A Lecture Companion

"Scaling Intelligence in Biology, Artificial Life, and Beyond" by Michael Levin

Michael Levin

Recorded on May 11, 2025

About this document

This document is a companion to the recorded lecture *"Scaling Intelligence in Biology, Artificial Life, and Beyond" by Michael Levin*, recorded on May 11, 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 ~20 minute very rapid talk reviewing ideas around the scaling of intelligence in unconventional substrates.

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.

Slide 1 of 20 · Watch at 0:00

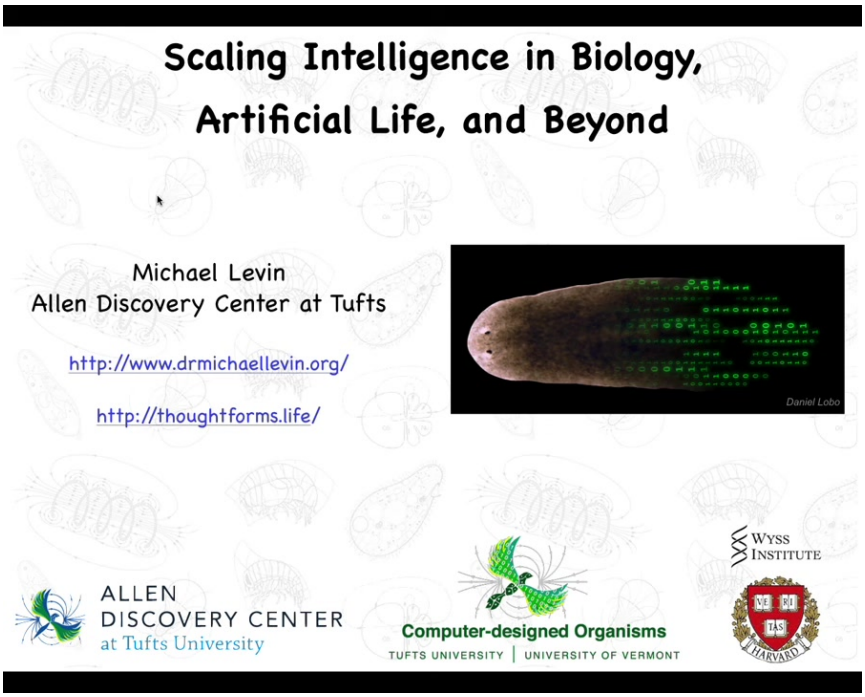

challenge in 20 minutes of transmitting some fairly unconventional concepts. I will do that. I can answer questions about the details of how it all works. At this website you can download all of the data sets, the software, the controls, everything is here. This is my own personal thoughts about what this all means.

Slide 2 of 20 · Watch at 0:22

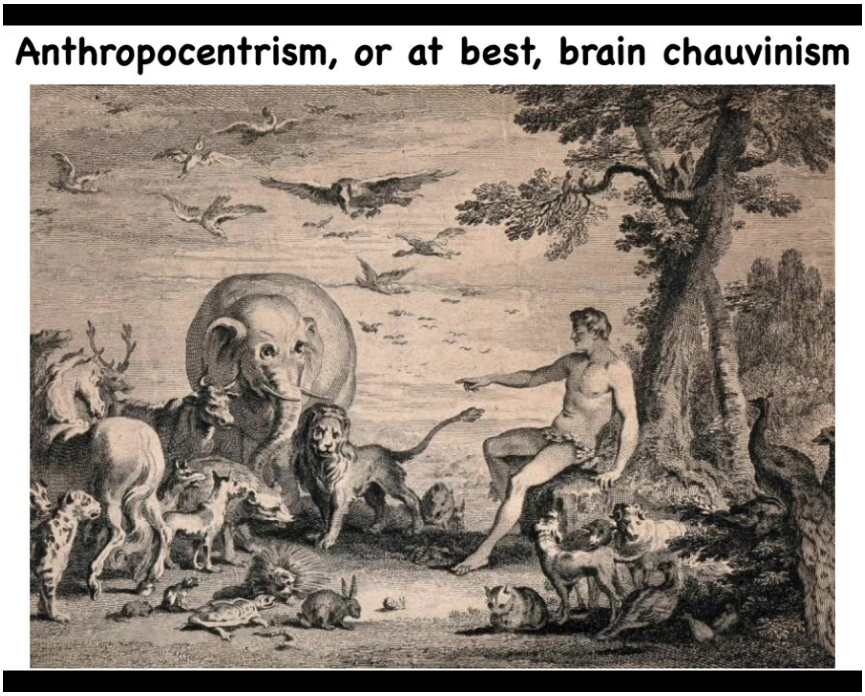

I want to start with this classic piece of art called "Adam Names the Animals in the Garden of Eden." This is a worldview that, even for people who are not overtly religious, really permeates modern science. It's the idea that there are these distinct categorical species. They're all different from each other. They're innumerable. We know what they are. And then Adam over here is different. This we're going to have to blow up and I'll show you the ways in which we dissolve that.

But the one thing that is interesting about this depiction is that, according to the old biblical story, it was on Adam to name the animals. God couldn't do it. The angels couldn't do it. It had to be Adam that names the animals. In these ancient traditions, naming something means that you've discovered its true inner nature. By discovering its name and giving it the right name, you've learned something profound about it. And that part is absolutely right about this. So we're going to try to get rid of this anthropocentrism and, in fact, brain chauvinism, but we are going to keep a part of this.

