

A LECTURE COMPANION

---

# **Bioelectricity, Biobots, and the Future of Biology**

---

Michael Levin

Recorded on November 25, 2023

## About this document

This document is a companion to the recorded lecture *Bioelectricity, Biobots, and the Future of Biology*, recorded on November 25, 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

A 30-minute talk (from April 2023) by Michael Levin given to an audience interested in longevity and the future of biotechnology research.

---

## 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).

**Design Challenge: Total Control of Growth and Form**

**Why we need it:**

- Birth defects
- Traumatic injury
- Cancer
- Aging
- SynBioMorph.
- Non-neural GAI

Problems of biological information processing

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

**Disclosures:**



Morphoceuticals

[morphoceuticals.com](http://morphoceuticals.com)



ASTONISHING LABS

[astonishinglabs.com](http://astonishinglabs.com)



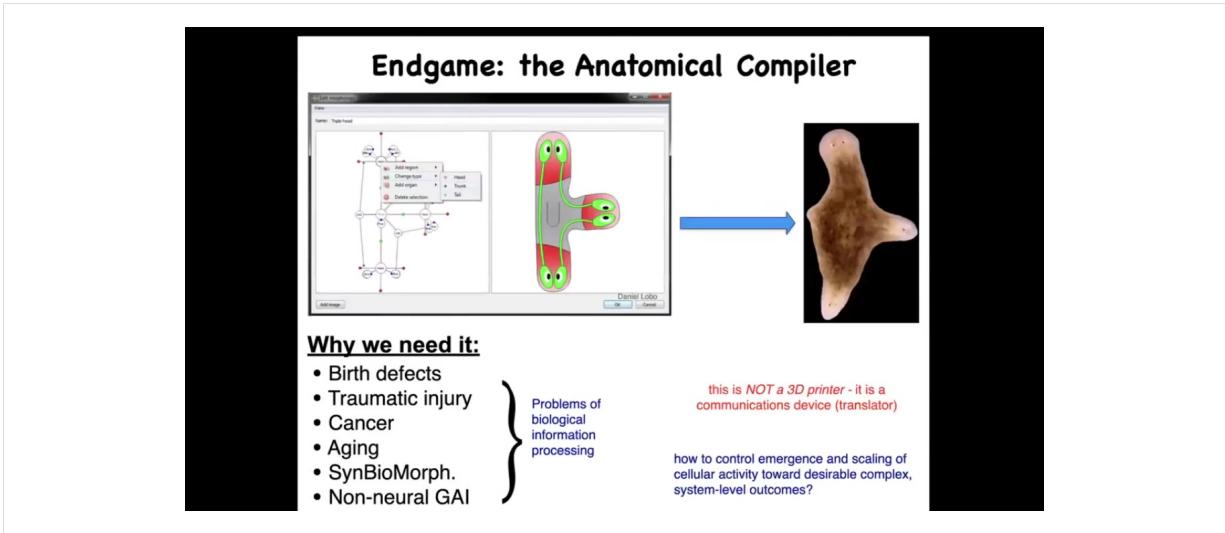
FAUNA SYSTEMS

is interested in any of the details, the software, the data sets, or getting in touch with me, this is the website where everything is. What we're going to talk about is the opportunities around bioelectricity and in general around cellular collective intelligence in bio-robotics, regenerative medicine, and so on.

The design challenge that we have can be summed up as the total rational control of growth and form. What I mean by that is that if we had the ability to tell groups of cells what to build in terms of three-dimensional anatomical structure, most medical problems would go away. Birth defects, regeneration after traumatic injury would be possible; cancer could be reprogrammed; aging and degenerative disease.

Not only biomedical problems, but synthetic biomorphology — not just reprogramming individual cells, but programming a large-scale shape and structure. Non-neuromorphic AI would become possible. All of these things boil down to one problem. How do we communicate our morphological goals to a group of cells? I have to do disclosures because I'm involved in three different spin-offs from my lab, which I'm happy to talk about afterwards: more for suitables, astonishing labs and fauna systems.

