

A LECTURE COMPANION

---

# **Neuroscience Beyond Neurons: bioelectricity underlies the collective intelligence of cellular swarms**

---

Michael Levin

Recorded on October 6, 2023

## About this document

This document is a companion to the recorded lecture *Neuroscience Beyond Neurons: bio-electricity underlies the collective intelligence of cellular swarms*, recorded on October 6, 2023. 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 talk I gave to a neuroscience audience for the Harley Hotchkiss Memorial Lecture at the Canadian Centre for Behavioural Neuroscience at the University of Lethbridge, Canada. 1 hr 4 minutes.

---

## 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](mailto:adi@aipodcast.ing).

The slide features a central title 'Neuroscience Beyond Neurons: How Bioelectricity Underlies the Collective Intelligence of Cellular Swarms' in a serif font. Below the title, the speaker's name 'Michael Levin' and affiliation 'Allen Discovery Center at Tufts' are listed. Two URLs are provided: <http://www.drmichaellevin.org/> and <http://thoughtforms.life/>. To the right of the text is a small image of a green, segmented worm-like organism with a grid overlaid on it, labeled 'Daniel Lobo'. At the bottom, there are logos for the Allen Discovery Center at Tufts University, Computer-designed Organisms (TUFTS UNIVERSITY | UNIVERSITY OF VERMONT), and the WYSS INSTITUTE at HARVARD.

Hello. Good morning, everyone. Welcome to today's Harley Hotchkiss lecture. As you know, the Harley Hotchkiss Memorial Lecture has been organized by the Canadian Centre for Behavioural Neuroscience at the University of Lethbridge since 2013. The series was designed to bring in neuroscientists from all over the world and southern Alberta to present their research and meet the department. The series honors Harley Hotchkiss, a Canadian businessman and philanthropist who supported scientific research. In the 10 years of its existence, it has brought in over 100 renowned scientists, making possible a most valuable exchange of ideas and information.

Today we have with us Dr. Michael Levin from Tufts University. Telling us a little bit about his career and his accomplishments, we have Dr. Arthur Lubzak.

It is a great pleasure to have Dr. Levin here. He got a bachelor's degree in computer science and in biology, and he got a PhD from Harvard University. Among many awards he has received, he is Distinguished Professor of Biology at Tufts. He's an endowed chair and he's director of the Allen Discovery Center at Tufts and Tufts Center for Regenerative and Developmental Biology. Today he will talk about how intelligence can emerge from cellular swarms. I would like to add my personal note that his work really inspired me in how I think about the brain — that it's not just a collection of smart, intelligent neurons interacting with each other to create our behavior. It's not only me: when I spoke with Brian Kolb about what he would do if he were just starting his career, he told me about one of Dr. Levin's videos. So pay attention because this is really exciting work. Welcome and thank you for joining us, Dr. Levin.

Thank you so much. Thank you for that kind introduction and thank you for the opportunity to discuss these ideas with all of you. If afterwards anybody is interested in downloading the primary papers, the data sets, the software, everything is here at this site. So today, I would like to transmit the following main points.

### Main Points:

- **Intelligence and navigation** applies well beyond conventional brainy animals in 3D space
- We are all **collective intelligences**; **neuroscience is not just about neurons**
- Morphogenesis is the **behavior of a cellular collective intelligence in anatomical morphospace**
- The cognitive glue of this collective intelligence is the same as it is for neurons in the brain: **bioelectricity** = the evolutionary precursor of CNS
- We can now read and write memories in this new collective intelligence with **applications in birth defects, regeneration, and cancer**
- This interesting invariance across fields means that **tools and concepts can be ported between neuroscience and developmental biology**, to the enrichment of both.

We will first talk about this idea that intelligence, and in particular navigation of arbitrary problem spaces, applies well beyond conventional brainy animals moving in three-dimensional space. I'm going to stress this idea that we are all collective intelligences, and try to stretch some aspects of neuroscience beyond neurons. I think neuroscience has a lot to tell us about things that are not neurons.

For a detailed example of this kind of thinking, I'm going to use morphogenesis and I'm going to frame it as the behavior of a cellular collective intelligence in anatomical space. Now, in order to have a collective intelligence, you need something binding the pieces towards goals in other problem spaces. That is exactly the same as for neurons in the brain; bioelectricity, or electrophysiology, is another way of putting it. I'm going to talk about how this is what I think is the evolutionary precursor of the central nervous system and its functions. I'm going to show you that we now have the tools to read and write memories into this unconventional collective intelligence. That gives us applications in birth defects, regeneration, and cancer.

Towards the end, we'll talk about this idea that there are interesting symmetries across fields, in particular developmental biology and neuroscience, which mean that tools and concepts can be ported in both directions. I think that's going to enrich both fields.

So the talk comes in three main sections. First, we will talk about some introductions and some ideas to broaden the usual focus on conventional animals, conventional brains, and so on.

The first thing I would like us all to remember is that we all started life as a single cell. So we all start life as a piece of what people call just physics. I don't like that term, but a lot of people will look at something like this and they will say, well, it's a little BLOB of chemistry and physics. This is an unfertilized quiescent oocyte. Then there's this amazing process of embryonic development, which results in one of these complex

creatures or even maybe something like this. All of us take this journey. All of us go from physics to having a complex metacognitive mind.

It is important to realize that during this process, there is no magic bright red line at which point you go from being a physical chemical system to having goals and memories and preferences and so on. This is something that has to develop slowly and gradually. So I'm interested in continuous views of this process.

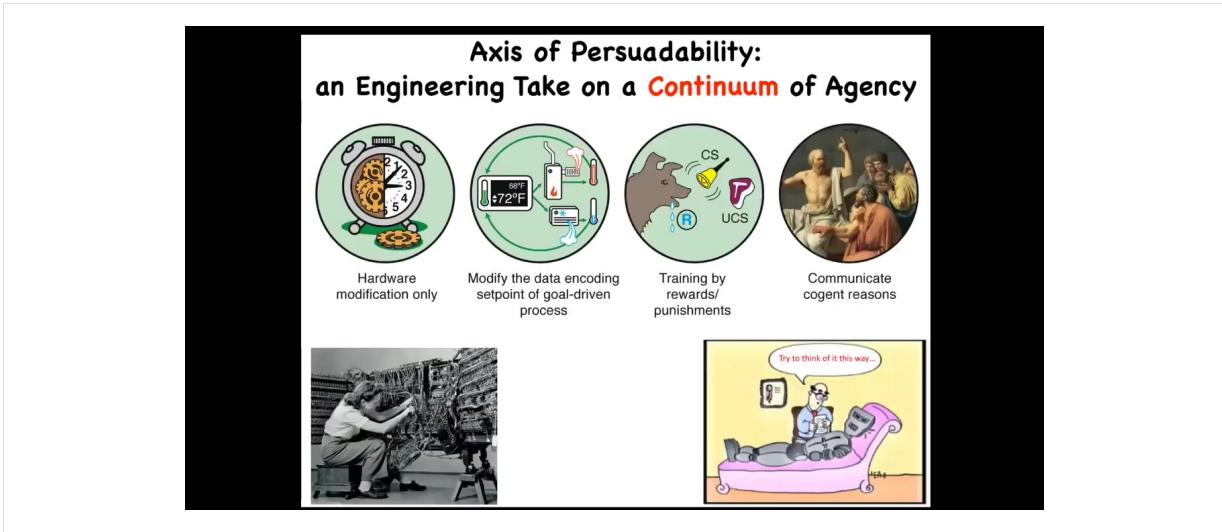
In particular, if we take developmental biology, evolution, and in fact bioengineering and augmentation seriously, we realize that it's a continuum in many ways, that there is no such thing as a standard, pure cognitive system, such as, for example, a human that people study. So that means we need frameworks for what is now called diverse intelligence.

I work in this field, and I'm trying to develop a framework that will allow us to recognize, create, and relate to truly diverse intelligences. This means not just primates and birds, and maybe an octopus or a whale, but also strange creatures, colonial organisms, swarms, synthetic organisms that have never existed before, artificial intelligence, both embodied and software, and maybe even someday exobiological agents. After all, the Earth is an N of 1 example of a lineage. Maybe there are very alien kinds of life out there.

So what I'm interested in is a framework that allows us to think about all of these kinds of things together. Of course, I'm not the first person to try for this. For example, here is an attempt by Rosenbluth, Wiener, and Bigelow back in the 40s to come up with some general waypoints along a continuum all the way from passive matter, using cybernetic notions, all the way up to the kinds of things that humans enjoy.

Our work in our lab spans from very basic philosophy of mind to conceptual frameworks to software implementations and benchwork to try to test these ideas and develop applications.

I really like it when these philosophical ideas give us new ways to move experimental work forward. I will show you some of those.



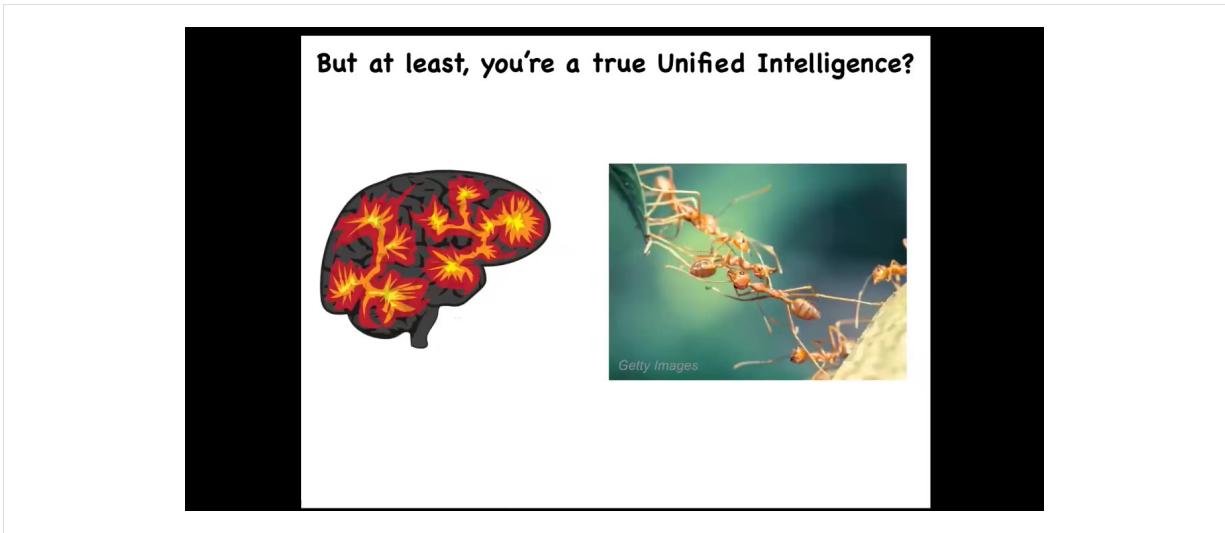
If we think about something I call the axis of persuadability, this is a continuum all the way from very simple machines all the way through to human and up. What we find is that different kinds of systems on this continuum are amenable to different kinds of interactions. Here, it's just rewiring. There's nothing else. You're not going to convince or reward a clockwork or anything like that.

Then you have some cybernetic systems where you can reset the set point, and you don't need to know all the details of how they work, but if you reset the set point, then they will do interesting things. Then we move up to systems that can learn, various types of reinforcement strategies, inference, and eventually very complex systems that can be motivated by very low energy stimuli that encode reasons to do things.

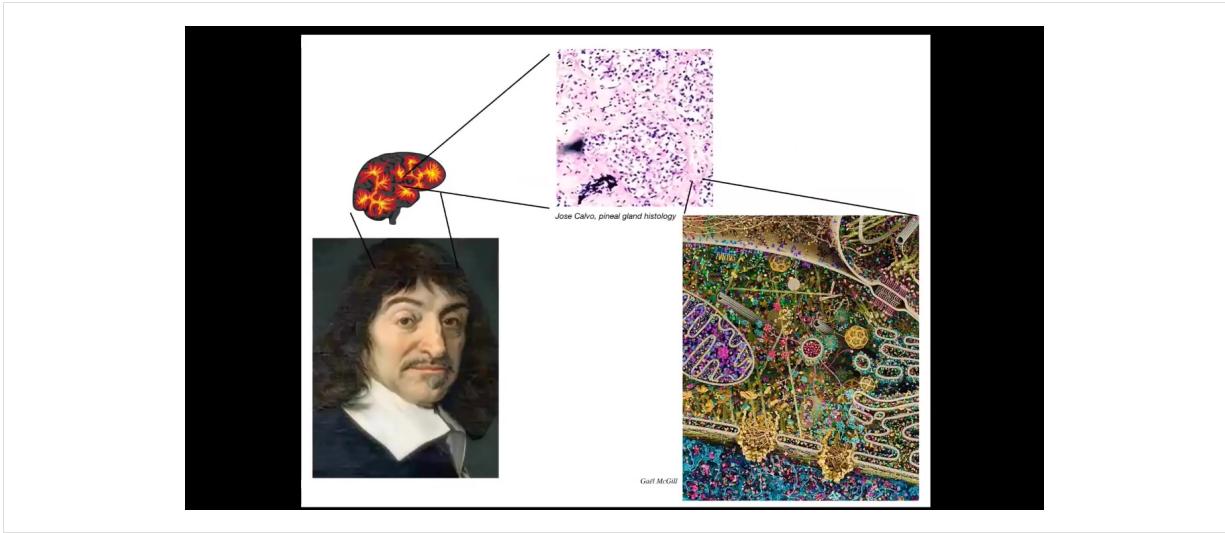
There's this whole continuum that is amenable to different tools. We use different tools down here than we use up here. That means that we should try to get it right. For any novel system, we should try to figure out where it lands on this continuum.

A lot of folks have assumptions, philosophical commitments to things like, they will look at some cells and say, "there's no way that cells are anywhere over here." Cells have to be simple physical, chemical machines. I think that is an empirical question. It's not something you can have armchair commitments to. You have to make a hypothesis about the problem space that these operate in and what kind of tools are going to let you have optimal control. You can do experiments and we can find out. There are different techniques for doing that.

Given that we arise from a single cell through this process of complex development, some people say we are a true unified intelligence.



We're not really a collective intelligence like ants and beehives. Maybe those can be called a collective intelligence, but they're not like us. We're a true unified intelligence.