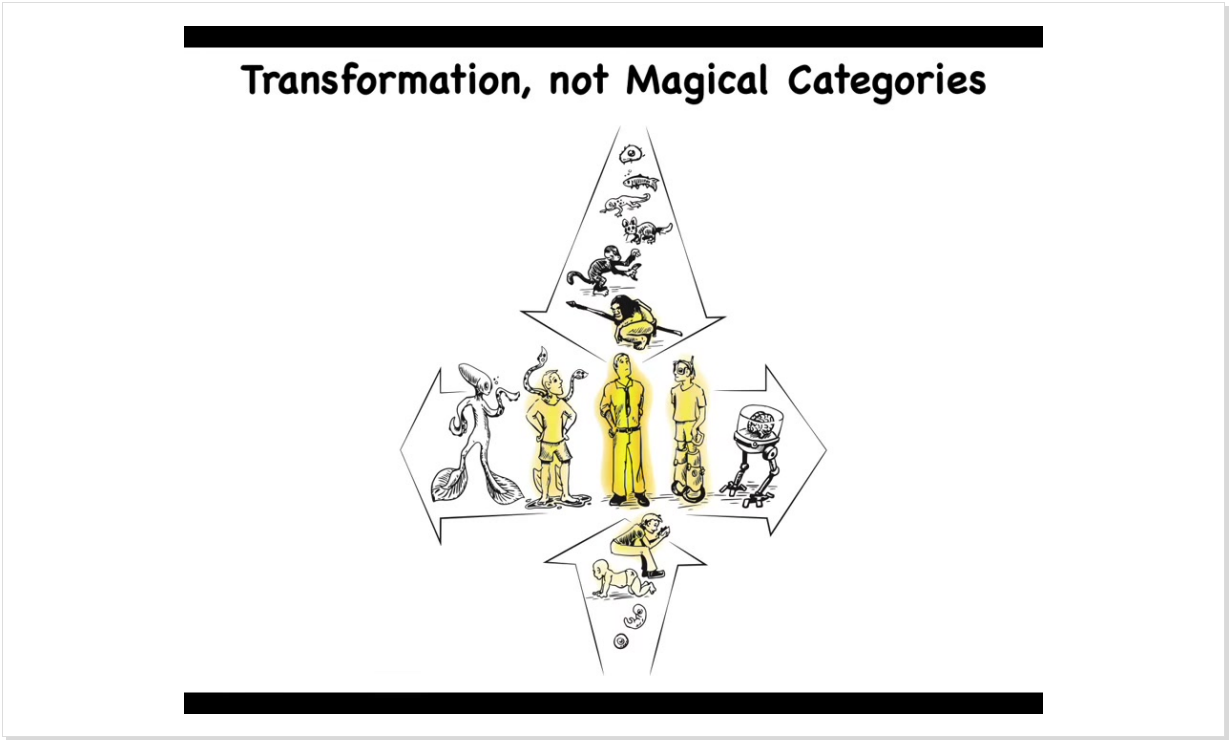

So the first thing is that we know now, from evolution and developmental biology, that this sort of standard modern adult human, which features so prominently in the stories of Philosophy 101, all of these things they talk about, humans do this and we do that. But actually, we know that we are at the center of several continua. These are continuous slow processes that got us from a single cell on an evolutionary scale or a developmental scale. And whatever you say about this human, as far as the ability to understand, the ability to have moral worth, to have responsibilities, to have credit and blame and all those sorts of things. Whatever this magical, a gentle glow that modern standard humans have, you would have to be able to tell a story of how it got here slowly and gradually from these continua.

More importantly, there's now an additional continuum, which means that both along the biological axis and the technological axis, you can make slow and gradual changes and get progressively further and further from this archetype. What we need in science and philosophy are stories of transformation, not magical categories such as a real human or proof of humanity certificates or any of that stuff. We are really at a place where we now understand none of that works.

The Power of Unification

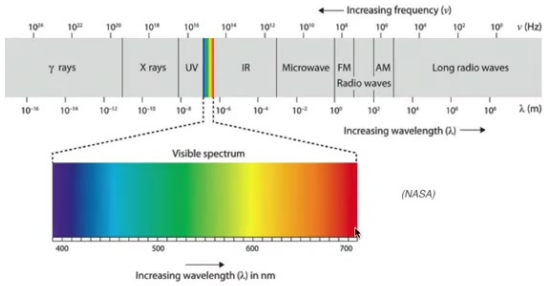

(NASA)

- understand how diverse phenomena are a continuum
- create technology to observe and detect parts of that continuum of which we were previously oblivious
- utilize this knowledge for a myriad of applications that improve quality of life

What I think is really interesting is an analogy to the electromagnetic spectrum. Back in the day, before we had a mature theory of electromagnetism, we had lightning, and we had static electricity, and we had light, and we had magnets and various other phenomena. We thought those were all different things. Not only did we think those were all different things, but we were only directly sensitive to an extremely narrow part of this whole spectrum. What happened? We acquired a good theory of electromagnetism, and that allowed us to do several things. It allowed us to understand that things that we think of as quite different are actually the same thing. They are different expressions of the same underlying dynamics. An incredibly powerful unification. Then it allowed us to build technologies and to realize that while we are directly only sensitive here, that's just a contingent fact of our evolutionary history. With tools, with the right equipment, we can become sensitive to all of this stuff. There are many applications. This is what I'm really interested in.

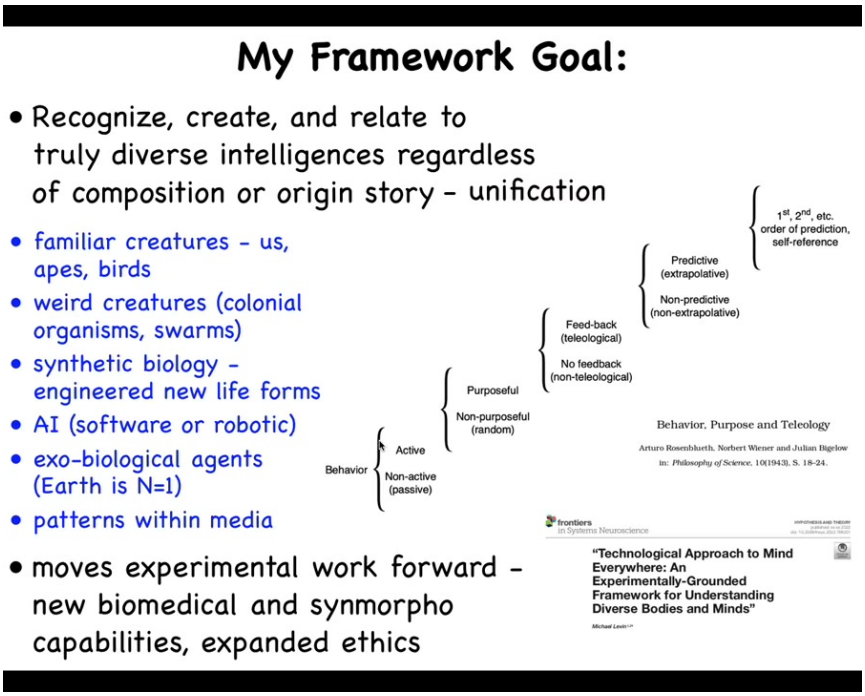

I'm really interested in creating tools, conceptual tools, and also practical bench tools, because the implications of the work that I'm describing here have very practical effects on things like birth defects, regenerative medicine, cancer medicine, bioengineering.

What we want are tools that allow us to see across the spectrum of minds. That is, to fight the kind of mind blindness that we have by default; because of our evolutionary history, we're hyper-fixated on three-dimensional space and medium-sized objects moving in three-dimensional space as intelligent beings.

I'm interested in developing a framework where we can see how you get all the way from passive matter, and by the way, I'm no longer sure there is any such thing, and all the way up to human level and beyond of metacognitive kinds of minds. These are the things I want to understand together as being on the same spectrum, not just mammals and birds, but all kinds of weird organisms, including things that are not even themselves physical, such as patterns within media.

I think we need this, not only for biomedical and bioengineering applications, but because this is, I think, a profoundly important step towards an ethics of a mature species that has some chance of survival long-term, demeaning us.

The first thing I want to remind us all of is that we all start life as a single cell. It's a little bag of chemicals that obey physical laws. Slowly but surely, this amazing, magical process of embryonic development takes the system from being the domain of physics and chemistry to being the domain of psychology and psychoanalysis, or at the very least, behavioral science.