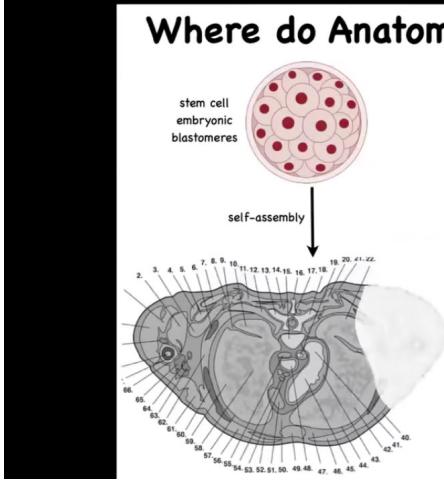
1



The solution to these problems we call the anatomical compiler. The idea is that at some point in the future, you will be able to sit down in front of a computer and draw the animal or plant that you would like to have. You would draw the three-dimensional structure, the anatomy, not the molecular biology of it, but actually the anatomy of what you want. This system would compile that anatomical description into a set of stimuli that would have to be given to cells to get them to build whatever you want. And in this case, here's this three-headed flatworm. I'll show you some flatworms in a minute.

There are two really important pieces to this. One is that this is not a 3D printer. The point is not to micromanage the position of every cell and try to assemble something complex piece by piece. It is also not anything to do with genomic editing or CRISPR. This is fundamentally the right way to think about this: as a communications device. This system and the AI that will make it possible are for translating our anatomical and functional goals to the goals of a collective intelligence of cells that will build this thing. We don't have anything remotely like that. We can only do this in very few specific cases, but I wanted to lay out the opportunity ahead. This is the roadmap.

## Where do Anatomies Come from?



The diagram shows a cross-section of a developing embryo. At the top, a cluster of red circles is labeled 'stem cell embryonic blastomeres'. An arrow labeled 'self-assembly' points downwards to the embryo, which is labeled with various numbers (2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100) representing different cells and structures.

- DNA specifies proteins; whence Anatomy?
- how do cell groups know what to make and when to stop?
- why do cells defect from target morphology (= cancer)?
- How to repair and remodel damaged or aging structures?
- How far can we push shape change? Engineers ask: what's possible to build given default genome?

Why don't we have an anatomical compiler? Now let's just remind ourselves that we all start life as a group of cells, and these embryonic blastomeres reliably build something like this.

This is a cross-section through a human torso. You can see the incredible complexity of all the tissues, organs. Everything is in the right place relative to each other. The question is, where does this shape come from? Where is this order stored? Where is it recorded? You might think the genome—everything's in the genome, but actually we can read genomes now and we know what's in the genome. What the genome encodes is the tiny protein hardware, the molecular hardware that every cell gets to have. There is nothing directly in the genome that says anything about size, shape, the number of organs, the types of organs.

The fundamental scientific challenge here is to understand how this group of cells exploits that hardware to build something reliably. How do they know when to stop? Why do cells sometimes defect from this plan, resulting in cancer? As workers in regenerative medicine, we'd like to know if something is missing, how do you convince the cells to regrow it? If, as engineers, we would like to say, what else can these cells build? With a standard genome, can you get the exact same hardware to build something completely different? These are the questions that we face.

I want to show an extremely simple example to remind us why genetic information is not enough. This is the larva of an axolotl. Baby axolotls have four little legs. This is a tadpole; baby frogs at this point do not have legs. In my lab, we can make something called a frogolotl.