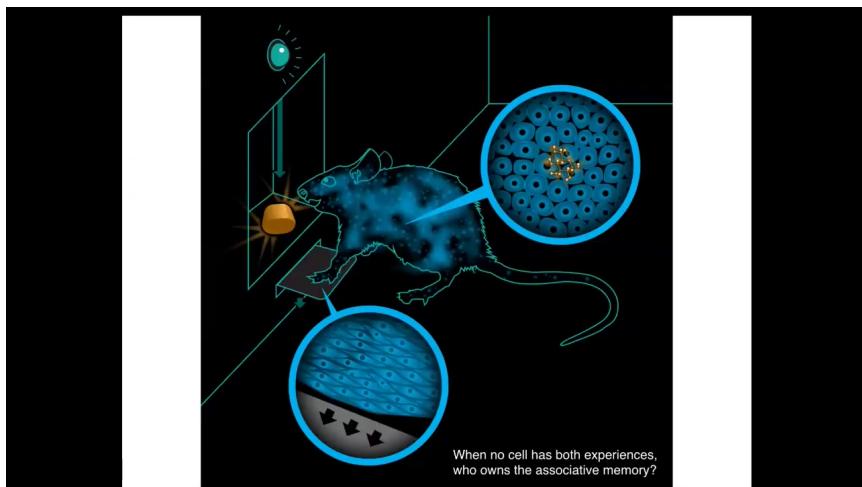
I want to stress the fact that this is not a good way of thinking about it. For example, René Descartes liked the pineal gland, because there was one of them in the brain, and he felt that the unified human experience should have a single kind of point of contact in the brain.

But if he had access to good microscopy, he would see that there isn't one of anything. This pineal gland, in fact, is made up of many cells. Each one of these cells has all of this stuff inside of it. It becomes clear that we are all collective intelligences.



We are all made of parts. There is no such thing, to my knowledge, as an indivisible diamond of intelligence. We all have parts.

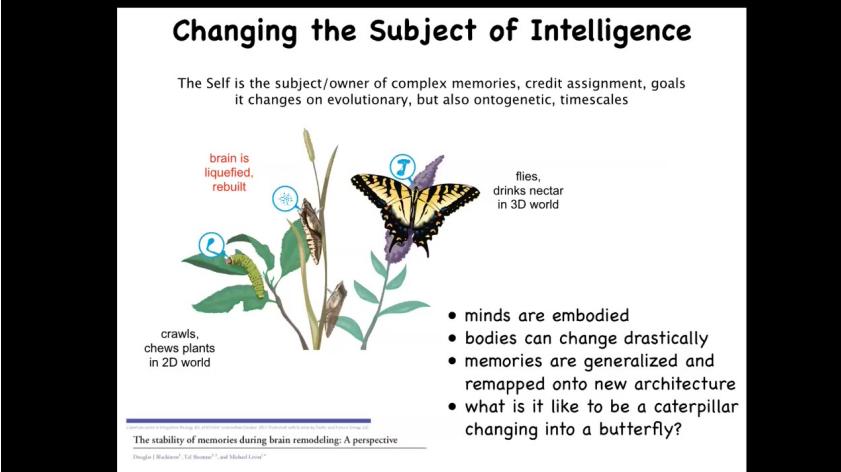
This is the thing that we are made of. This is a single cell, and this happens to be a lacrimaria. It's a free-living organism. It has no brain. It has no nervous system. It handles all of its anatomical, metabolic, physiological, and behavioral needs at the one cell level. So here it is. And it's a rather amazing little creature. This is the material that we are made of. We're not made of passive matter. We're made of this kind of agential material that used to be independent organisms and have all sorts of competencies and agendas.



One interesting thing that happens is that when you train a rat, when I press a lever, I get a reward. You might ask, who is the owner of this amazing associative memory that it formed? The cells at the bottom of the feet were the ones that interacted with the lever. The cells in the gut are the ones that received the sugar reward. Who owns this association between this and this? Here you need to start thinking about the whole animal. It's the rat. There is no single cell that had both experiences. No individual cell had both experiences in a way that would allow it to form this association. Now we have to bind these cells in some functional way so that you can have this emergent creature, this rat, which knows things that none of the individual components know. This is what we want to talk about. We want to talk about that cognitive glue that enables the subunits to form emergent beings that have memories and preferences and behavioral repertoires in different problem spaces.

## Changing the Subject of Intelligence

The Self is the subject/owner of complex memories, credit assignment, goals  
it changes on evolutionary, but also ontogenetic, timescales



The diagram shows a caterpillar crawling on a leaf, labeled 'crawls, chews plants in 2D world'. Above it, a brain is shown with a thought bubble containing a caterpillar, labeled 'brain is liquefied, rebuilt'. To the right, a butterfly is shown flying and drinking nectar, labeled 'flies, drinks nectar in 3D world'. A thought bubble above the butterfly contains a butterfly, labeled 'T'.

The stability of memories during brain remodeling: A perspective  
Douglas J. Wiklund<sup>1</sup>, Adi Shamir<sup>2</sup>, and Michael Lach<sup>3</sup>

- minds are embodied
- bodies can change drastically
- memories are generalized and remapped onto new architecture
- what is it like to be a caterpillar changing into a butterfly?

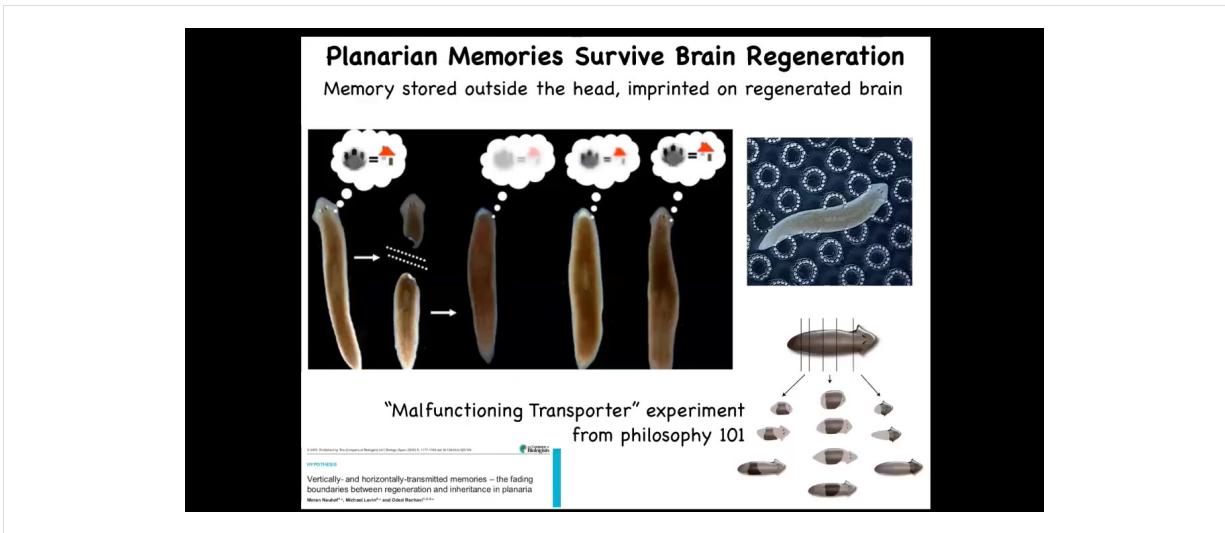
I want to talk briefly about a couple of unusual examples, just to warm up, because many people spend their lives studying mammalian brains. I want to talk about some unusual scenarios that can help us think in broader directions.

First of all, there's this idea that we have a significant metamorphosis in certain creatures that helps us to think about how behavior and bodies come together and how they intersect. This is a caterpillar. A caterpillar is a soft-bodied creature that requires a specific controller when there are no hard elements on which to push. So running a caterpillar-type body is a specific task. It has a brain that's well-suited for that process. It eats leaves and lives in a two-dimensional world.

What it has to turn into is a butterfly that's a hard-bodied creature, a completely different controller. It lives in a three-dimensional world, flies around and drinks nectar, and doesn't care anything about leaves. During this process, the brain is mostly taken apart. Many of the cells are killed. Most of the connections are broken. Radical reorganization of the brain to get from here to there.

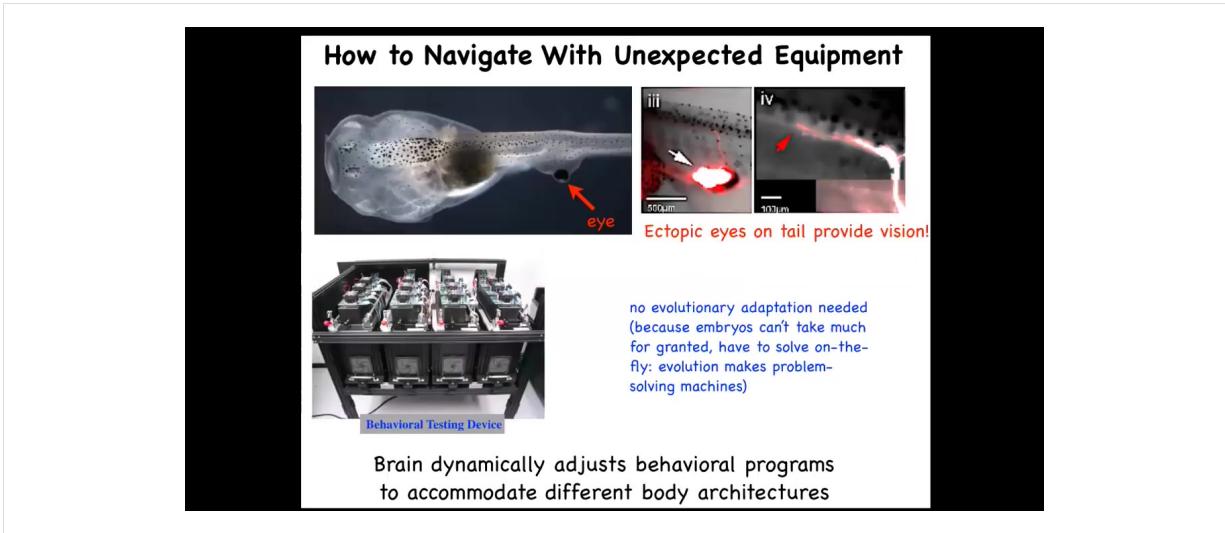
If you test butterflies and moths, they still have the memories that you trained the original caterpillar on. This work has been done and it's reviewed here. They keep the memories. Not only does memory persist through this remarkable reorganization of the substrate; when I talk about this to computer science audiences, it is a huge challenge. We don't have any substrate for computation that could do this.

Not only have the memories persisted, but, in an important way, they're generalized. Because when you've trained the caterpillar to find leaves on a particular color disc, when the butterfly goes there, it's not because it remembers the leaves. It doesn't care about leaves. It doesn't want leaves. What it's generalized is the concept of food. So both leaves and nectar are now food. There is this generalization that happens. All of that happens through a complete reorganization of the brain and the body.



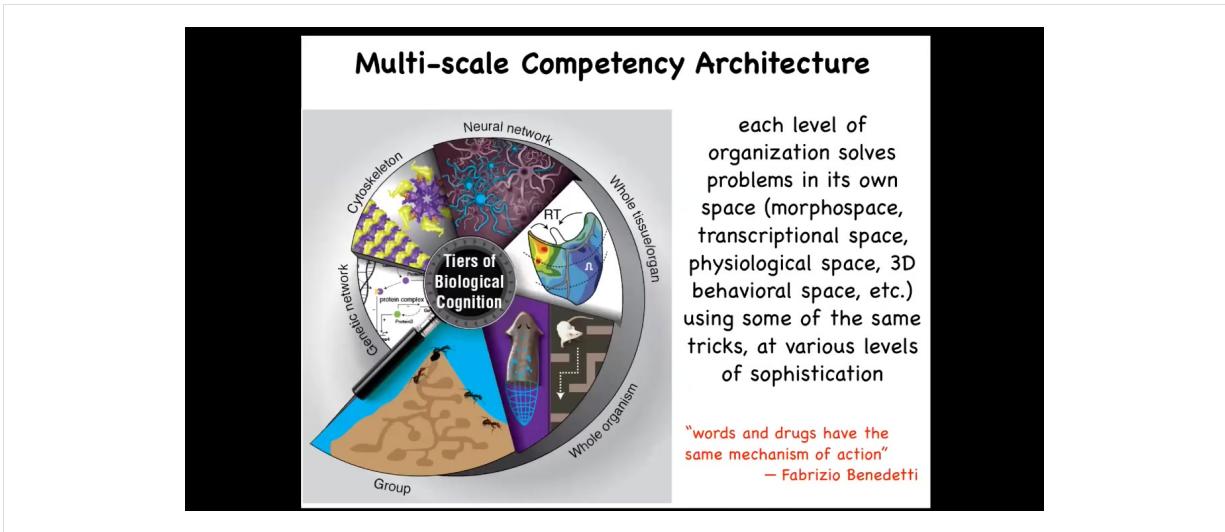
It gets even more impressive in animals like this. These are planaria. These are flatworms that regenerate every part of their body, including their brain. What McConnell discovered in the 60s, and then we later showed using modern techniques about 10 years ago, is that if you train them to find food on this particular little substrate, this bumpy substrate, it's place conditioning, you can remove their heads. The tail end here sits and does nothing. It has no behavior until the brain is formed. They form a centralized brain and their head and everything is reformed. At that point, you can show that they've retained the original information. So here you have information, presumably encoded somewhere outside of their brain, imprinted onto the new brain as it develops, so that it can run behavior. All of that is discussed here.

That brings up philosophical questions: if you chop them into pieces and these pieces keep the original information, then have you really created multiple individuals which are in a sense identical to the original? This is kind of the old malfunctioning transporter experiment from philosophy of mind 101.



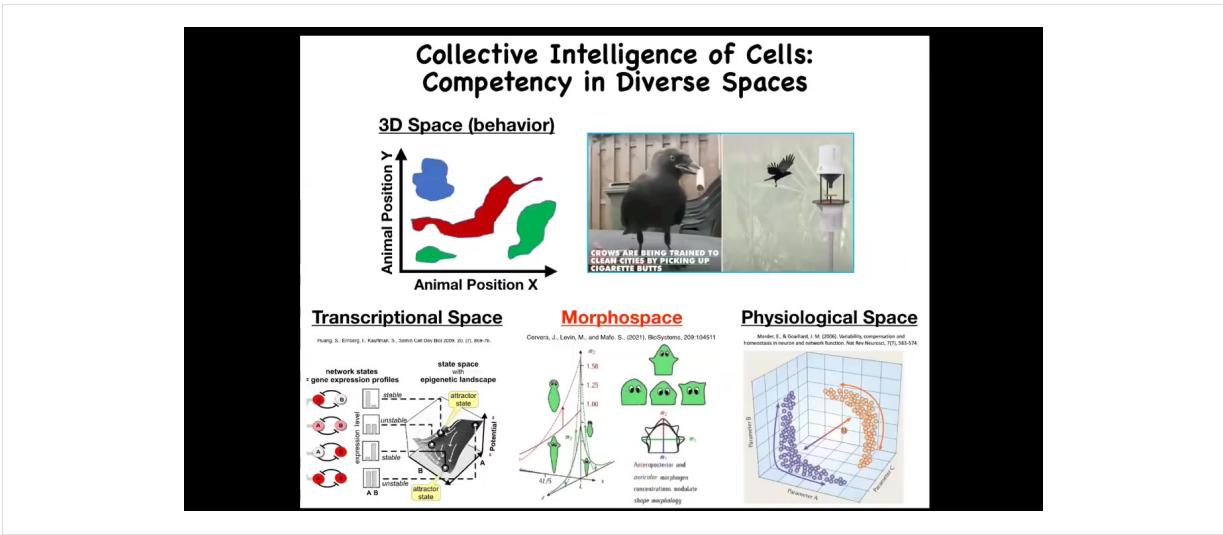
I want to show you another example of cognitive plasticity, this time in a vertebrate model. This is a tadpole. Here's the brain, here are the nostrils, the mouth. What you'll find is that there are no eyes here. We've prevented the eyes from forming, but what we've done is put an eye on its tail. I'll show you in a few minutes how we do that. When you put an eye on its tail, it makes one optic nerve, that optic nerve comes out. It does not go to the brain. It synapses sometimes on the spinal cord here, sometimes in the gut, sometimes out here somewhere. In these cases, the animals can see perfectly well. We've created a machine that tests them on visual assays, visual learning cues. It turns out they perform quite well. There is no evolutionary adaptation needed. We create these animals. Their sensory motor architecture is completely novel. Instead of having the eyes here where they belong connected to the brain, there's this weird patch of tissue on their back that is light sensitive and feeds information somewhere into the central bus. Then this brain can use that in its behavioral repertoires, no problem.