We know from developmental biology that there is no special magical point where these things kick in and you become a mind as opposed to previously being just physics. We need to understand the scaling of this, of mind as it goes from these very simple things to something much more complex.

There are some other paths along this journey. I won't have time to talk about it today, but if anybody wants to ask me, we can discuss what happens later in terms of a dissociative identity disorder of your body, which is basically cancer. There are some other even weirder things that can happen, which I will describe in a moment, where your cells can have an entirely new life, even after the death of the original human donor as anthrobots.

The key to understanding this process is that it works very differently from how we build our computers and robots.

Slide 6 of 20 · Watch at 6:49

There are some interesting differences. We are made of an agential material. The reason that robots and computers, at least for now, don't get cancer is because they are made of passive parts, and then hopefully the whole system has some degree of intelligence, but life isn't like that.

Every part of our body has agendas. It has competencies. This is a single cell. This is known as a Lacrymaria. Just one cell, no brain, no nervous system. This thing has massive competency within its own tiny little cognitive light cone. It only cares about things with a very short spatiotemporal horizon. But put together, they can form much greater things.

A lot of folks will say that you're saying that intelligence goes down to the single-cell level.

Slide 7 of 20 · Watch at 7:43

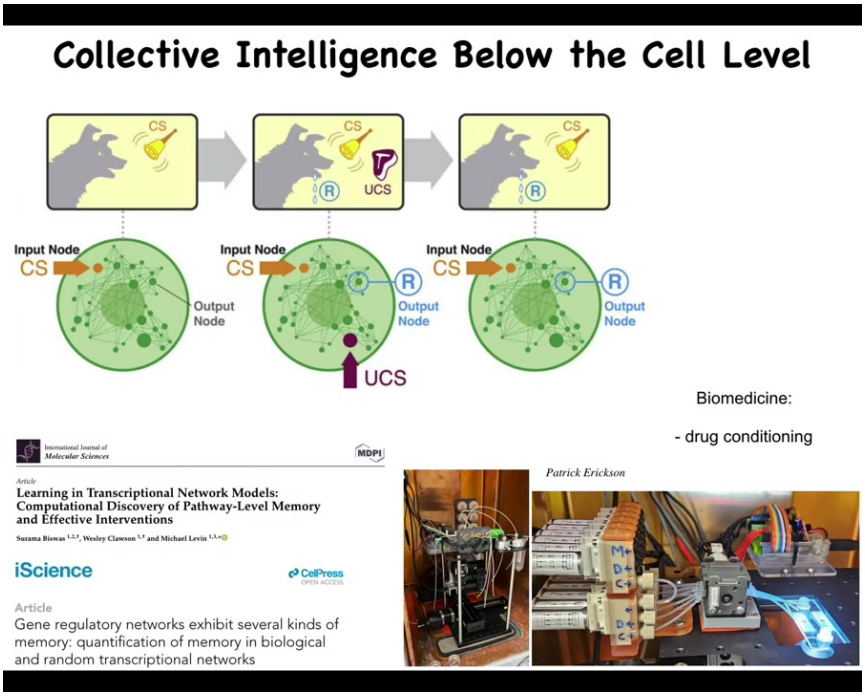

That's weird enough, but actually it's well below that because even the molecular networks inside of cells already have six different kinds of learning they can do, including Pavlovian conditioning.

And they have, if you're interested in things like causal information theory and Tononi's models of phi and integrated information that some people think is associated with consciousness, those already exist here at the molecular network level.

Inside of that cell, the material, the tiniest material, already has aspects of learning and integrated information.

And we're making use of that by developing applications and drug conditioning to train these pathways.

Slide 8 of 20 · Watch at 8:32

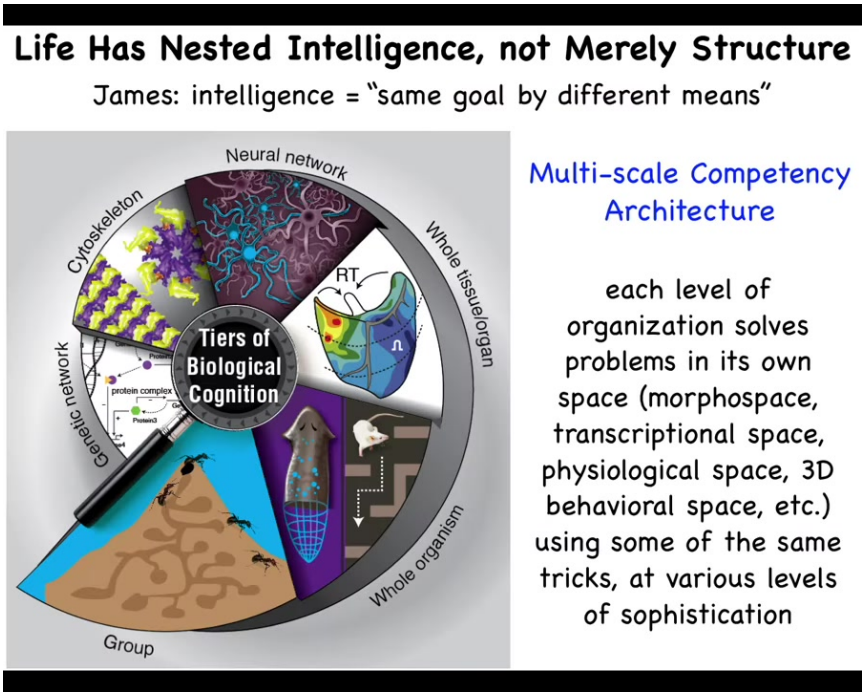

And so our bodies have this amazing nested intelligence I call it a multi-scale competency architecture where every level of organization has the ability to learn and to solve problems in different spaces.

William James defined intelligence as "the ability to get to the same goal by different means." That's a very good definition because it doesn't say it has to be a brain. It doesn't say what problem space you're operating with. It just says you have a capacity to navigate that space in a way to solve problems.

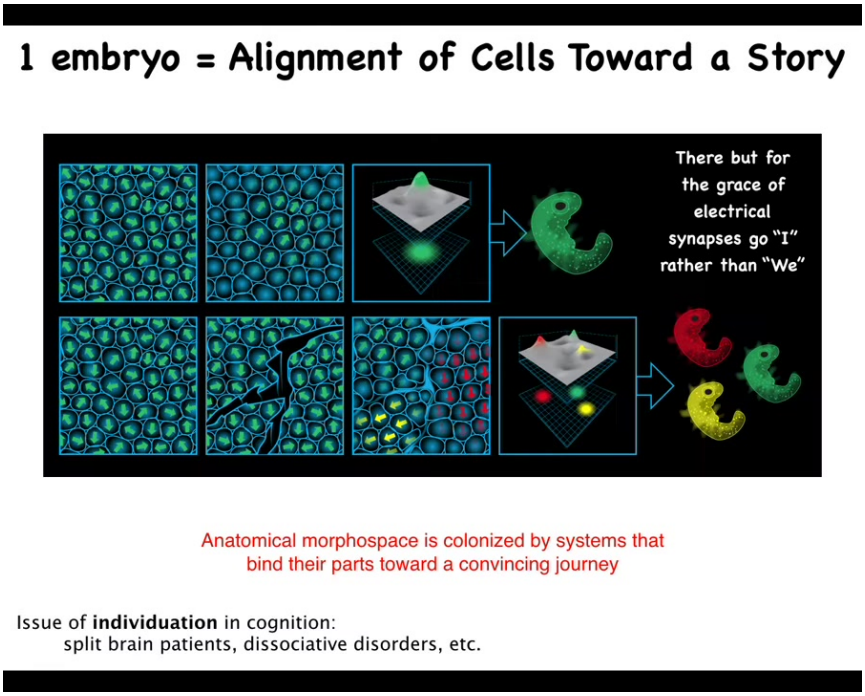

The very first problem that the collective has to solve is to get aligned towards a specific journey in anatomical space.