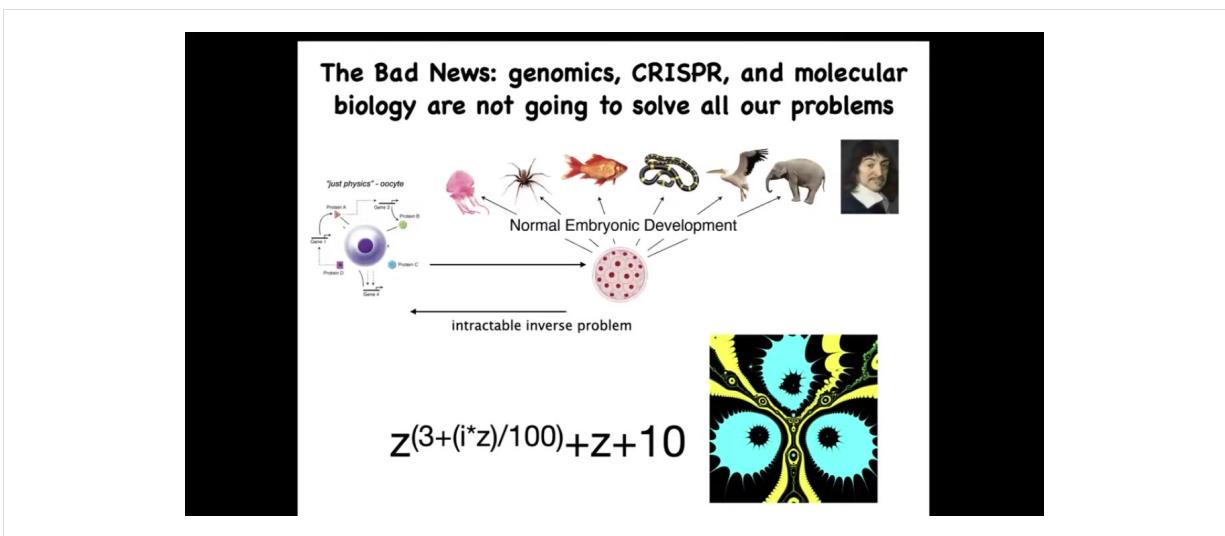
What's a frogolotl? It's a chimeric embryo. That's part cells from an axolotl, part cells from a frog. I can ask a very simple question. We have both genomes. The axolotl

genome has been sequenced, the frog genome has been sequenced. Could anybody tell me whether frogolots are going to have legs or not? The answer is no.

There is no existing model that will go from genomic information to being able to predict anatomy in cases like this. We can't go from the genome to a description of what the anatomy is going to be without cheating, comparing that genome to something where you already know what it is. There is this fundamental disconnect between the molecular hardware and the large-scale anatomical decision-making of the collective of cells that's going to decide to build this large-scale structure, possibly involving cells from the frog.

This is the frontier. The frontier is understanding how large-scale collective decisions are being made that determine our anatomy and our function.

Slide 4 of 29 · Watch at [5:19](#)