This plasticity and the ability of the body to change both in the developmental time scale and an evolutionary time scale is something that's really important and we really need to understand.



We think that it is intimately tied to our fundamental multi-scale competency architecture. This is the idea that our body, from the lowest molecular level up through the level of multiple individuals and various swarms, every layer is competent in its own space. In fact, some may be in multiple spaces. They all solve problems. They're not merely structural nested dolls; they solve interesting problems.

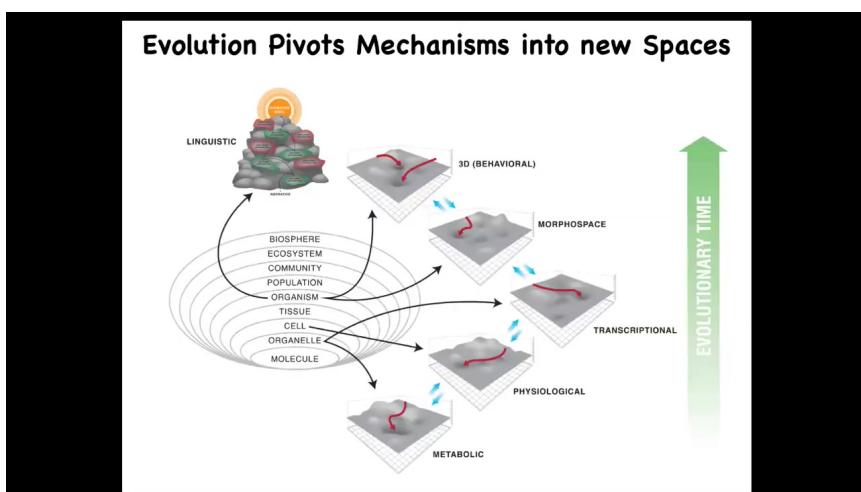
In particular, I really like this quote from Fabrizio Benedetti, who studies the molecular mechanisms of the placebo and nocebo effects. He says, "Words and drugs have the same mechanism of action." The idea is that this multi-scale architecture—these levels are not independent. They are functionally linked. Things that happen at the high cognitive level eventually filter down to the molecular level and vice versa. There's a constant flow of information up and down.



When I say intelligence in various spaces, what I'm referring to is the idea that we as humans are pretty good at detecting intelligence in medium-sized objects moving at medium speeds through three-dimensional space: birds and mammals. That's because most of our sense organs face outwards and we have a very characteristic spatiotemporal scale on which we are able to perceive intelligent behavior.

Now, there are other spaces: the space of gene expression, the space of physiological states, and anatomical morphospace, which is the space of possible anatomical configurations. You can think that if we had evolved with a sense organ facing inwards, able to test 20 different parameters of our blood chemistry, we would have evolved with a primary intuitive feeling that we live in a 23-dimensional space, that our liver and our kidneys are intelligent agents navigating that space in a way to solve various problems. This is very alien thinking to us because we do not perceive any of these spaces directly. Our various body components, in fact, solve problems in these various spaces.

For example, even molecular networks, something as simple as a deterministic gene regulatory network or a chemical pathway, can have six different kinds of learning, including Pavlovian conditioning. Even in the lowest forms of the body's components, some of the components that are tractable to techniques of behavioral science and neuroscience are there with us.



One hypothesis that we've been working on is this idea that what evolution basically does is pivot some of the same tricks, and there's this fundamental trick about navigating various spaces, but it pivots this through different spaces, maybe making some specializations, but basically all the way from metabolic spaces at the beginning of life all the way through to eventually linguistic spaces and the need to keep a straight narrative in a complex linguistic space. This traditional behavior is what most people in the field study. What we just talked about is this idea that we really need to broaden some of our ideas around what implementations might support different kinds of intelligence and this plasticity that is underlined by the fact that all of the components are actually solving problems as well.

What I want to do next is talk about this example of a really unconventional agent. This is a collective intelligence that lives in morphospace. Just very briefly, I'll mention evolution, but we could talk for a long time about that. This is going to be our example of how do we use techniques and concepts from both bench science and neural and behavioral science to understand really a field that is mostly considered to be quite a different field.

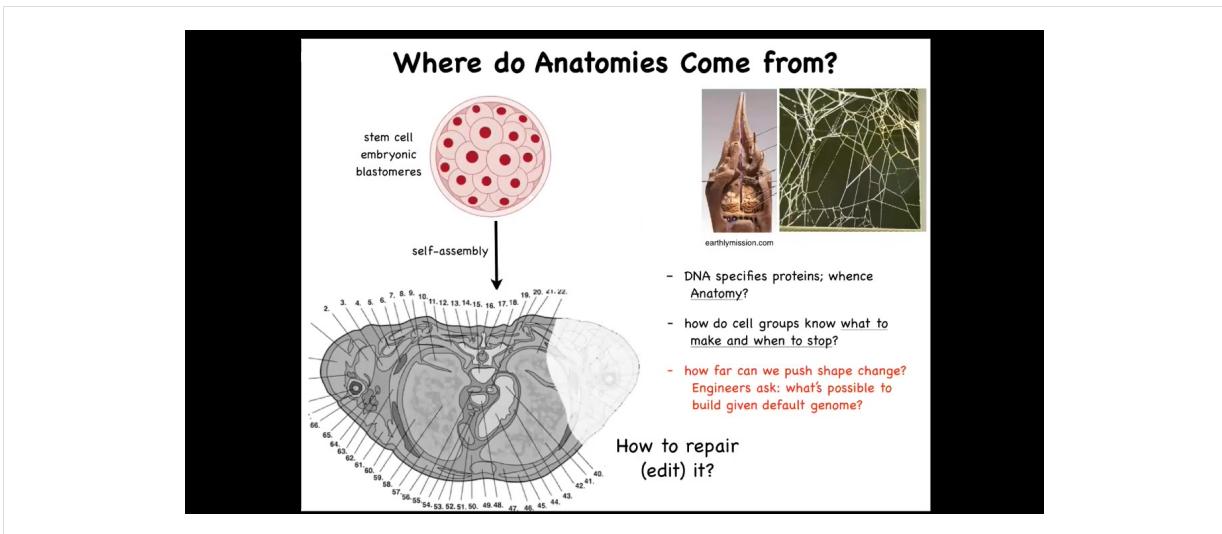
Alan Turing needs no introduction, but he was very interested in minds broadly conceived. He was interested in intelligence, in fact, in minds, in unconventional substrates and computation, and what is the essential element of being a mind and being able to solve problems. He specifically was interested in problem-solving machines and intelligence through reprogrammability.

It turns out that he also wrote this paper, which is really interesting. He wrote a paper on the chemical basis of morphogenesis, an early mathematical model of the kinds of processes that might organize a living body from scratch. This idea of symmetry breaking, emergence of complexity, and regulation of patterns.

You might wonder why somebody who is interested in computer science and fundamental mathematics would be looking at models of diffusion. I think it's because he saw something very profound, which is that there is a really strong symmetry between these two fields. This idea that the processes by which the body comes to be are an important mirror of the processes by which minds come to be. The idea that understanding intelligence and understanding the scaling of cognition in the universe may be quite similar to how we understand where bodies actually come from. I think that's very profound.

He didn't talk about this directly as far as I know in any of his writings. He didn't live to explore it, but I think he was on to something.

Slide 14 of 47 · Watch at [21:53](#)



Let's ask ourselves, where do bodies come from? This is a cross-section through a human torso, quite reliable in most births. You can see this incredible complexity, this order, everything is in the right place, the right shape and size relative to the right stuff. You might wonder, where is this pattern encoded? Where does this pattern come from? If it's so reliable, where is it? Of course, we all start life as a bunch of embryonic blastomeres, and eventually you get this.

A lot of people will say it's in the DNA. It's in the human genome, but we can read genomes now. We know what's in the genome, and what's in the genome is nothing directly related to this. It doesn't specify anything like this. The genome specifies the

micro-level hardware that every cell gets to have. The proteins are what the genome specifies. Now we need to understand how this works.

We now know the genome doesn't directly specify this any more than the termite genome specifies directly the shape of the termite mound or the spider genome specifies the geometry of the web that it's going to build. These are all emergent features from the computations, the physiological computations that hardware performs.

Now we need to understand how these cells know what to make? How do they know when to stop? In particular, what else can these cells build? How far can you push them? Could you make them build something completely different? Regenerative medicine asks, what would you have to do to repair it? In other words, how do you convince cells if something is missing or damaged? How do you convince cells to go back to the normal pattern?

**Biomedical Endgame: Anatomical Compiler**

**Why we need it:**

- Birth defects
- Traumatic injury
- Cancer
- Aging

Problems of information processing

how to control emergence and scaling of cellular activity toward desirable complex, system-level outcomes?

this is *NOT* a 3D printer - it is a communications device (translator)

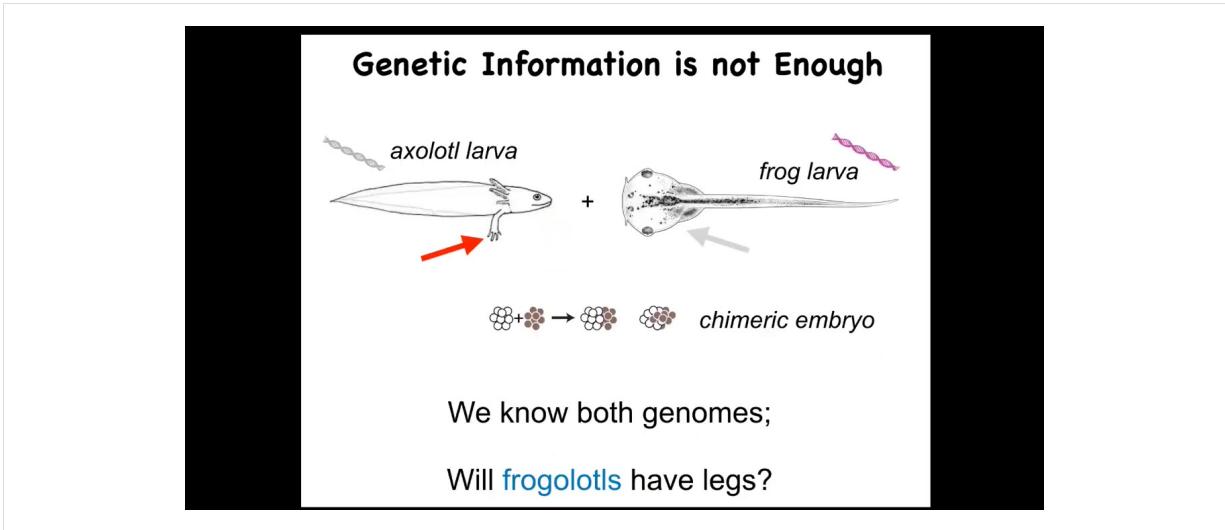
Philosophical question: if we had it, would it mean we *understood* totally?

We're very interested in this thing, which I call the anatomical compiler. I thought of this when I was thinking about the end game of our whole field, what are we going for such that when it's done, we can all go home. The idea is that at some point in the future, you'll be able to sit down in front of a computer system and draw the plant or animal that you want, or the organ or the weird biological robot. It doesn't matter.

If we knew what we were doing, the compiler would convert that anatomical description into a set of stimuli that would have to be given to cells to cause them to build that particular structure.

Why do we need this? First of all, most biomedical problems would go away. Everything related to birth defects, traumatic injury, cancer, aging, degenerative diseases, all of that would go away if we had the ability to tell a collection of cells what to build, because then you could build and rebuild organs and other structures at will.

The important thing about this is that this is not meant to be a 3D printer. The point isn't to micromanage and put cells where you want them to be. It's more of a communications device. It's a translator between your anatomical goals and the set points that guide cell behavior. Why don't we have this? Why are we so far away from this?

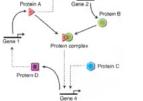


I want to point out a very simple example. This is a baby axolotl. Axolotls have little four legs at this stage. Here's a frog larva. They do not have legs at this stage. We make something in my lab called a frogolotl. This is just a chimera of an axolotl and a Xenopus embryo. Both of these are sequenced. You've got the axolotl genome, you've got the frog genome. Now I ask a simple question: are frogolotls going to have legs or not? The fact is we have no idea. We can look empirically, but nobody can make a prediction. There are no models, despite the availability of genetic information, that will make a prediction on an experiment like that. In fact, the genomic data can't even be used to guess what the actual shape of the organism is going to be unless you cheat by comparing that genome to something else that you already know.

### The State of the Art: biological information

We are very good at manipulating molecules and cells;

We are a long way from control of large-scale form and function







biomedicine has been largely at the hardware level of analysis and control, we do not understand higher-level information processing in biology (intelligence)

Here's where we are today. The field is very good at manipulating cells and molecules so we can get information like this. We're really quite a long way away from large-scale form and function control. I'm going to argue that this is because we are currently where computer science was in the 40s and 50s.

This is what programming looked like in the 40s and 50s. You literally had to interact with the hardware. You had to physically rewire it. But the reason that nowadays you don't get out your soldering iron and start rewiring your laptop when you want to go from "Microsoft Word" to "PowerPoint" is because we have figured out that there's a ton of power in high-level interventions that reprogram and take advantage of a certain kind of hardware.