When you look at an early embryo, here's an embryonic blastodisk. There might be 100,000 cells. We look at this and we say there's an embryo. What are you counting when you say there's one embryo? What is there one of? Because there are 100,000 cells and within each one of those there are organelles and chemicals. What are you counting? I'm going to say that what you're counting is alignment. You are counting the fact that, under normal circumstances, all of these cells are committed to the same story, to the same model of where in anatomical space they are going to go. Anatomical space is the space of all possible geometric configurations of the body. All of these cells are going to collaborate on building one particular structure, meaning they're going to get from the point of a single cell to the configuration of this complicated embryo, because they have all bought into the same story about where they're going to go. What keeps these things aligned is a self-model that they all accept.

Under normal circumstances. I used to do this in duck embryos as a grad student. You can make little scratches in this blastoderm, and every one of these little islands that's formed, until it heals up, doesn't know about the existence of the others. They form their own embryo. Eventually you get twins and triplets.

The number of individuals in this embryo is not set by the genetics. It is not obvious. It is not determined up front. They can self-organize by all aligning towards different aspects of that story, and they will all complete the journey on their own.

The very first thing that has to happen is a bioelectrical mechanism that aligns them all together. It's a bioelectrical network that allows them to remember what pattern they're supposed to build in the first place. That is what makes an embryo an individual rather than millions and billions of cells.

This has many implications in cognitive science. One of the things that my lab has driven is the development of this parallel between cognition and morphogenesis, between the formation of the body and the formation of the mind. They're the same problem. When you look at this kind of dissociation, there are many things that we learn here about split brain patients and dissociative identity disorders.

Slide 10 of 20 · Watch at 11:45

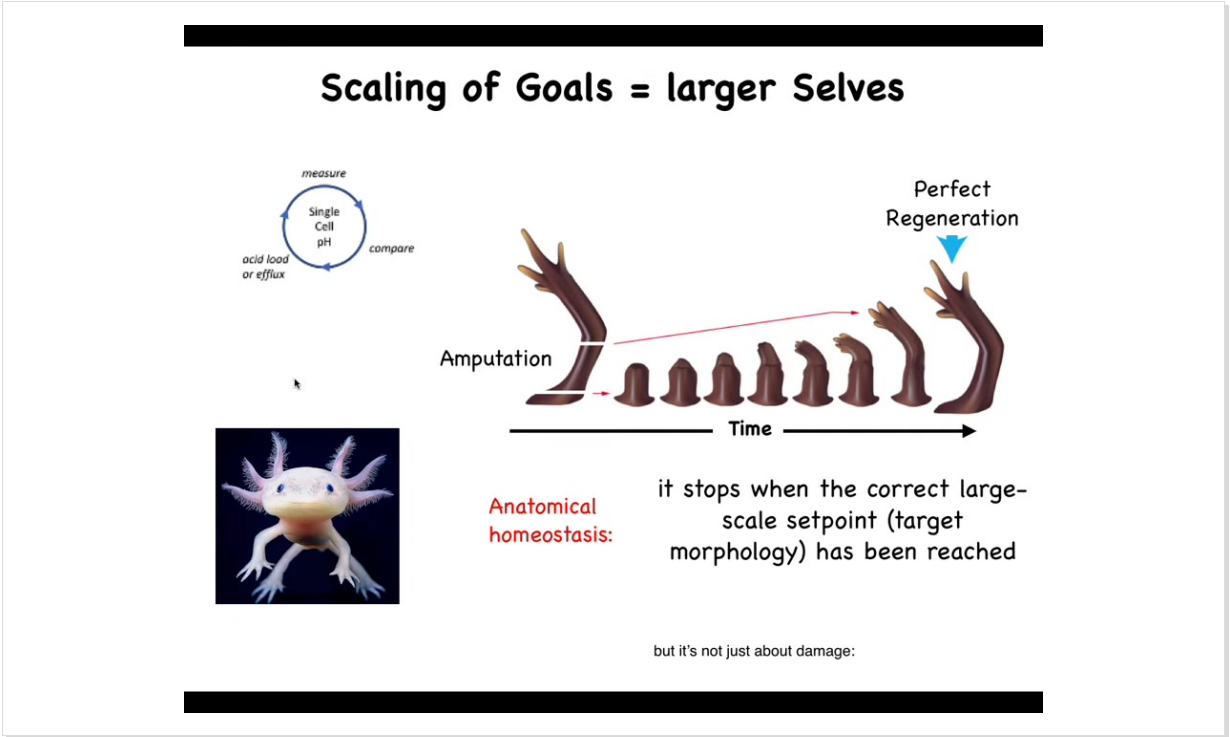

What's happening in this process that's interesting is the scaling of goals. I coined the cognitive light cone, which is meant to be the boundary of the self. It's meant to distinguish the self from the outside world. What it is is the scale of the largest goal you can pursue. Not the range of your senses, not the reach of your effectors, but the size of the largest goal state you can remember and pursue.

In single cells, they have tiny little goals. They have short memories, short anticipatory power in the future. Their goals are things like pH. They're scalars. They're single

numbers about pH within the cell. When they get together into networks, that electrical network allows them to store grandiose goals.

Here's an axolotl. These amphibians regenerate most of their body parts, including limbs, eyes, jaws, spinal cords. If you were to amputate anywhere along the plane of this limb, these cells would immediately sense that they've been taken away from their goal. They would work hard to rebuild, and when they get there, they would stop. How do they know when to stop? They stop when a correct salamander limb has been completed. They stop when they get back to their homeostatic state. No individual cell knows what a finger is or how many fingers you're supposed to have, but the collective absolutely does. It does a means-ends analysis here to get back to where it needs to be once you've deviated it. This whole thing is not just about damage. It's not about fixing this kind of surgical defect.