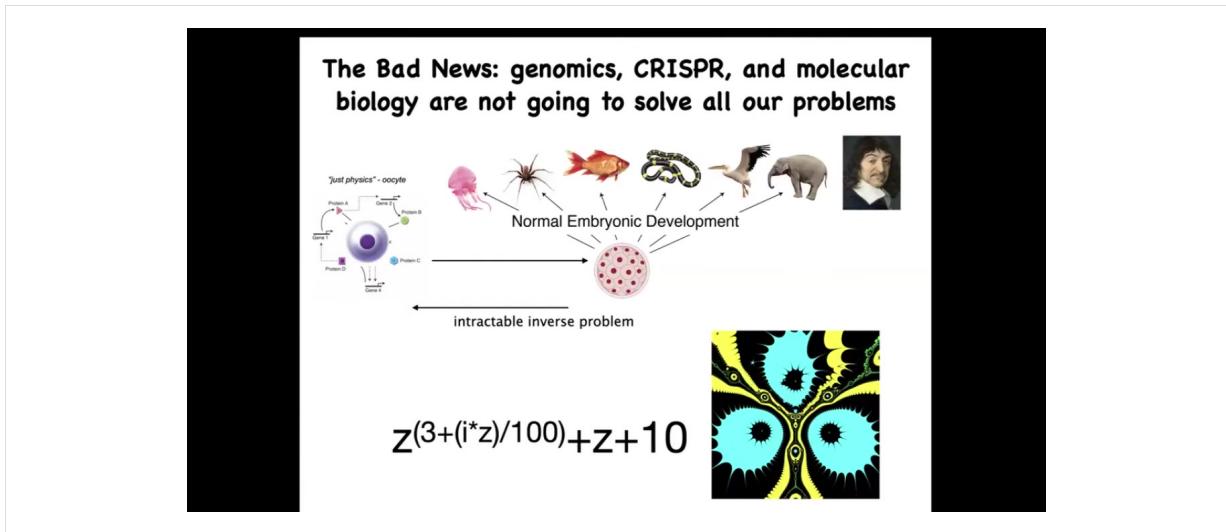


Now the bad news is that by themselves, these very important technologies, such as genomic editing and molecular biology, are not going to give us everything that we want. Because this process of going from a single cell egg to one of these complex bodies — this morphogenetic process — is incredibly complex and recurrent. In other words, all of these cells are doing various things and together they end up building something like this. But reversing this is extremely hard.

To give you a very simple example from the fractals: this is a simple mathematical formula. It's very short, very simple. If you iterate it and plot a version of the Julia set, you will get this very complex pattern. That's great. Going forward is easy. Starting

with the rules and getting something complex is very simple. Going backwards, if you were a worker in regenerative medicine, this thing isn't quite symmetrical. I don't want three of these. I want four. How would you have to change the formula?

Slide 5 of 29 · Watch at [6:21](#)



There is no way to do that. This is an intractable inverse problem. Figuring out what you would have to change on the molecular or genetic level to make large-scale changes and repairs is extremely difficult. It'll be possible for some cases, but it needs to be complemented with a top-down approach.

**State of the Art: biological information & control**

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)

Now, here's where we are today. We're very good as a community in manipulating molecules and cells, and there are amazing kinds of advances all the time in terms of synthetic biology and protein engineering. We're actually still a very long way away from control of large-scale form and function.

Injuries, cancer, birth defects, these things are largely not yet addressable. I think that's because we are still fundamentally in molecular medicine; this is where computer science was in the 1940s and 1950s, where in order to reprogram the machine, you have to physically interact with the hardware. You have to physically rewire. This is where all the excitement is today: gene regulatory circuits, novel molecular interactions and so on. But in addition to that, we have the opportunity to do something amazing, which is to take advantage of the high-level information processing, AKA intelligence in the layers of biology.

**The Good News:**

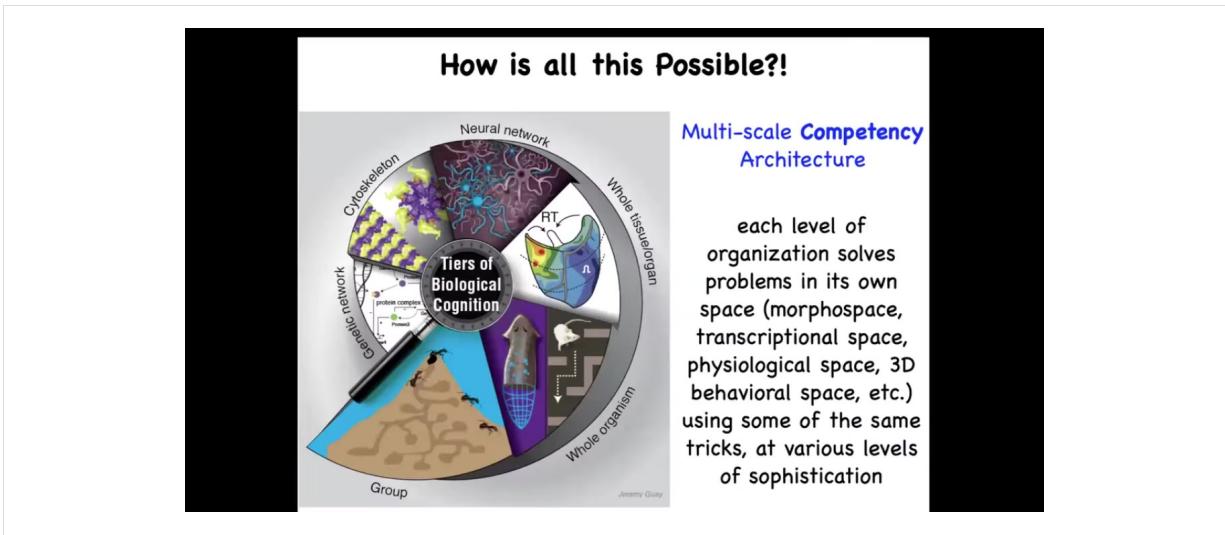
**“Intelligence is the ability to reach the same goal by different means.”**