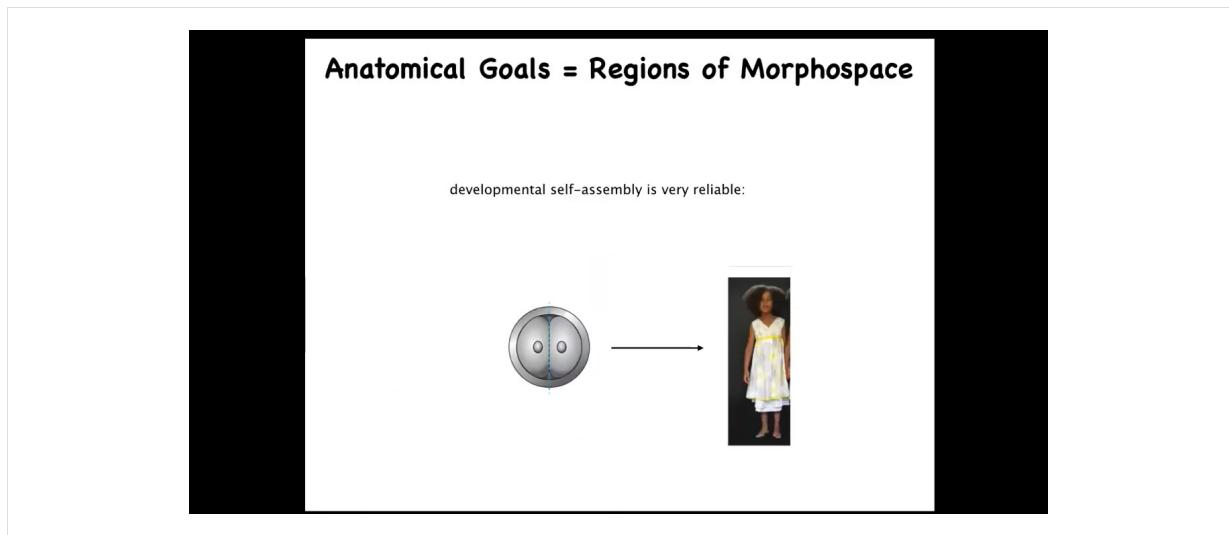
I'm going to argue that biology is exactly that kind of hardware, which, for neurosciences, is not exactly news. For molecular and developmental biology, this is a different way of thinking. This idea is that we've only begun to scratch the surface of what's actually possible. We have not really understood the amount and the type and the tractability of information processing and intelligence or problem solving at different levels of biology.

“Intelligence is the ability to reach the same goal by different means.”

- William James

*what kind of collective intelligence do cellular swarms deploy?*

And what do I mean by intelligence? I like William James's definition. William James said that intelligence is the ability to reach the same goal by different means. It's a cybernetic definition. It doesn't talk about brains or whether you were engineered or evolved or anything like that. It talks about the difference between two magnets trying to get together, which if faced with a barrier, will simply stand there pressed up against the barrier. They're unable to walk around, temporarily get further from their goal in order to finally reach it. They have no delayed gratification ability, a very simple system. Romeo and Juliet, on the other hand, have all kinds of complex abilities, including planning, to try to get their goals met. In between, you have all sorts of things, cells and tissues and self-driving vehicles and animals of different types. This ability to reach the same goal, despite novelty, is a kind of measure of this intelligence. It's important because these competencies, these capabilities, can be taken advantage of in therapeutic contexts if we understood how to deal with them. What kind of collective intelligences do cellular swarms have? I'm going to show you a few examples.



One is that while development is very reliable, almost always you start with a human egg and you end up with a human organism. It's not hardwired because you can take early embryos, cut them in half or in many pieces, and you don't get half bodies, you get perfectly normal monozygotic twins. There's this ability to navigate to the ensemble of goal states corresponding to some kind of normal human target morphology from different states.

I'm collapsing a very complex anatomical morphospace onto just two dimensions to show how they all navigate to that set point, avoiding various local maxima along the way.

That kind of regulative development is not just for embryogenesis.

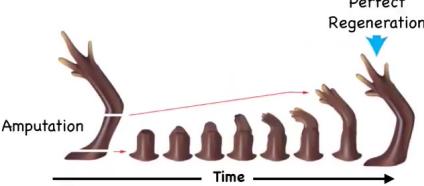
**Same anatomy, from different starting states**

- get to the same outcome
  - despite perturbations
  - from diverse starting positions
  - via different paths



**Anatomical homeostasis:**

it stops when the correct large-scale setpoint (target morphology) has been reached



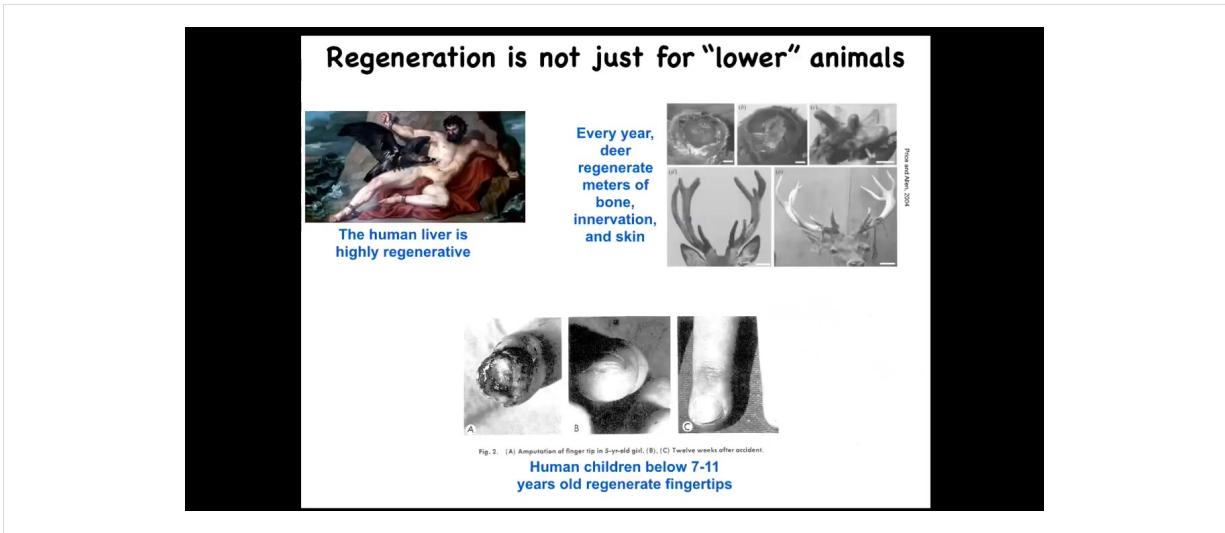
Amputation

Time

Perfect Regeneration

The diagram shows a sequence of limb stumps over time, with a red arrow pointing to the first stump labeled 'Amputation'. The stumps increase in size and complexity. A blue arrow points to the final regenerated limb labeled 'Perfect Regeneration'. A horizontal arrow below the stumps is labeled 'Time'.

It also works in some animals. It works during their whole lifespan. This is a salamander. They regenerate their limbs, their eyes, their jaws, their ovaries, their spinal cords, and so on. What's interesting about this process is that if you cut it anywhere along, for example, this limb, they grow exactly what's needed, no more, no less, and then they stop. The most amazing thing about regeneration is that it knows when to stop. It stops when a correct salamander arm has formed. This is a kind of anatomical homeostasis. You can start off at any configuration and it will keep going in the right way until you get to the right thing.



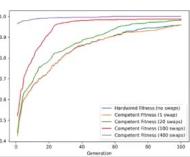
Humans can do this a little bit. This isn't just for worms and salamanders. Humans regenerate their liver. Even the ancient Greeks knew that. Human children will regenerate their fingertips up to a certain age. Deer regenerate huge amounts of bone, vasculature, innervation, and skin every year. Up to a centimeter and a half of new bone growth per year—per day when they're growing back their racks.

### Evolutionary Implications

- Hierarchical control = search the nicer space of inducements, rewards, signals, incentives, etc. not the chaotic space of microstates
- Competencies of modules smooths the evolutionary landscape and buffers negative effects of mutation, giving the process patience (improves evolvability)
- Competency hides genetic information from selection, resulting in intelligence ratchet: further improvements by raising the IQ of the agential material, vs. micromanaging the hardware (genome)



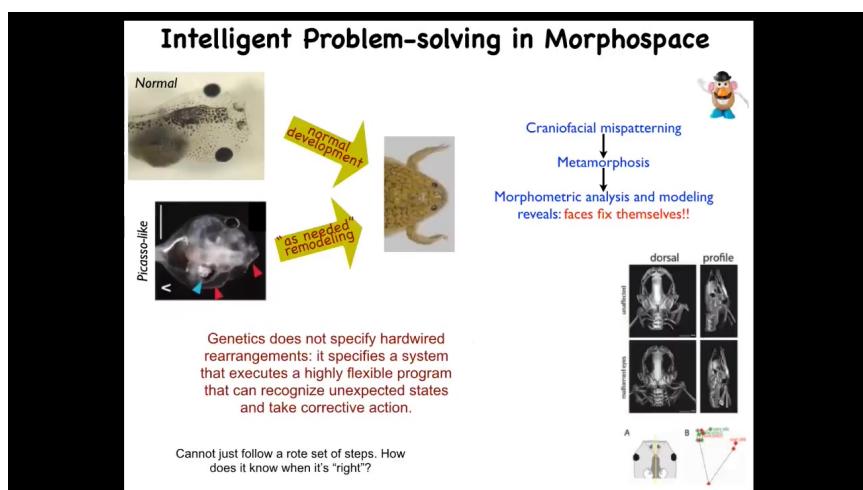




This idea that what evolution is working with is a kind of a competent material where cells will correct for all kinds of defects results in a very interesting intelligence ratchet, because that competency makes it hard for selection to see whether the fitness of the animal comes from the quality of the genome or from the fact that the actual structural genome wasn't great, but the competency made up for it. And the more that happens, the more effort from evolution actually goes into ratcheting the competency mechanism. So it's a very interesting thing. We can talk about that afterwards.

So now that we've seen that collectives of cells can solve various problems, remember, the problem that they solve is how to walk a particular path through anatomical amorphous space that starts out as a single cell, for example, or a healthy organ in the case of regeneration, and then is deviated from that to a different location and then knows how to come back despite various perturbations. There are remarkable examples of this; they're all discussed in our papers.

Now that we see that ability to navigate morphospace is the thing we want to explain, there's an obvious thing to look for, which is: what is the cognitive glue? How are the cells all bound together to the same construction project? No individual cell knows what a finger is or how many fingers you're supposed to have, but the collective clearly does. We need to understand how the individual cells are bound. So we're going to talk about bioelectricity.



Now, just to remind you of a simple experiment that forces that kind of question: something that we discovered a few years ago. So here's a tadpole, normal: brain, nostrils, mouth right here. These tadpoles have to become frogs. In order to become frogs, they have to rearrange their face. So everything has to move. The eyes have to move, the jaws, the nostrils, everything has to move around. It was thought that this is a very hardwired, meaning low intelligence process because all tadpoles look the same, all frogs look the same. So as long as you somehow encode specific movements, in other words, direction and amount for each of these organs, you will end up from a normal tadpole to a normal frog.

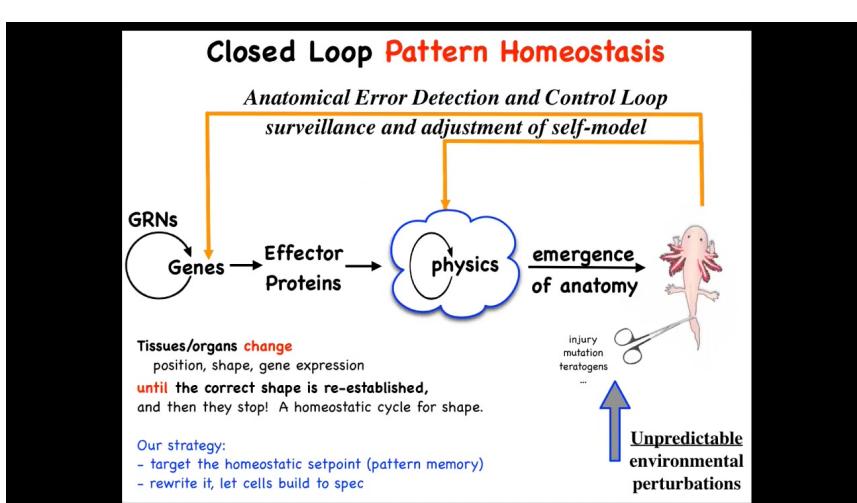
We decided to test that. The way you test the competency of these kinds of systems is by perturbative experiments. You cannot do it just by observation. You cannot do it by only modeling. You have to perturb them by, as William James said, putting various obstacles in their way and seeing what they do to respond. So what we did was we created so-called Picasso tadpoles. Everything is in the wrong place. The eyes on the back of the head, the mouth is off to the side, the whole thing is scrambled.

What we found out is that these animals actually end up being pretty normal frogs, because all of these organs do not just go in the same path that they do in normal animals. They will go in novel abnormal paths. In fact, sometimes they go too far and they have to double back until they get to the correct frog face and then they stop. So what the genetics actually gives you is an error minimization scheme, a system that can recognize that it's in the wrong state and take corrective action to get where it needs to go to complete that journey in morphospace despite being started in a very abnormal configuration.

Again, this configuration is at the large scale. So it's not that the individual eye cells are sprinkled throughout the whole head. The eye as a whole has to move, the mouth and

everything else. This is truly a collective decision-making problem. But it raises a fundamental question. How does it know what the correct pattern is? What is the set point? How does it know what the frog face is supposed to look like?

Slide 24 of 47 · Watch at [33:56](#)



You can think of it as, first of all, this is the standard kind of story that you're told in developmental biology textbooks, which is that there are gene regulatory networks. They make proteins. They interact with each other in various ways. There is this process of emergence, and it's a very complex process because lots of simple local rules are being carried out again and again, and out comes something complex like this. So this is a feed-forward, open-loop process. This certainly does happen.

There are many simple processes that give rise to complexity. This is true. But this is only part of the story. I don't think even the most important part. Because what happens here is something very interesting. If you deviate from this state by injury, mutation, teratogens, many other ways, what happens is that activity kicks in both at the level of physics and at the level of genetics that tries to reduce the error and get you back where you need to go. So this is not really an open loop process. This is a process of anatomic homeostasis.