Slide 11 of 20 · Watch at 13:19

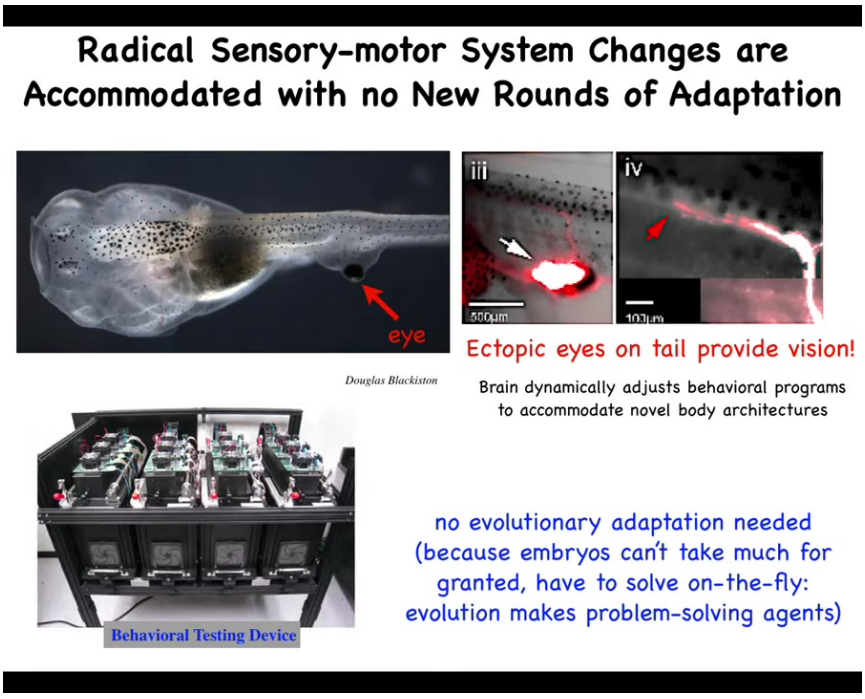

It also allows you to do these amazing things. Here's a tadpole that we made. What you'll notice is there are no eyes where the eyes belong; here are the nostrils, here's the mouth, here's the brain, the gut. We put an eye on its tail.

But what happens with this eye is that it makes an optic nerve. That optic nerve does not go to the brain. It synapses on the spinal cord or sometimes on the gut, sometimes nowhere at all. The most amazing thing about it is that these animals can see. We know because we built this device that trains them for visual learning tasks, and they can see perfectly well.

This is shocking. Why does it not take additional rounds of mutation and selection to radically change this animal's sensorimotor architecture and make things work? You don't need it. It works out of the box.

And it works out of the box because that process, where all of the cells—every single time, beginner's mind—have to solve the problem of what are we, what are we going to build, and where are we going to go, is never obvious to them or assumed that they are what they are. The story that the genetics tells you what you're going to be is not the right story at all. Instead, what the genetics builds is a problem-solving agent that is very creative in interpreting its environment and interpreting the genetics that have been passed down to it.

And so that's why you don't need additional rounds of mutation because the material never thought it was going to be a perfect tadpole in the first place. It's able to adapt to all kinds of novel manipulations.

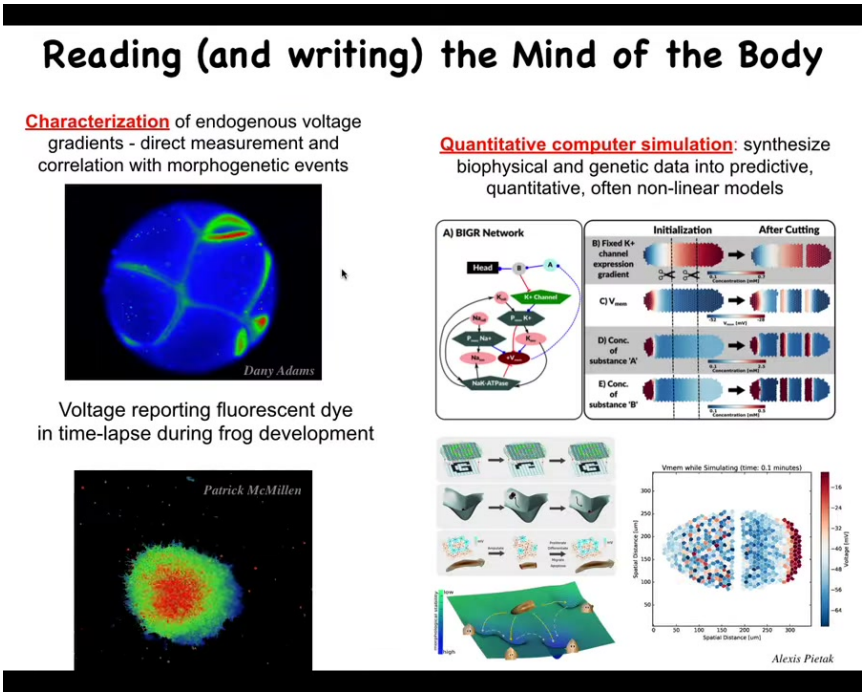

I kept mentioning this word bioelectricity. What we've developed is our methods to read and write the mind of the body the way that neuroscientists have done in the brain. Here you're seeing voltage-sensitive fluorescent dyes; the color corresponds to voltage. What you're seeing is a time lapse. Here are some cells. This is a frog embryo. What you're seeing is all the electrical communications that allow this thing to be more than the sum of its parts.

The bioelectricity is the cognitive glue that binds your neurons together into not just a pile of neurons, but into you. The bioelectricity is also the cognitive glue that binds individual cells into a collective that can remember what a tadpole is supposed to look like. We have lots of tools for doing simulations that we've created. We're trying to merge those with all kinds of connectionist ideas about attractors in memory networks that can do pattern completion and that have the ability to navigate the space.

Let's get back to the beginning of the talk. What I was pointing out is that there are some really weird intelligences out there that we are not familiar with. We use morphogenesis as one example. Groups of cells are an unconventional collective intelligence that navigates a really strange space that we cannot visualize. It's a high-dimensional anatomical amorphous space. Our goal is to learn to communicate with that intelligence, to predict its behavior and to ask it to do different things in biomedical contexts.

In order to do that, we've developed this bioelectrical interface. Now we can directly read the memory states of this intelligence, and we can try to rewrite, give it new ideas.

