we learn from, do all the heavy lifting.

I will stop here. I'll thank you and take questions.

Thank you for reading.

More lectures

You can find more of my lectures here.

Follow my work

Twitter • Blog • The Levin Lab

Want one for your lecture?

Want something like this for your own talk? Reach out to Adi at adi@aipodcast.ing.