On the one hand, this is maybe not so surprising. Of course, biologists know all about homeostatic loops. On the other hand, there's a couple of weird things here. One is that typical homeostatic loops have a scalar as a set point, a single number. So maybe hunger level or pH. In this case, the set point is basically a complex anatomical

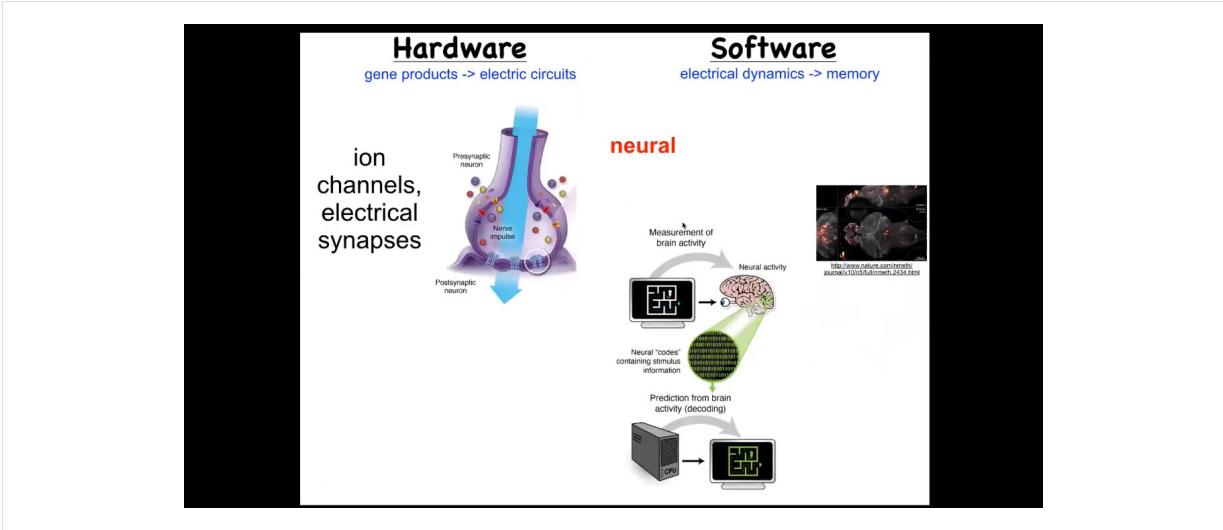
structure, not to the individual cell level, but really a fairly complex data structure here. That's first.

Second, certainly in behavioral science, you can talk about at least some kinds of animals having goals and doing things for reasons. But typically in cell and molecular biology, we are told to avoid that. You're not supposed to think about systems having any kind of a future goal; you're supposed to feed forward from chemistry and see what happens. This is a weird way of thinking about it, but it makes a very strong prediction.

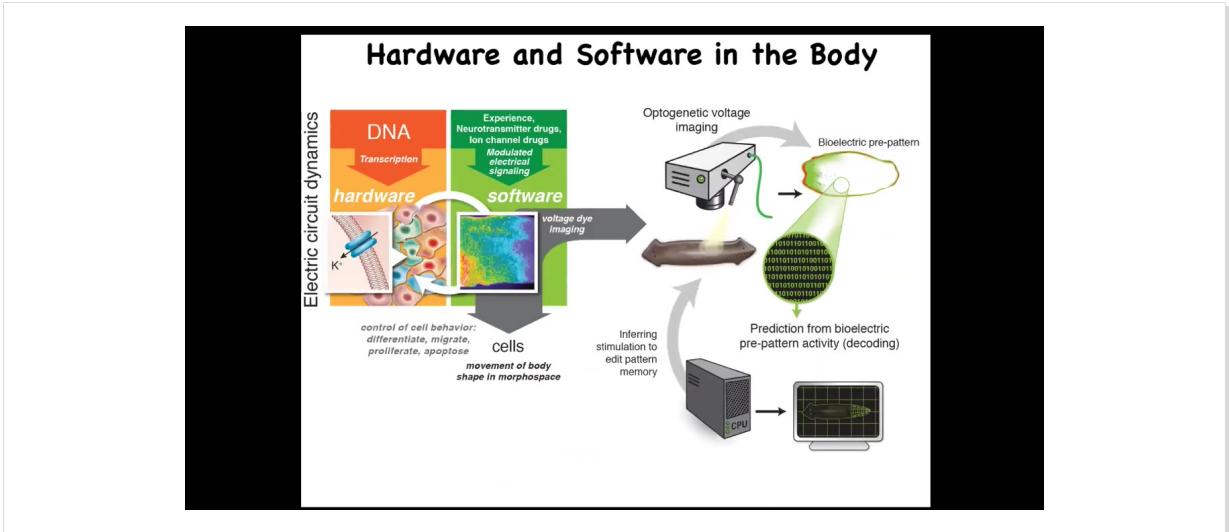
The prediction is that as a good homeostatic system, you ought to be able to make changes by rewriting the set point. In other words, you should be able to find the biophysical mechanism of the set point. How does the system know what the correct pattern is? You should be able to decode it. You should be able to rewrite it.

Those three things are important in applications because currently, the only game in town is down here. If you want to make changes up here, we are told that you'll have to work out what changes in the lowest level rules are going to give you system-level changes that you want. This process is fundamentally irreversible. It is a terrible inverse problem. There's no real way to reverse it.

So that limits the applicability of things like CRISPR and genomic editing, because, with the exception of some low-hanging fruit, it is completely unclear what genes you would edit to make large-scale changes here. Could we actually find a set point? This is where we've been looking.

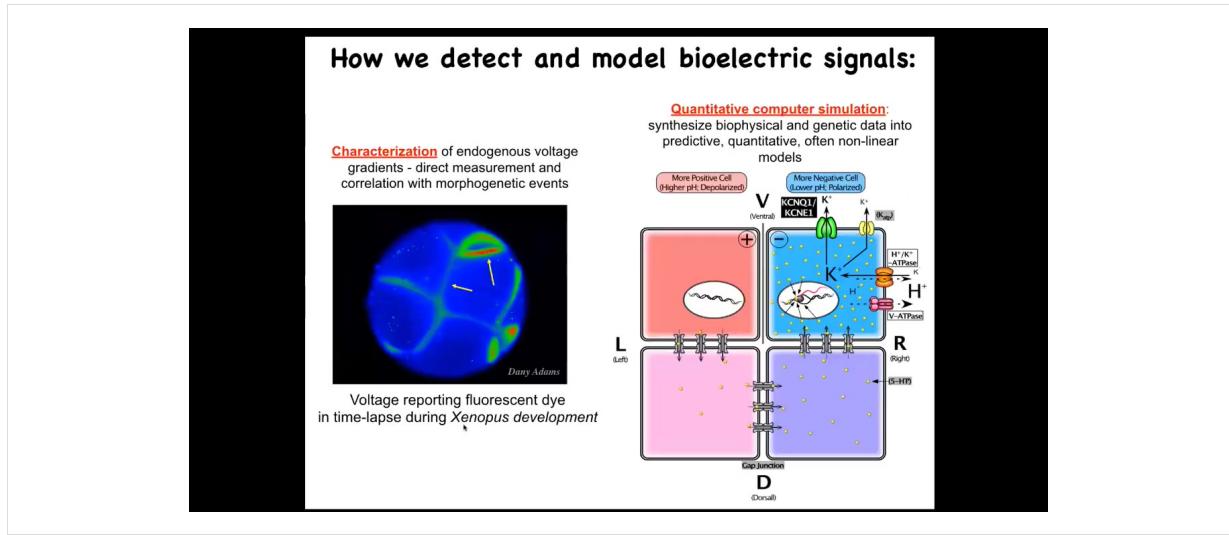


Now, in thinking about where to find the set point, we said, what's the one example, the one non-controversial example of a biological system that really encodes goals and the pursuit of goals? That occurs in neuroscience. We know what the hardware is. The hardware are these networks of cells that achieve a certain voltage potential by means of ions moving through the ion channels. They may or may not communicate that voltage to their neighbors through electrical synapses known as gap junctions. And so that's the hardware. The software here that's mediated by this can be read out. This group made this amazing video of a living zebrafish brain. You can see all the electrical signaling. It's the commitment of a lot of neuroscience that this is really the encoding of the cognitive content of this individual. There's this process of neural decoding that if we were to record all the electrical activity, and if we knew how to decode it, you would be able to read the animal's memories, preferences, goals, potential behaviors. That's the process of neural decoding. This physiology is where we think all of that is sitting. The amazing thing is that actually every cell in your body does this. Every cell in your body has ion channels. Most cells have gap junctions to their neighbors. From the earliest time of development, your tissues are an electrical network that is processing information. Now, we can then try the same kind of decoding research program. Could we decode the bioelectrical signals in non-neural tissues and ask what are they thinking about? What problems are they solving? In particular, much like here, we're tracking the voltage with a voltage reporter. This is an early frog embryo in time-lapse, reporting the voltage so we can see all the electrical conversations that these cells are having with each other as they sort out who's going to be head, tail, and so on. There's this basic kind of diagram where what the bioelectric networks in your nervous system are doing is controlling your muscles to move your body through three-dimensional space. This can be decoded and modeled.



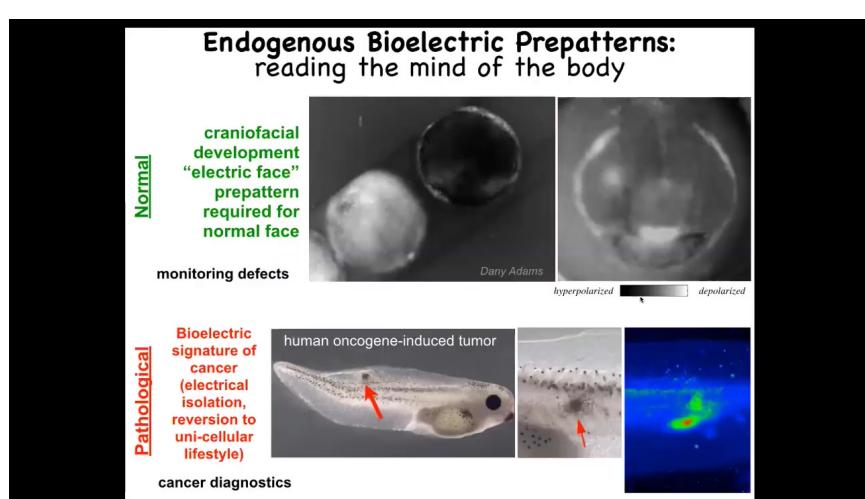
And my argument is that this amazing trick is really an elaboration of an evolutionarily much earlier system in which all of the cells were running an electrical network to move the configuration of your body through morphospace. And what evolution has done here is a couple of things. One is to greatly speed up the process. So we're talking milliseconds in neuroscience as opposed to minutes or hours in developmental bioelectricity. And it pivoted space for time, where this is mostly about the control of space. One can argue that nervous systems are mostly about the control of behavior as a function of time. Evolution caught on to this system as early as bacterial biofilms. This is incredibly ancient and it long predates nerves and muscles.

So we then went on to try to develop some of the first tools, the first molecular tools to read and write this cognitive information out of non-neural cells. No surprise, we took most things from neuroscience because as it turns out, the tools cannot tell the difference.



We try to maintain that difference through different journals and different funding bodies and different educational materials and courses and departments. But all of the tools, both the bench tools and the concepts, poured over very nicely. Some of that distinction might be quite artificial.

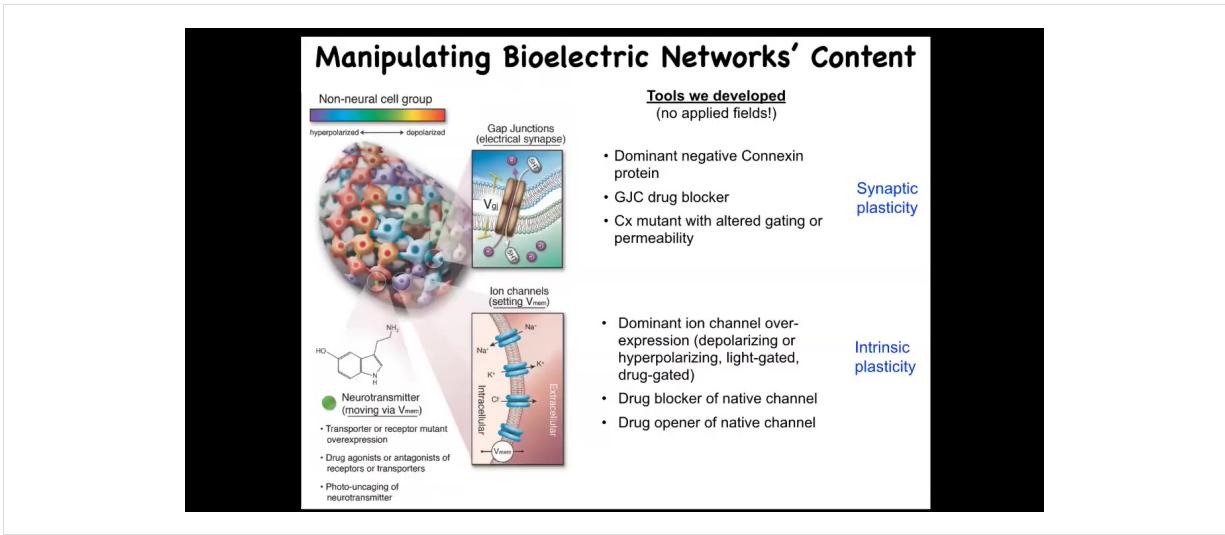
We developed some tools to non-invasively read the bioelectric pattern. We also do a lot of computational modeling to try to understand how these bioelectrical patterns result from the activities of channels and pumps and gap junctions and so on.



Let me show you a couple of such patterns. This is an early frog embryo putting its face together. You can see, this is a voltage-sensitive fluorescent dye in time-lapse. This is one frame out of that video. We call it "the electric face". You can see here that in that frame you can already see, here's the animal's right eye, here's the mouth, here are the placodes. The major structures of the face are already laid out in this subtle bioelectrical pre-pattern before the genes come on to regionalize that face and before the anatomy actually starts to form. As it turns out, this is absolutely an essential pattern memory that is required for forming a normal face. If you move this around, as I'll show you in a minute, the genes and the anatomy will faithfully follow. That's an endogenous pattern that you might use for monitoring defects.

We'll come back to cancer a little bit at the end, but here's a pathological pattern. When you inject a human oncogene, eventually you'll get a tumor. It will start to metastasize. But before that, you can already tell where it's going to be. You can tell where the regions are that are going to go off of the plan and revert back to an amoeba-like environment where they're unplugged electrically from the rest of the network. They're not receiving the signals about how to build various organs. They're just amoebas as far as they're concerned. The rest of the body is external environment.