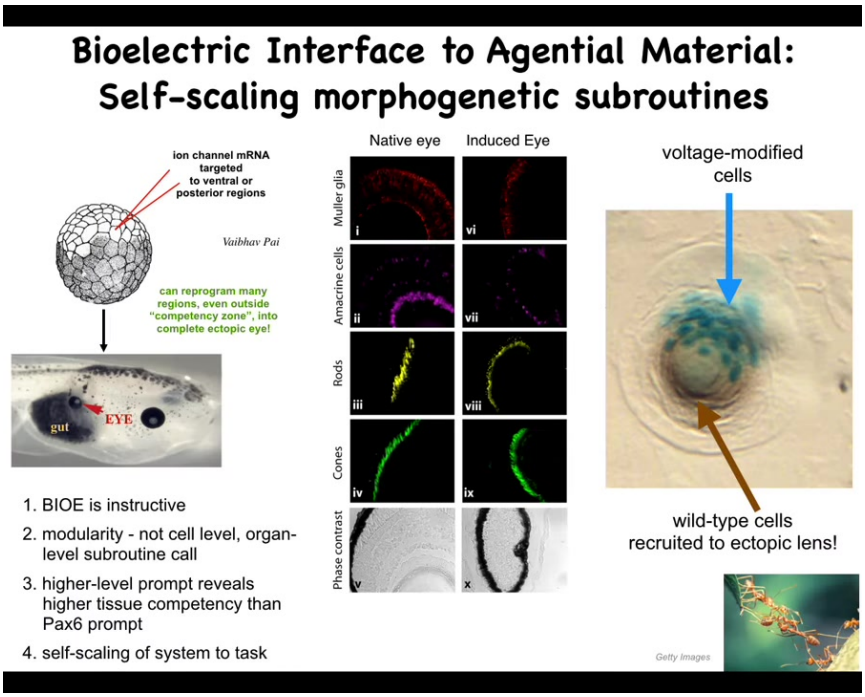

What kind of idea could you give it? One thing you might say is that a proper tadpole should really have an eye on its tail, on its gut in this case, and so the way you would do that is by injecting RNA that encodes a particular channel, ion channel protein, in this case a potassium channel. What you would do is establish a little voltage state here that says to the surrounding cells, you should build an eye. How did we know? 10 years of work trying to understand how the cells interpret these voltage gradients. So when you do that, you make an eye. These eyes have all the right lens, retina, optic nerve, they have all the right things.

They can do another neat trick, which is that if you only inject a few of them, so the blue ones are the ones we injected, they will recruit their neighbors to make this lovely lens that's sitting out in the tail of a tadpole somewhere. We didn't have to touch these cells. All we said is, you guys should make an eye. They take it from there. They say, there's not enough of us. Let's get our friends to help. They convince. And it's a process of trying to infect them with a better world model of what they should be doing. There's a debate that goes on and the cells resist. But when they win, you get this beautiful eye. And of course, there are many other collective intelligences that scale to problem size on their own.

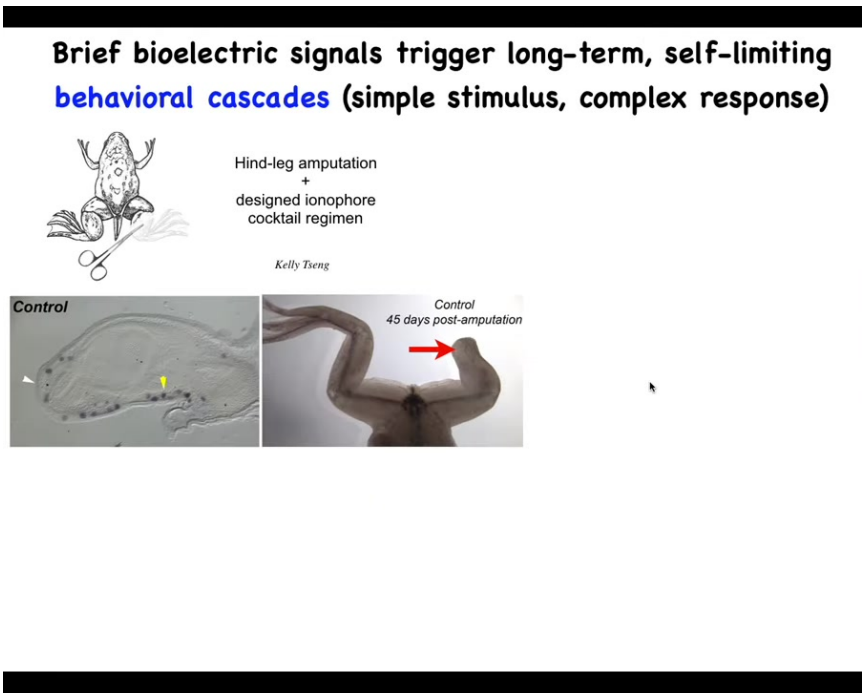

This is a useful application where we're looking for limb regeneration. Frogs, unlike the axolotl, do not regenerate their limbs, nor do we. Forty-five days later, after losing a leg, there is no regeneration normally.

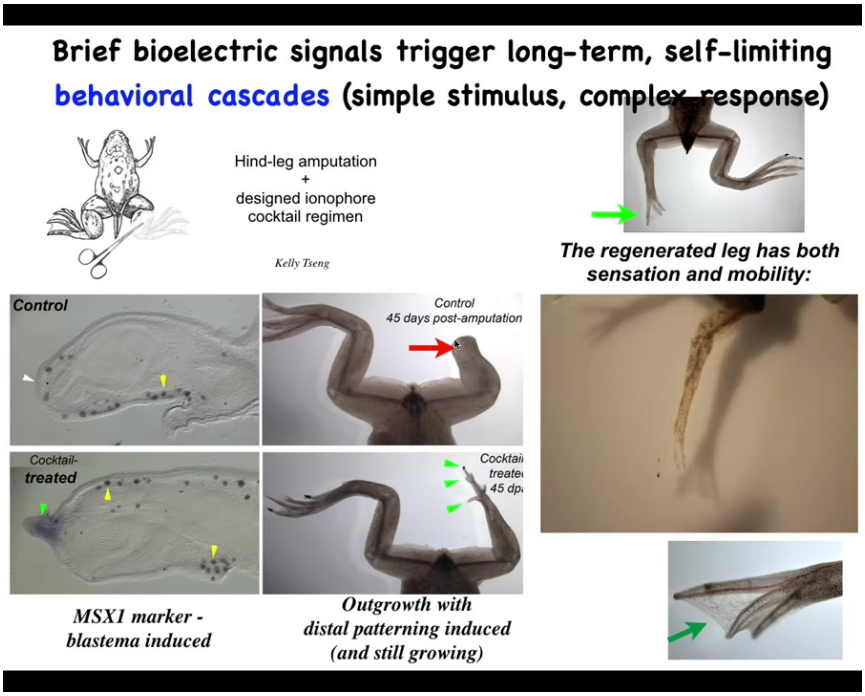

But if we give the cells an early signal, it's a wearable bioreactor with some ion channel payload. It immediately tells the cells to get going. 45 days later, you've got some toes, you've got a toenail, eventually a touch-sensitive and functional leg.

The most important thing here is that with any good cognitive system, you do not micromanage the molecules. When I'm talking to you, I'm not worried about reaching into your brain and having to arrange all the synapses so that you remember what I've said. I'm giving you information on a very thin communications channel, and I'm trusting you as a high-end cognitive system to do all the biochemistry downstream that's required for you to react to what I'm saying. The same thing is true here. This signal was present for 24 hours, and after that, we have shown a year and a half of leg growth, during which time we don't touch it at all. The goal is not to micromanage it. It is not to tell the cells what to do. It is not to 3D print scaffolds for stem cells, none of that. The goal is to convince it on day one that this is the path you should go. The leg building path is the right one, not the scarring path. There you go.