- William James

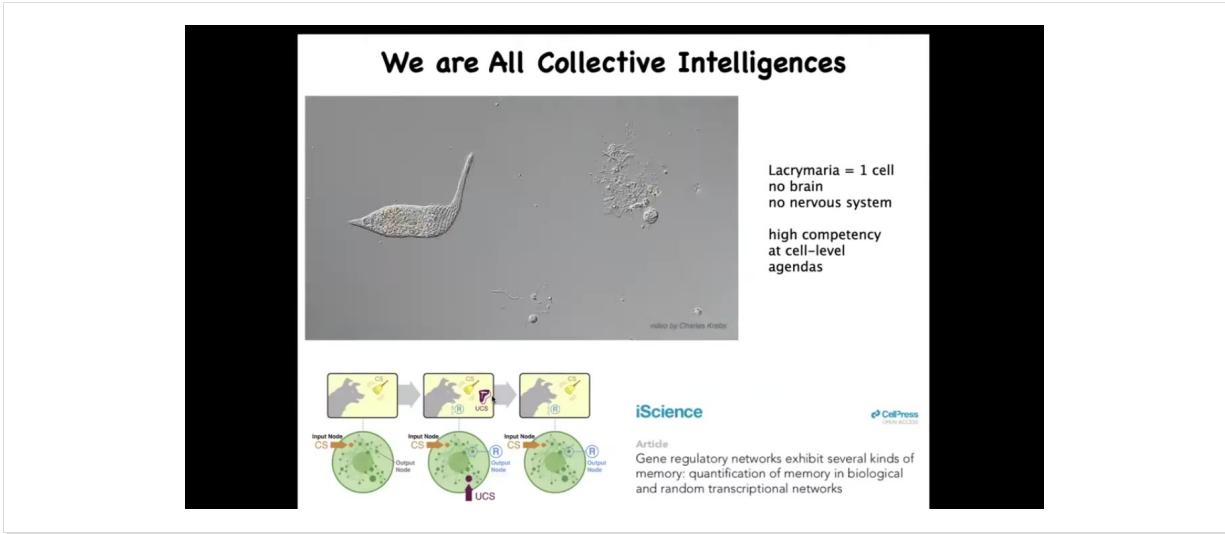
use AI as translator between us and the collective intelligence of cells navigating morphospace

When I say intelligence, I don't mean a human-level, second-order "I know that I know" kind of thing. I mean something else that might be at different locations along this continuum. This sort of continuum is due to Norbert Wiener and colleagues. William James defined intelligence as the ability to reach the same goal by different means.

I don't have time to show you many examples, but I'll show you a few. Biology is really good at this. What we now have the opportunity to do is to use AI as a kind of translator between us and this robust intelligence of cells and tissues that are navigating something we call morphospace. Morphospace is a virtual space of all possible morphogenetic or anatomical outcomes. We have the ability to control how the living system navigates that morphospace.