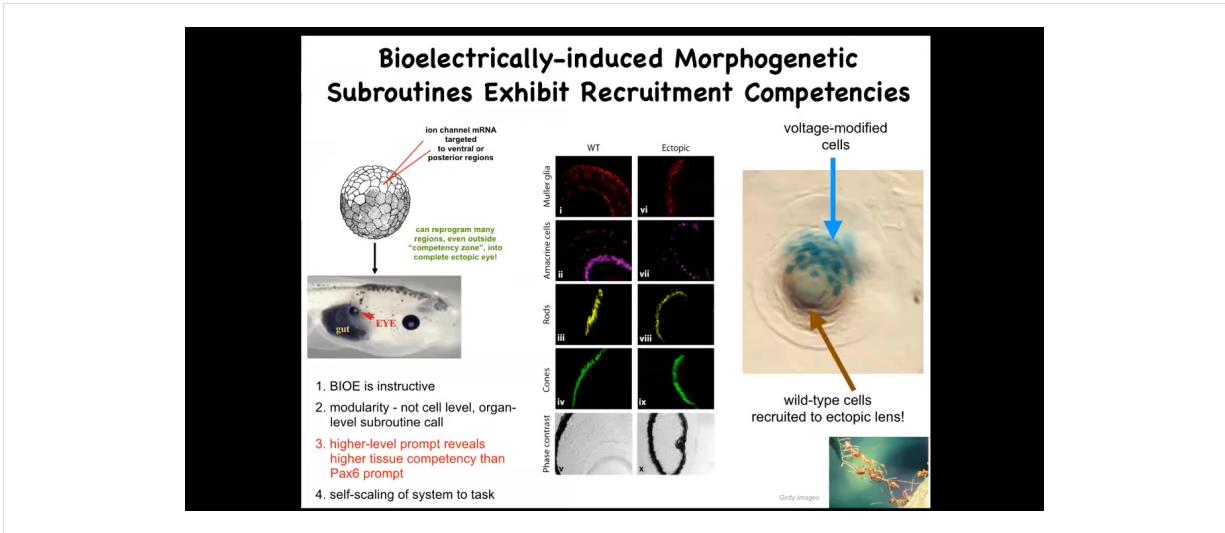
These are just some examples. I show you this because this is the easiest pattern to decode. We have many other patterns that are quite hard, but this is pretty obvious what you're looking at.



In addition to the imaging, the most important thing is functional. Can we manipulate these bioelectrical patterns and show that they're functionally critical for what the morphogenetic collective does next? I have to stress, we do not use any applied fields. There are no magnets, there are no frequencies, no electromagnetic waves, no electrodes.

We use all the tools of molecular neuroscience and physiology. We use pharmacology to open and close channels and gap junctions. We can mutate the channels to change their properties in a known way. We can use optogenetics.

The idea is to take any region in the animal and use these tools to alter the topology and the activity of the bioelectrical network in a very specific way to try to change its decision-making as it navigates the anatomical morpho space. Whereas the neural decoding was the part with the dyes that I just showed you a minute ago, this is equivalent to the inception of false memories, for example, and people have done this in mice. I want to show you what happens when you do this.



Remember that little eye spot that I showed you in the bioelectric face? You might ask, what happens if you introduce that same kind of spot somewhere else? You take a memory of eye formation and you falsely reproduce it in other regions of the body that are actually supposed to be doing something else. If we inject specific ion channels that can induce that same type of eye spot that I showed you a minute ago, and introduce it into cells that are precursors of things like the gut, those cells will happily build an eye.

These eyes can contain all of the right layers, the lens, optic nerve, retina. They all have them.

First, it tells you that the bioelectric pattern is instructive. This is not about messing up the pattern and making monsters. This is about showing that those patterns can cause specific organ-level changes in morphogenesis.

We're not telling individual cells to differentiate a different way. We're telling the whole region to make an eye. This is extremely modular. We don't have to provide all the information about how to make an eye. We don't know how to make an eye. This is a very high level subroutine call and this only works with high level competent systems because in very simple machines, you have to be in charge of every aspect of the functionality. You have to program every aspect of what it does. Here, we don't need to do that because the system has built-in competencies.

The interesting thing about these competencies is that when you're analyzing a system, any system, to make an assessment over what it can do, you're really taking an IQ test yourself because it is not obvious. If you don't prompt it in the right way, you don't know.

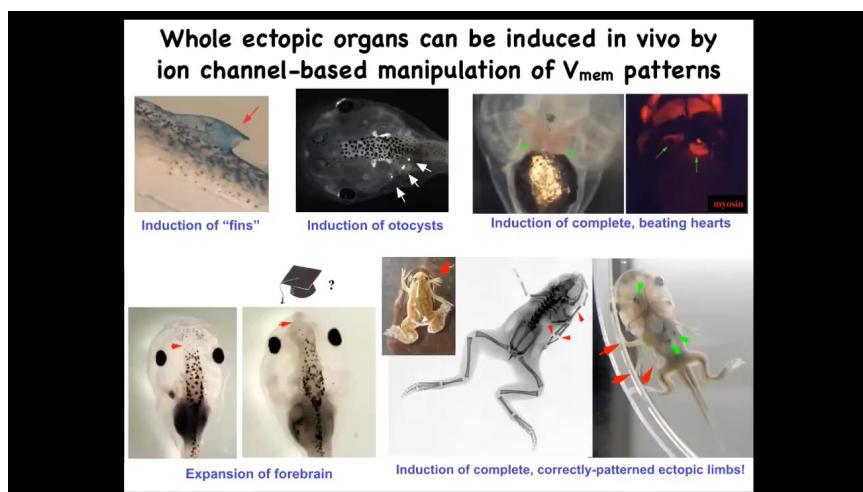
For example, the textbook will tell you that only the anterior neuroectoderm up here is competent to make an eye. The classic experiments were done with a master eye gene, PAX6. If you misexpress PAX6, only these cells are competent to make an eye. But that

isn't the actual best master regulator. If you do it with a voltage state, you find out that the competency is much more than you thought it was. Every cell in the body is able to participate in this, including endoderm. Endoderm is not supposed to make eyes.

It's important to be humble about when the limitations that we see are a result of our own lack of knowledge about what is a meaningful signal to the surrounding tissues.

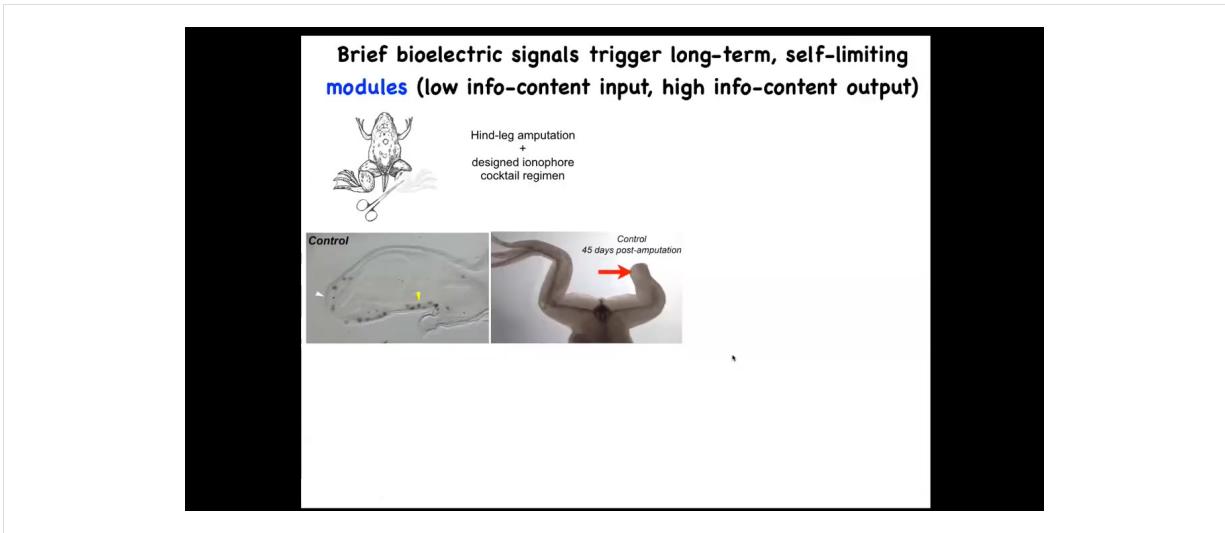
This is a lens sitting in the tail of a tadpole somewhere, an extra lens. The blue cells are the lineage label of the cells that we injected with our channel. All of this other stuff here was never manipulated by us. Why does it participate? These cells can do what many collective intelligences do, like ants and termites. When there's not enough of them to do the job, they recruit their neighbors.

There's a secondary instruction event taking place where these cells tell these other cells, you need to work with us to make this eye. We don't have to program that in both the eye induction and the scaling to the task. All of that is a native competency of these collectives if you know how to activate them.

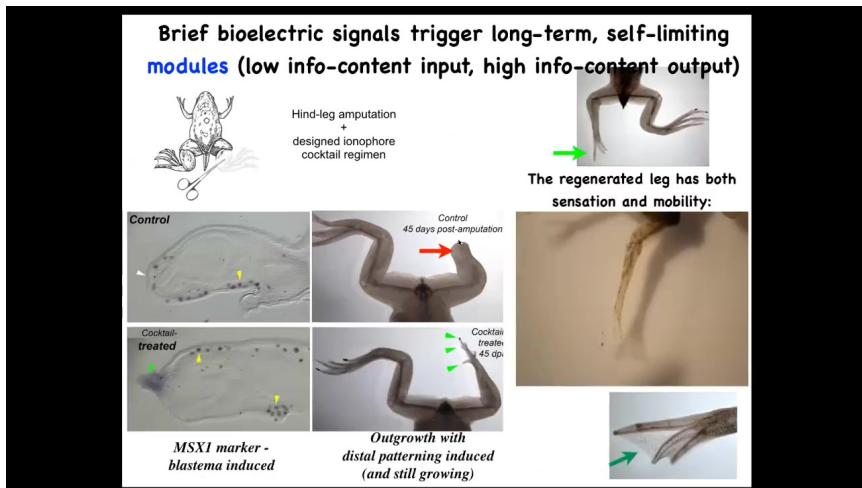


We can make many things. We can make ectopic forebrain, ectopic hearts, ectopic otocysts, and extra legs. This is actually an optogenetic frog that has an extra leg growing out of its mouth. We can make fins. That's weird. Tadpoles aren't supposed to have fins. It's more of a fish thing, but we'll get to that.

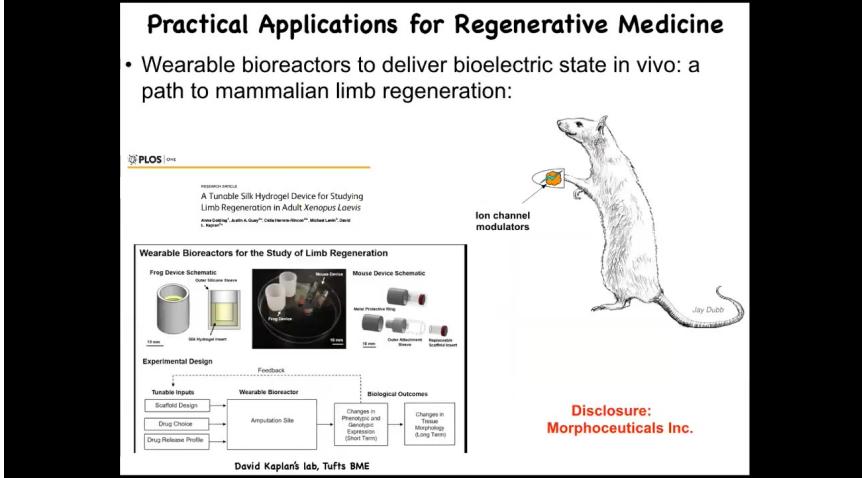
What you can see is that it is possible to rewrite these bioelectric pattern memories in ways that make the collective navigate to different regions of that anatomical morphospace.



And we're applying this in our regenerative medicine program. So frogs normally, unlike salamanders, do not regenerate their limbs. 45 days later, there's nothing.



We came up with a cocktail that within 24 to 48 hours triggers pro-regenerative genes like this *MSX1* in this blastema. By 45 days, they've got toes, they've got a toenail, and eventually a pretty respectable frog leg that's both touch sensitive and motile. The thing about this is that this is not micromanagement. We're not 3D printing anything. We're not controlling stem cells. We're not interacting with this thing at all; 24 hours of stimulation, and then for a year and a half this leg will grow. So it's the decision-making of the collective very early on. It's not trying to micromanage it.



**Practical Applications for Regenerative Medicine**

- Wearable bioreactors to deliver bioelectric state *in vivo*: a path to mammalian limb regeneration:

**Wearable Bioreactors for the Study of Limb Regeneration**

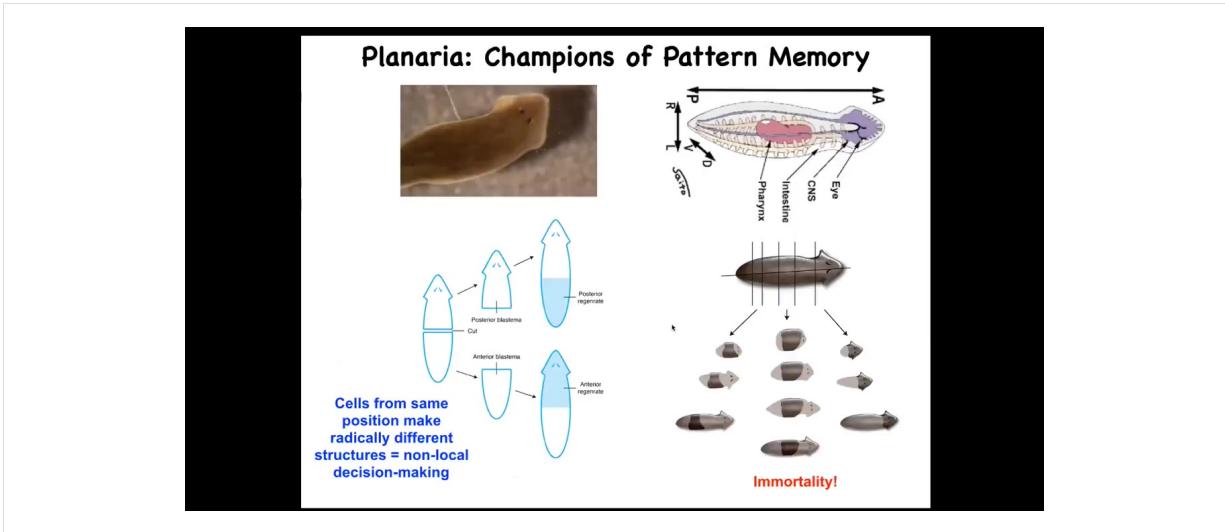
**Experimental Design**

**Feedback**

**Biological Outcomes**

**Disclosure:**  
**Morphochemicals Inc.**

I have to do a disclosure, because Dave Kaplan and I have a company called Morphoseuticals Inc., where we're trying to apply the same technology, a wearable bioreactor and a payload of modulators that convince the cells that they need to rebuild. We're now moving into rodents.