In the last two minutes, what I want to show you is this. So far, what I've been telling you is that we can convince living tissue to repair or remake or reposition normal organs that they already make. I want to show you something even further from this, the remarkable plasticity of life.

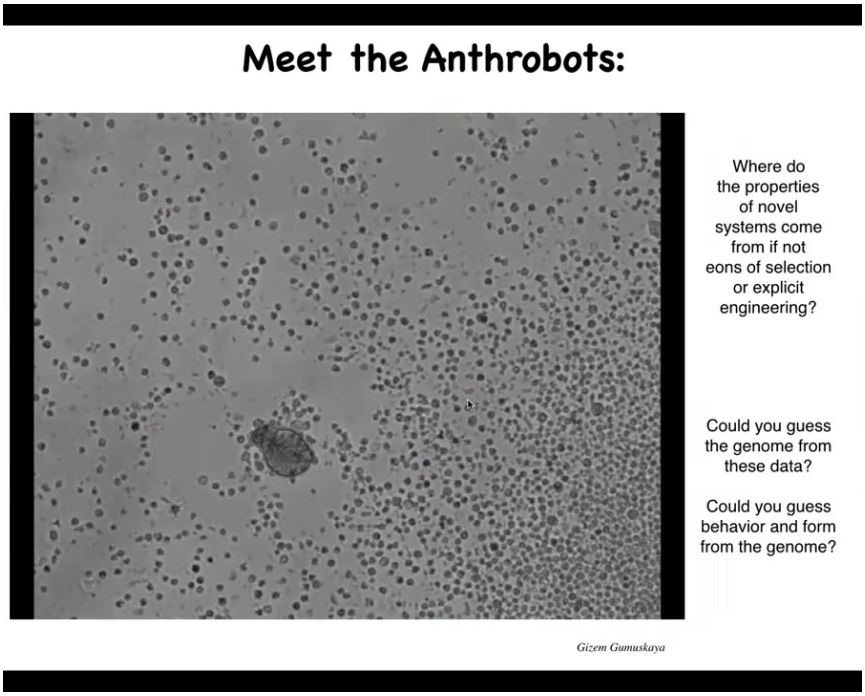

If I show you this video and I ask you what you think this is, a reasonable guess would be that it's a primitive organism that we got from a pond somewhere. And if I ask you what the genome would be, you would guess that it has one of these ancient genomes with these tiny little creatures. I can tell you that the genome here is 100% Homo sapiens. These are perfectly normal adult human cells. They have not been manipulated with any synthetic biology circuits. There are no scaffolds here. There are no transgenes. There's no genomic editing.

What this is is taking cells from an adult human patient, tracheal epithelial cells from their airway, and giving them a chance to have another lease on life, to reboot their multicellularity. The original patient may or may not be alive, but the cells in a slightly different environment, not that different, but liberated from the rest of the body, could have this completely novel life.

You would not know that by looking at it, because this doesn't look like any stage of human development. This is completely novel.

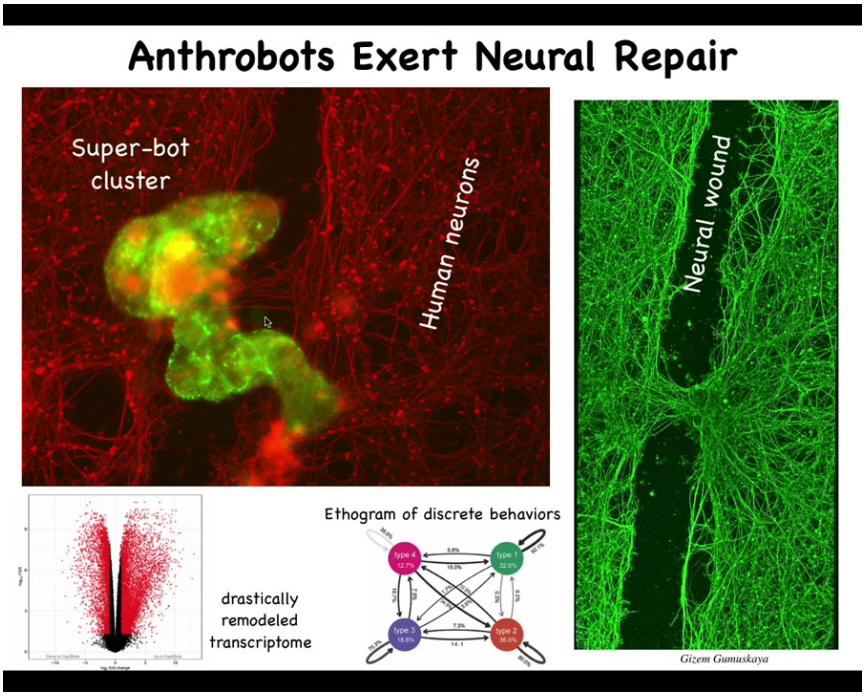

They have novel capabilities. Half of their genome is expressed differently. Each one of these red dots is a gene that it expresses differently than it would have if it had stayed in your airway. About half the genome, 9,000 genes, are completely altered because genetics don't drive what you are. Genetics are a resource book that active systems dip into as affordances.

One of the cool things they can do is if you place a bunch of human neurons and you put a big wound through them, the bots will gather together into what we call a super bot cluster. What they do is start knitting the two sides together. Here you can see when you take it off what happened. Who would have thought that your tracheal cells that sit there quietly for long periods of time just dealing with mucus have the ability to self-assemble into a novel proto-organism with its own life, with its own gene expression, with its own set of behaviors, and these kinds of capabilities that we're only now beginning to scratch the surface of, because this is going to be personalized in-body therapeutics. These biobots made of your own cells don't need immune suppression. You can put them back in the body, and we're now working up the full list of what they can actually repair.

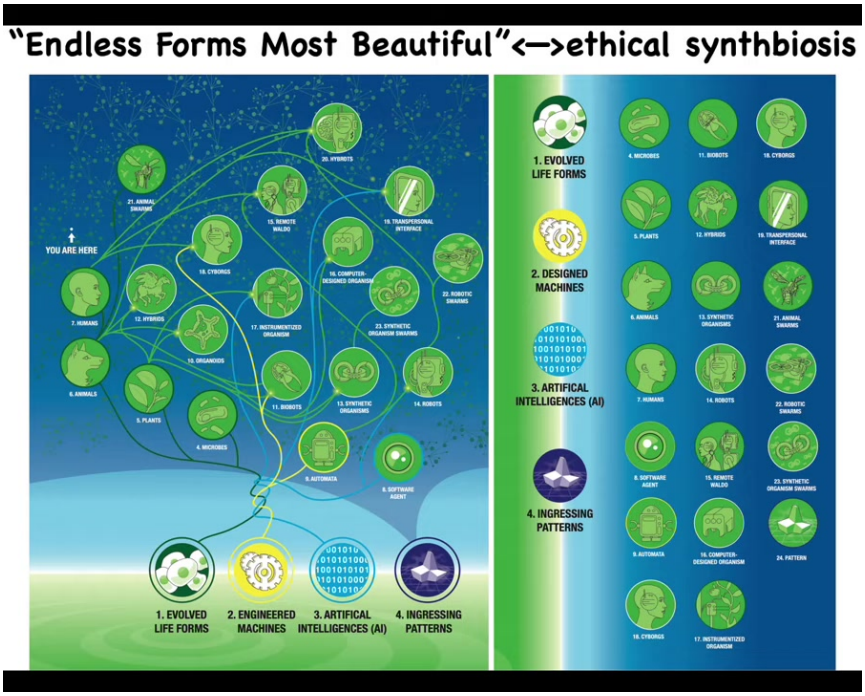

To end, because life is a problem-solving process from the beginning. Almost any combination of evolved material, engineered material, and software is some kind of viable agent.