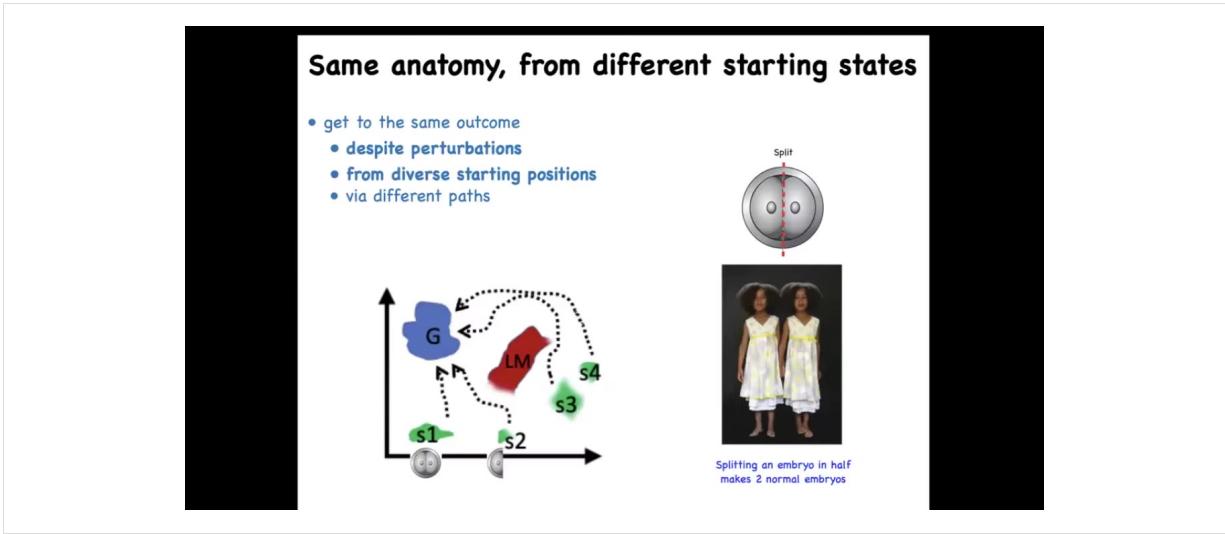
And it's possible because biology, unlike most of our technology, I sometimes give a different talk called "Why Robots Don't Get Cancer." The reason that biology is so special, and there's no reason why we couldn't emulate this, is that it uses a multi-scale competency architecture. The idea is that at every level, from molecular networks to cells, to tissues, to whole bodies, whether in behavioral space or anatomical space, or even collectives, each of these layers solves its own problems. They are not just structural. They are problem-solving systems that operate or navigate in the space of gene expression, the space of physiology, the space of anatomy, and of course, the three-dimensional behavioral space, in our case also linguistic spaces. And they all have different levels of sophistication and different competencies to do this. But the key is that rewiring them physically is not the only game in town. We also have the opportunity, because their intelligence is higher than a simple hardwired clockwork, to communicate with them, to reset their set points, and to have various other kinds of interesting interactions with them.



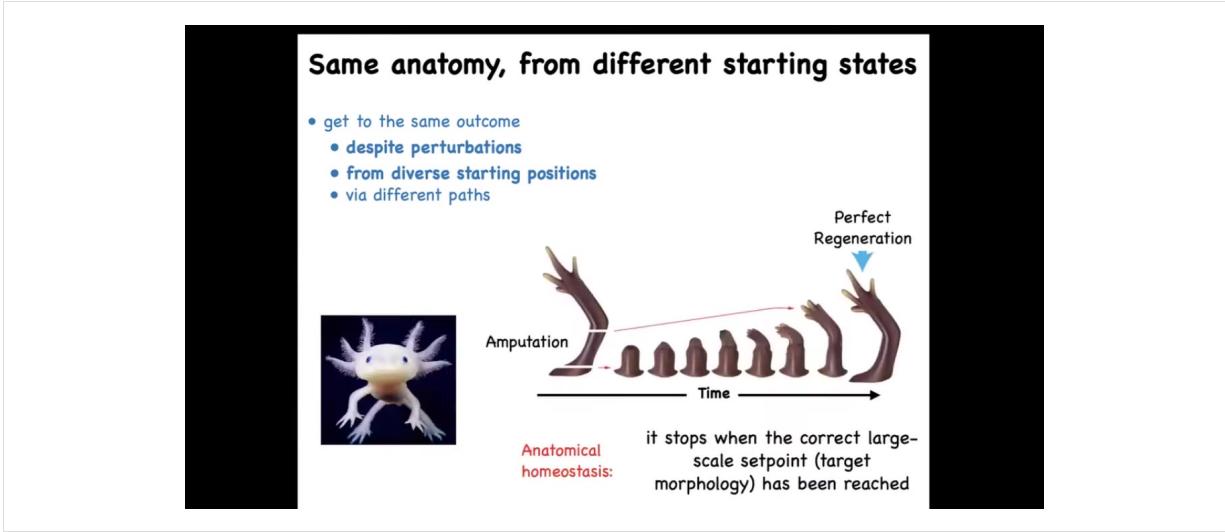
So the thing to recall is that fundamentally, we are all collective intelligences. There's no such thing as an indivisible diamond of intelligence. We're all made of parts. And the parts that we are made of are individual cells.