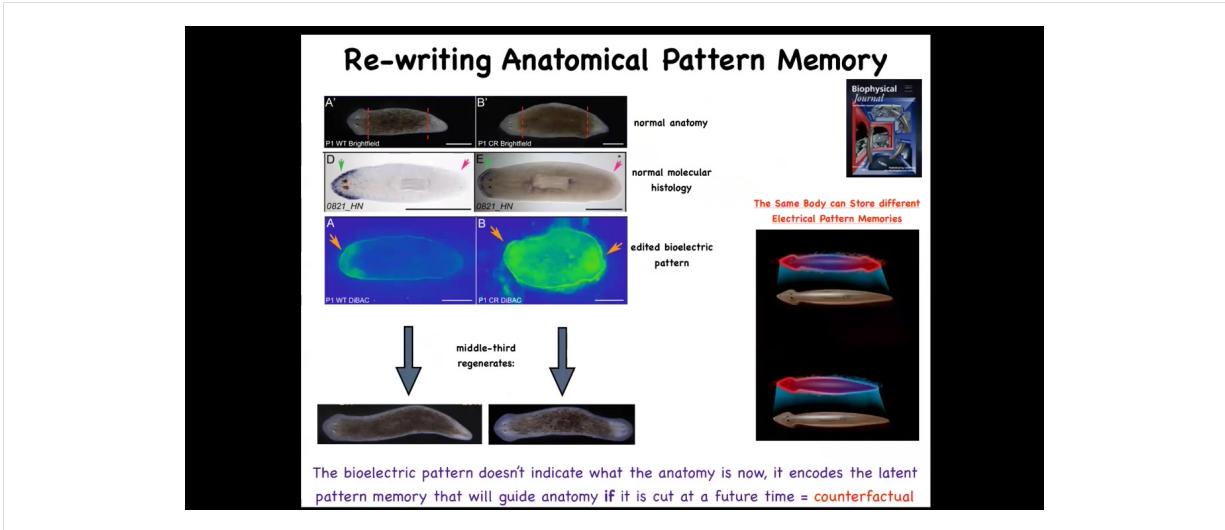


The next story I want to tell you has to do with planaria, because I want to hammer this idea of bioelectrical pattern memories being literal representations of navigation in this anatomical space.

These planaria have an amazing capacity: you can cut them into pieces. 276 is the record. Every piece makes a perfect little worm.

We asked a specific question. Once you've cut a piece, how does it know how many heads it's supposed to have?

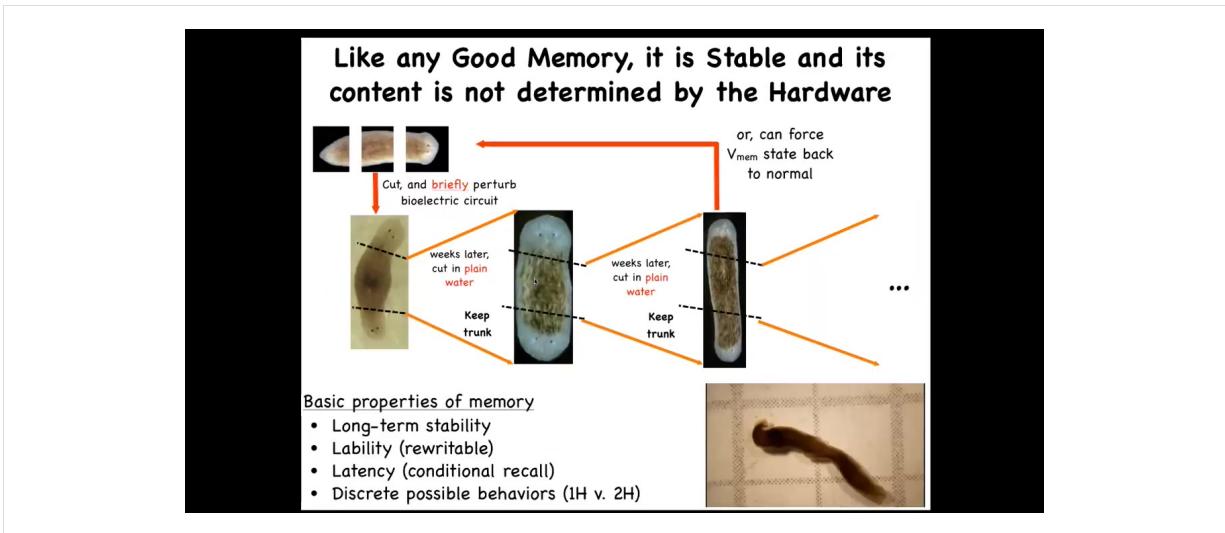
What we found out is that there's a voltage, an electrical circuit that sets a voltage gradient in these pieces that determines how many heads it's supposed to have. If you take a normal one-headed worm, this worm has anterior genes in the head, not in the tail, and it reliably makes this one-headed worm.



But that's because it has this bioelectrical pre-pattern that says one head. And so what we were able to do is take an animal, give it a different bioelectrical pre-pattern. This is a mess; the technology is still being worked out, but you can see here now it's very clear. There's now these two regions. And this animal still has its anterior genes in the head, not in the tail. But if you cut him, he gives you a two-headed worm. And this is not Photoshop, this is real.

Something important to stress here is that this bioelectrical pattern is not a map of this two-headed animal. This is a readout of this one-headed animal, which is anatomically and transcriptionally normal, and it will be that way. It will stay normally one-headed until you injure them. This is a readout of a counterfactual memory. It's what I will do if I get injured. It's what a correct planarian looks like, aka the set point of that anatomical homeostatic cycle.

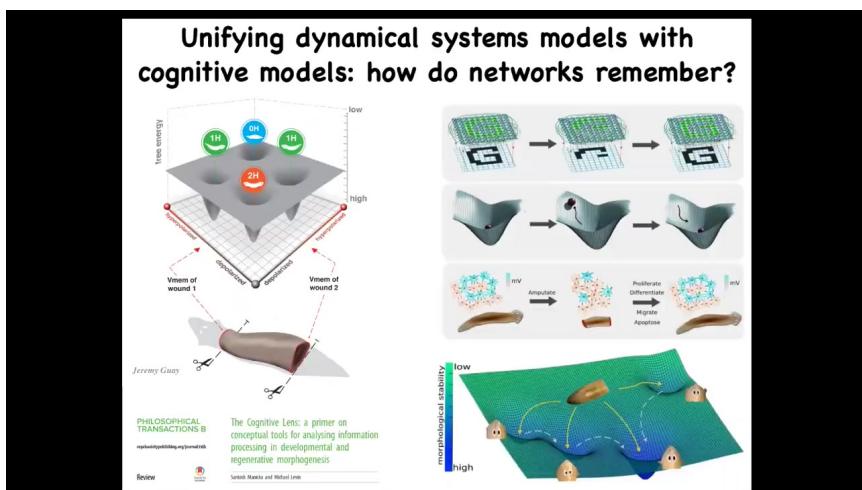
So I promised you at the very beginning that if we are to understand how these systems solve problems in anatomical space, we have to know how do they remember what is the correct pattern. And what you're looking at is one simple example of how they encode goal state information. This is literally a precursor to this amazing ability of human brains and other animals to time travel, to think about things that are not true right now, things that used to be true or that will be true in the future. This bioelectrical state does not describe the anatomy at this moment. It describes a goal state that will be pursued if the animal is injured.



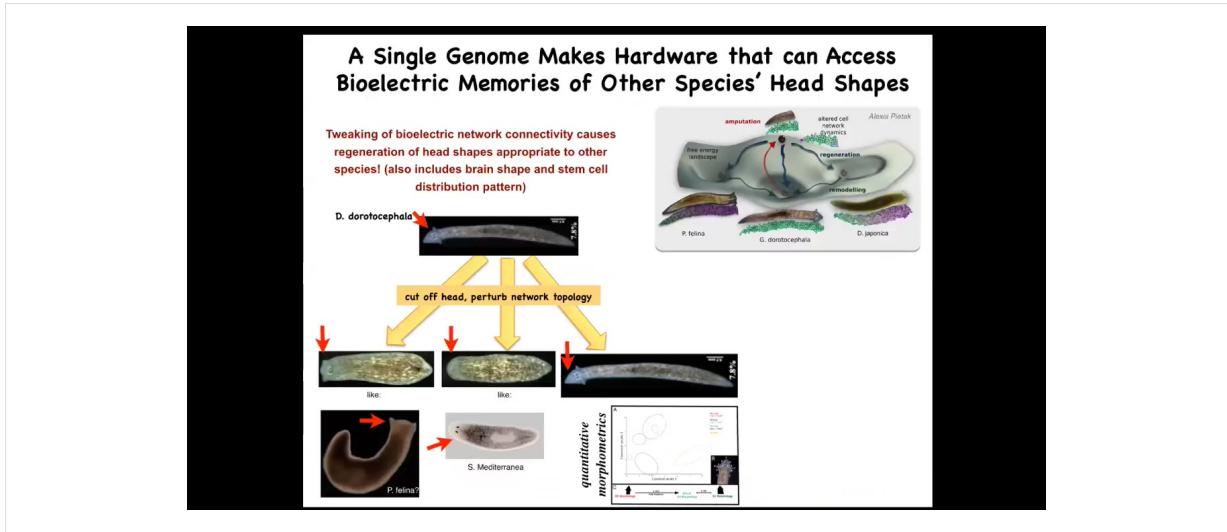
I keep calling it a memory Y, because if you take these two-headed animals, you can cut them again in plain water. So you cut off the ectopic primary head, you cut off the secondary head, and with no more perturbation, they will continue to generate two-headed animals in perpetuity.

Now remember, the genome is normal. There's no genomic change here. This is a very counterintuitive result because it was thought that if you cut off the extra head, this genetically normal fragment is going to regenerate the normal thing, what the genome says. This is why two-headed worms were first seen in 1903. As far as I know, nobody cut them until we did it in 2008, because it was too obvious what would happen. In fact, that's not what happens.

What you see here is that the number of heads is not directly nailed down by the genome. The genome gives you a machine that by default generates a memory of 1 head, but you can rewrite it. It's reprogrammable. Without editing the genome, you can physiologically give them a new memory and they will continue to be two-headed. It's long-term stable, it's rewritable, and it has conditional recall, which I showed you a minute ago. You can see here this video of these guys hanging around.

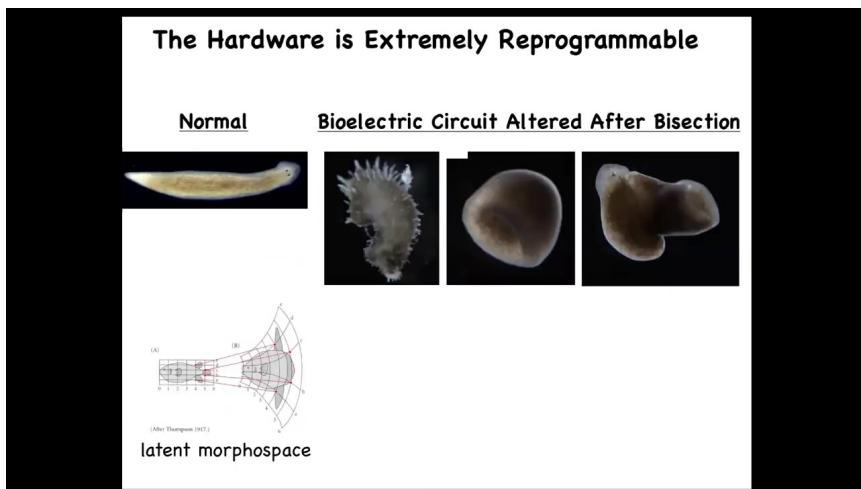


We are trying to put together different ways of thinking about systems that have different attractors. The state space of the bioelectrical circuit corresponds to one head, no head, two head, and so on. There's lots of thought in connectionist AI and neuroscience about how electrical networks can store memories of patterns and regenerate those patterns upon damaged or partial inputs. There is a lot of work to be done to integrate these, but we think they're both different views on the same set of processes.



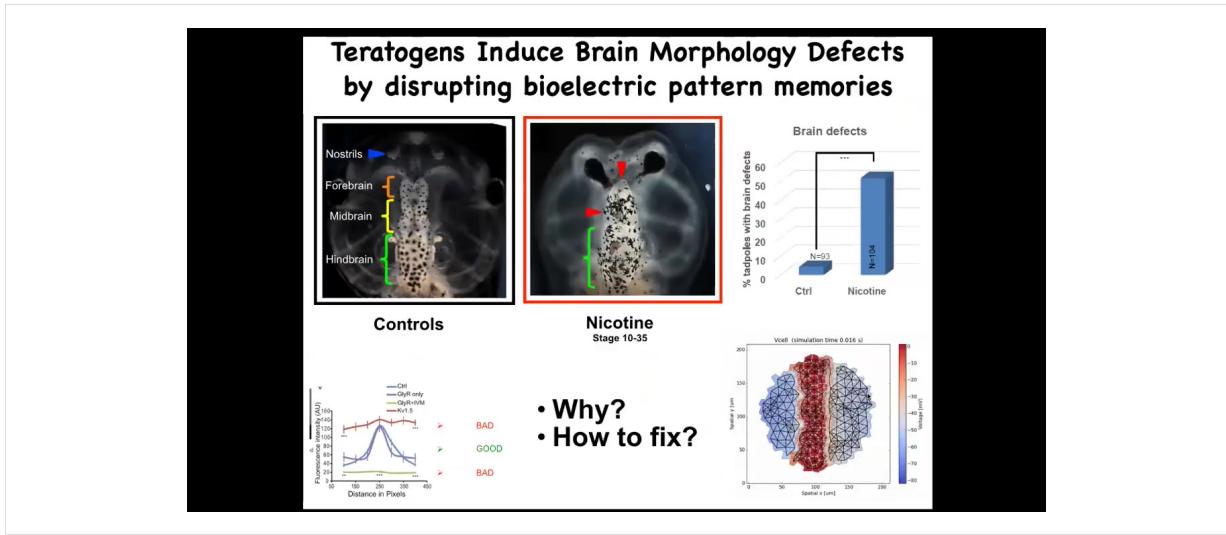
It's not just the number of heads, it's also the shape of the head. So this triangular-shaped planarian, by manipulating these bioelectrical patterns, could be made to make a flathead like a piphalina or a round head like an *S. mediterranea*. No genomic change, but 100 to 150 million years' evolutionary distance between this and this. It also changes correctly the shape of their brain and the distribution of stem cells like these other species.

This same hardware is perfectly able to visit the attractors in morphospace belonging to these other species. That's normally where these other species go, but this hardware can go there too, if properly guided in morphospace.