Life is incredibly interoperable. Hybrots and cyborgs and chimeras of every kind all make use of these ingressing patterns, which is a whole other thing we could talk about from mathematics and computation that create viable beings.

Everything that Darwin saw when he said "endless forms most beautiful," meaning natural life, is a tiny little blip here on this enormous space of beings with whom we are going to share our world. Many of these things already exist. Many more are forthcoming.

What we need is a word GPT-4 came up with for me: "synthbiosis." This idea that we are going to have to get better at recognizing other minds and ethically relating to them so that we can all have a more positive embodiment.

Slide 19 of 20 · Watch at 23:06

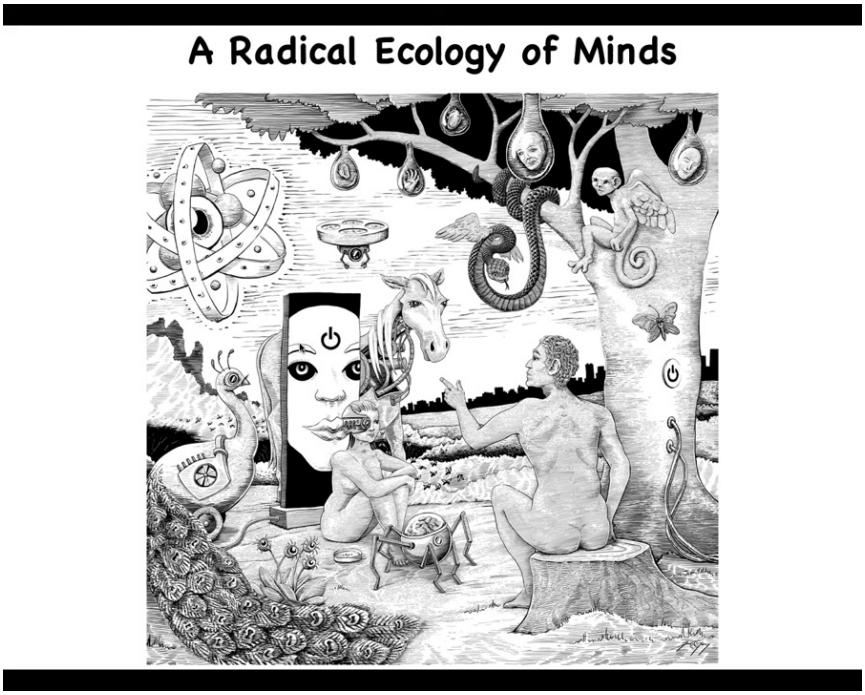

This is what I think the future of the Garden of Eden is going to look like. It's going to be very weird. It is on us to understand what we're dealing with. Xenobots and anthrobots and augmented humans and chimeras of all kinds; we are going to have to raise our game because the things that used to work—what do you look like and where and how did you get here? Meaning factory versus trial and error of evolution, those categories are not going to be any good anymore.

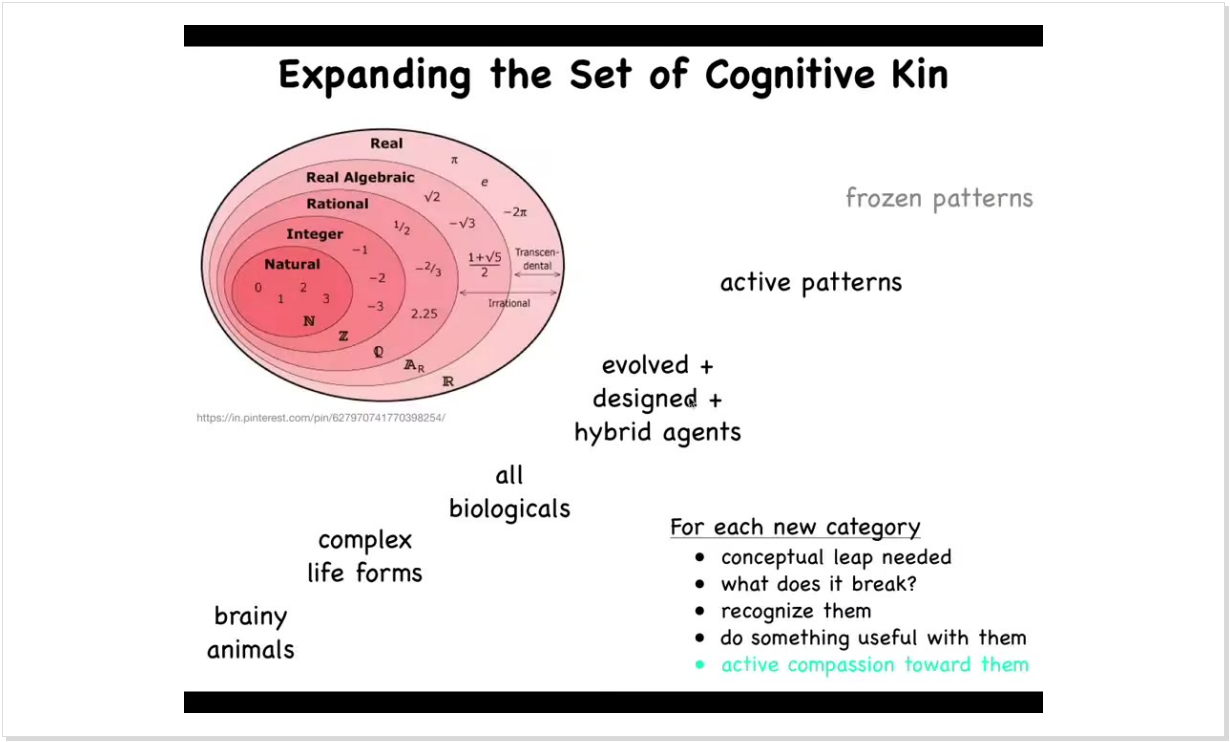

I'm going to skip all this. Point out that there are some amazing people who need to get the credit for all the things that I showed you today.

These are my postdocs and my grad students. We have lots of remarkable collaborators, our funders, disclosures. These are three spin-off companies from our work that support our research. All the biggest thanks go to the model systems because they do all the heavy lifting and teaching us about this stuff.

I will stop here and thank you for listening.

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.