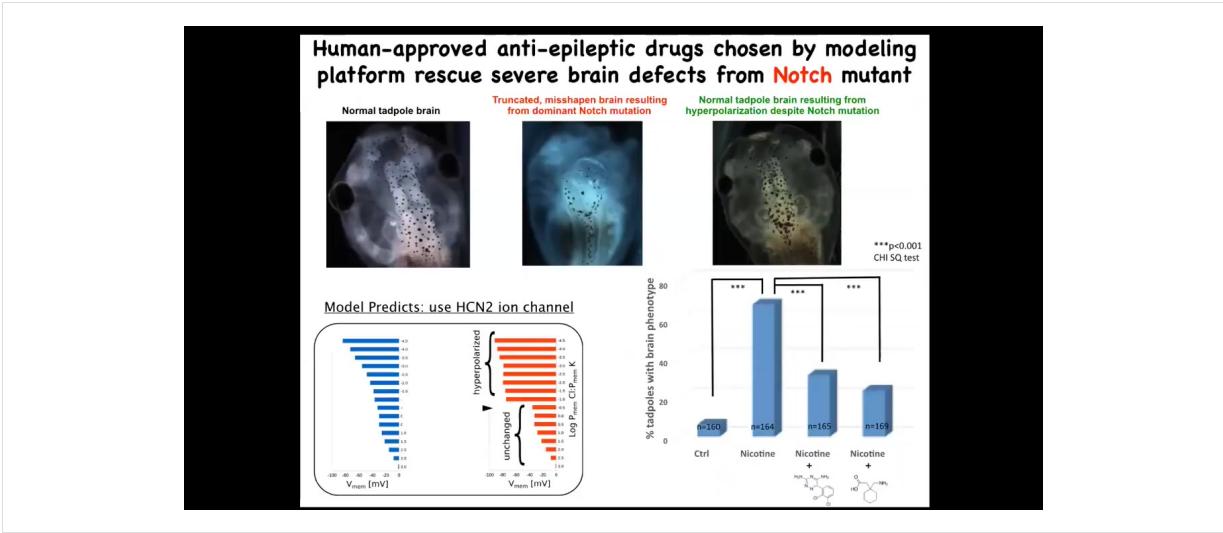


In fact, you can go way beyond things that even look remotely like planaria, and we've seen planaria make things like this: spiky forms, these kinds of cylindrical shapes, and this combination hybrid shape. I think this is what D'Arcy Thompson in 1917 saw as a latent morphous space around each animal and plant form. He drew these axes and it was entirely unclear what these axes were, but I think he's representing the internal encoded map of the collective intelligence of these cells and that we now have the ability to stretch and deform that internal representation.

So what we're trying to do is the same thing that multi-scale neuroscience wants to do, which is to go all the way from the molecular level where different channels and pumps are expressed, to the spatial bioelectrical patterns, to the organ level and the whole-animal level anatomical order, and ultimately to algorithmic models that describe this behavior so that we can rationally understand and manipulate what's going on. So a full stack approach here.



I'm going to show you one example where we've been able to do that, which is this in the frog brain. Here's the normal frog brain, forebrain, midbrain, hindbrain. There are a variety of teratogens, both chemical and genetic, that screw up the brain. We decided to check whether we could use a computational approach to try to repair this very complex structure. We made a computational model of the bioelectrics around the pre-pattern that tells this brain the shape and the size that it should be, having characterized this pre-pattern years before that.

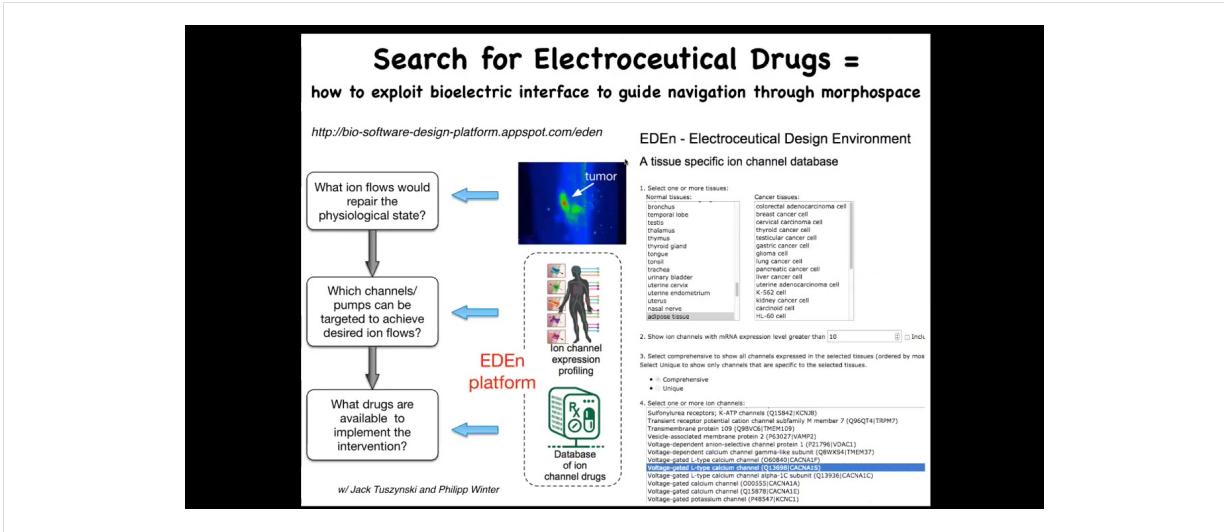


What we were able to do was create a model. We've done nicotine and alcohol and different things, but here's the most extreme case.

Most of you will know the Notch gene, which is a very important neurogenesis gene. When Notch is mutated to a dominant active form, what you see is that there is no forebrain. The midbrain and hindbrain basically are a big bubble. These animals have no behavior and are profoundly disabled.

We created a computational model of this. We asked the model, "What channel could we open or close to get back to the correct bioelectrical state?" It focused in on this. It's a very interesting channel with some specific properties, called HCN2. We found that when we upregulated HCN2, either by injecting RNA encoding it or with HCN2 openers, you basically get back to a normal brain, anatomically normal in terms of gene expression, and they get their IQs back. They have indistinguishable learning rates from control animals. All of this while still having that Notch mutation.

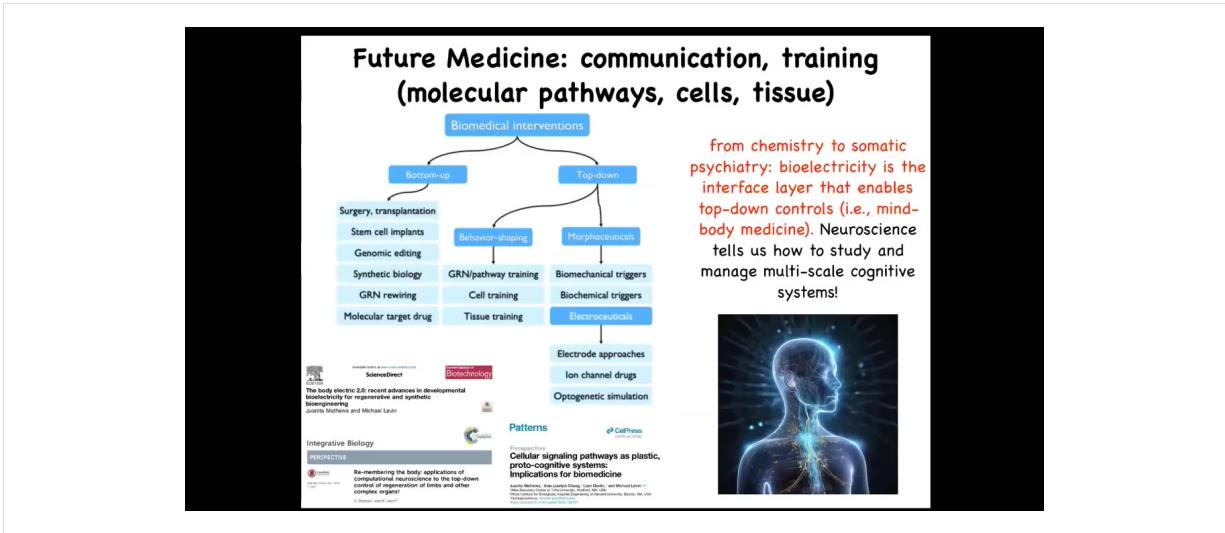
So at least some types of genetic defects, meaning hardware defects, can be fixed in software. I'm not saying they all will be like that. There will be many that are not like that. But it really is amazing how much control you can have at the software level and override what really is a pretty profound hardware defect.



What we're doing now is we're trying to create this platform, and you can play with an early version of it here online, that ultimately is going to have different tissues and organs, both normal and transformed.

You can say I have an abnormal bioelectric pattern in some region, and it will computationally model it, then give you a list of channels and pumps that are expressed in that region, which we know from scraping transcriptomics profiles. It will give you targets to try to get the incorrect pattern into the correct pattern. Then you can pick drugs, because something like 20% of all drugs are ion channel drugs. This is a huge toolkit of potential electroceuticals if we understood the modeling.

The idea, again, is not to micromanage the state of any one cell. The idea is to communicate with the collective and try to get it to move to different regions of that amorphous space and exploit the different competencies that it has for repair, for scaling of activity.



I'm almost done here.

All the details are in these papers. I want to put forward a view of biomedical interventions as fundamentally in two categories. There are the bottom-up ones that seek to interact with the hardware and micromanage the symptoms and functionality. And then there is this stuff, which is just beginning to be exploited. The idea is that top-down, you can take advantage of the many competencies of your material.

There's something special here because neuroscience brings something extremely important. It's a unique field for all of this. We've been doing these applications in birth defects, normalizing cancer, and regeneration, because in neuroscience it's the only field where you can see how electrophysiology mediates from reasoning, so from high-level cognitive intent to molecular causes.

What I mean is that if I told you that with the power of my mind alone I could depolarize 30% of my body cells, you might think that's crazy, or you might think that's some special mind-body medicine thing, maybe some biofeedback. But it's a normal 24/7 event that happens during voluntary movement. In the morning, when you make this high-level executive decision that you're going to go to work, that information flow has to trickle down to the opening and closing of ion channels in your muscle cells to get them to fire so that you can get out of bed and go to work.

That incredible architecture our bodies have links all the way from high-level cognition through the circuits you study and down into the molecular control of cells. It is a unique multi-scale understanding still missing in other fields and can be imported to great use.

My gut feeling is that all of this means that future medicine in these areas, so regenerative medicine, is going to look more like a kind of somatic psychiatry where you're literally communicating with a collective and manipulating its prior memories,

its measurements, its beliefs about the energy landscape that it's in these different spaces. We have a lot to learn from how multi-scale neuroscience does this.

Slide 45 of 47 · Watch at [1:00:51](#)

### Summary:

- **Intelligence and navigation** applies well beyond conventional brainy animals in 3D space
- We are all **collective intelligences**; neuroscience is not just about neurons
- Morphogenesis is the **behavior of a cellular collective intelligence in anatomical morphospace**
- The cognitive glue of this collective intelligence is the same as it is for neurons in the brain: **bioelectricity** = the evolutionary precursor of CNS
- We can now read and write memories in this new collective intelligence with **applications in birth defects, regeneration, and cancer**
- This interesting invariance across fields means that **tools and concepts can be ported between neuroscience and developmental biology**, to the enrichment of both.

So what I've told you today is that intelligence and navigation of problem spaces is, I think, a generic concept that's applicable far outside of standard behavioral science. We're all collective intelligences.

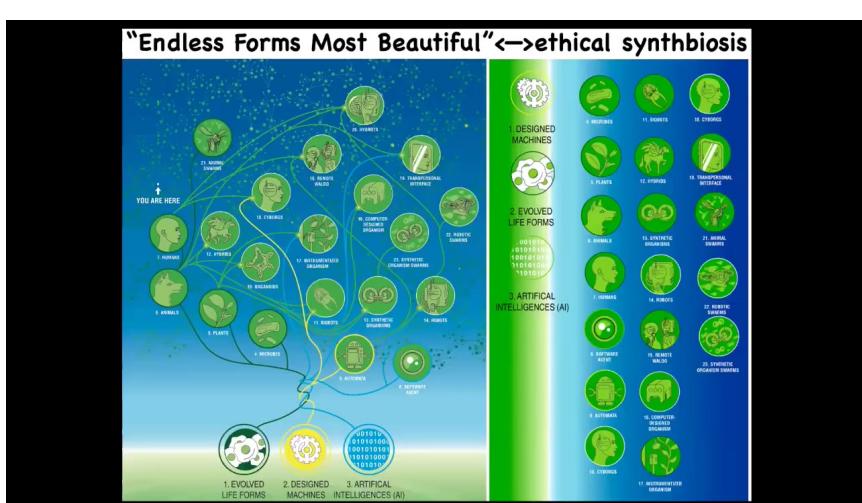
We use everything from active inference and perceptual control theory and bi-stability, multi-scale processing in the retina. All of these things we use outside of the nervous system to help us achieve better control in regenerative medicine. I've talked about the behavior of a collective intelligence of cells in morphospace as underlying these amazing properties of development of regeneration. I really think that this is where the brain learned its amazing tricks.

The last thing I want to show you very briefly is this idea that there's this really interesting invariance across fields that can enrich both developmental biology and neuroscience. For many years, I was saying to my students that you can take almost any neuroscience paper and do a find replace. Every time it says neuron, you say cell. Every time it says millisecond, you say minute or hour. Then you've got a nice potential developmental biology paper full of hypotheses to explore.

Finally, with two collaborators, Thomas O'Brien and Joel Stremmel, we made an AI tool, and you can all play with it here, that does this in a very smart way by mapping a few terms to each other. You can take any abstract in a neuroscience paper and see

what the corresponding developmental biology paper looks like. This is a beautiful paper by Panotsis and Miller. You go on here and you feed this abstract, and what it gives you is a very plausible set of claims that is testable in developmental biology. We are developing these kinds of tools to help exploit the invariances between fields to enrich both sides of this equation.

Slide 46 of 47 · Watch at [1:03:05](#)



And the final thing is that because of this amazing interoperability of life and the fact that cells and tissues are a kind of material that are happy to work with, all sorts of novel configurations and new materials and so on. All of existing biology, everything that Darwin meant when he said "endless forms most beautiful," is like this tiny corner of option space. In the future, meaning the next couple of decades, we're going to be surrounded by all of these things that already exist, especially in neuroscience.

Cyborgs, hybrids, different kinds of chimeras and combinations, all kinds of bodies and minds that we've never seen before that have never been on the tree of life with us. We are going to be looking for insights from diverse intelligence research to expand the way that neuroscience and behavioral science thinks about agency, cognition, memory, and how we are going to ethically relate to these creatures that actually don't look anything like us. You can't use the familiar kinds of things we use when talking about brain organoids in a dish. What does the brain look like? Is it like a human? It doesn't need to be like a human. There are going to be all sorts of wild combinations that we need to think about.

All of this is discussed in these various papers.



I just want to thank the postdocs and the students that did all the work, our many collaborators and all of our technical support. Here are three companies that support our work and many other funders, and most of all the model systems, because the animals do all the heavy lifting. So thank you so much and I'm happy to discuss.

**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](mailto:adi@aipodcast.ing).