Now, here's a cell that happens to be a free-living organism known as a lacrimaria. One cell. There's no brain, there's no nervous system. This thing is incredibly competent at very local, tiny goals, metabolic, physiological, and other states. You can see it here feeding in the environment. Anybody who's into soft robotics should be drooling at this. We don't have anything remotely like this level of control. And again, just one cell.

And it is made of molecular networks that also have competencies. For example, the gene regulatory circuits and pathways inside cells are capable of at least six different kinds of learning, including associative conditioning. And this is all without changing their structure at all. So this is not about rewiring GRNs with gene regulatory networks through experience. This is *in situ* exactly how they are. They're capable of learning.

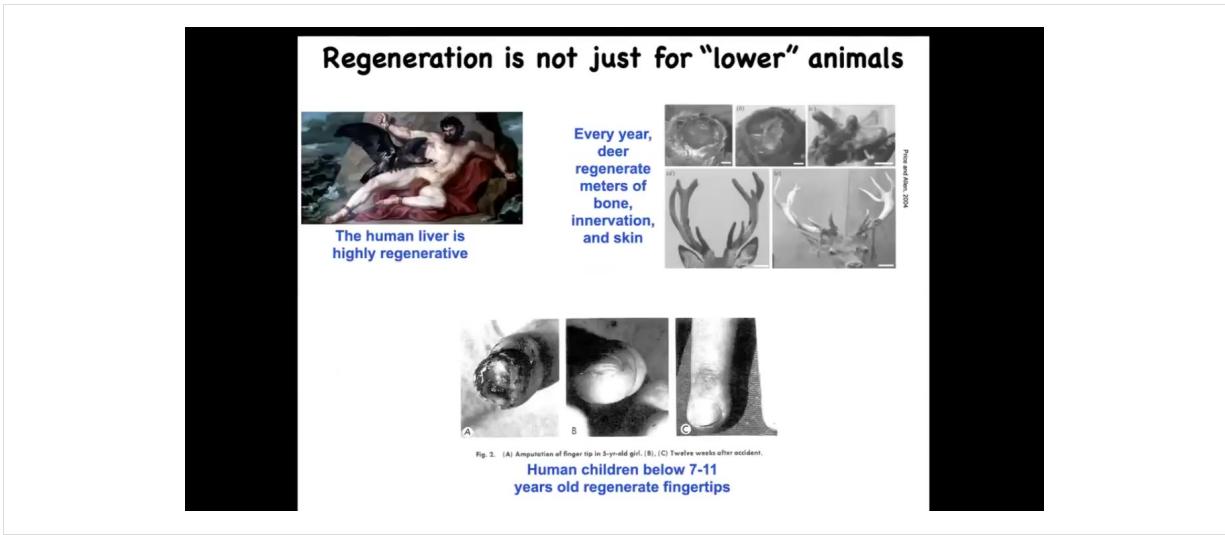


This multi-scale architecture of competency, meaning that we are made of cells that used to be independent organisms and didn't drop all their intelligence when they joined together, means that we have lots of different competencies in morphogenetic space. For example, when you take an early embryo, a human embryo, and you cut it into halves or quarters, you don't get 2 half bodies, you get 2 perfectly normal monozygotic twins because each half can recognize that the other half is missing and rebuild what it needs. In fact, you can think of embryonic development as a kind of regeneration of the entire body from just one cell. There's this ability of different systems, here's the normal embryo, here's the embryo cut in half and so on: the ability of different systems to navigate this space to reach the ensemble of goal states that we associate with normal target morphology or normal anatomy for that species, and in fact, in many cases, going around avoiding various local maxima that would trap a less intelligent system. This is one example of regulative development.



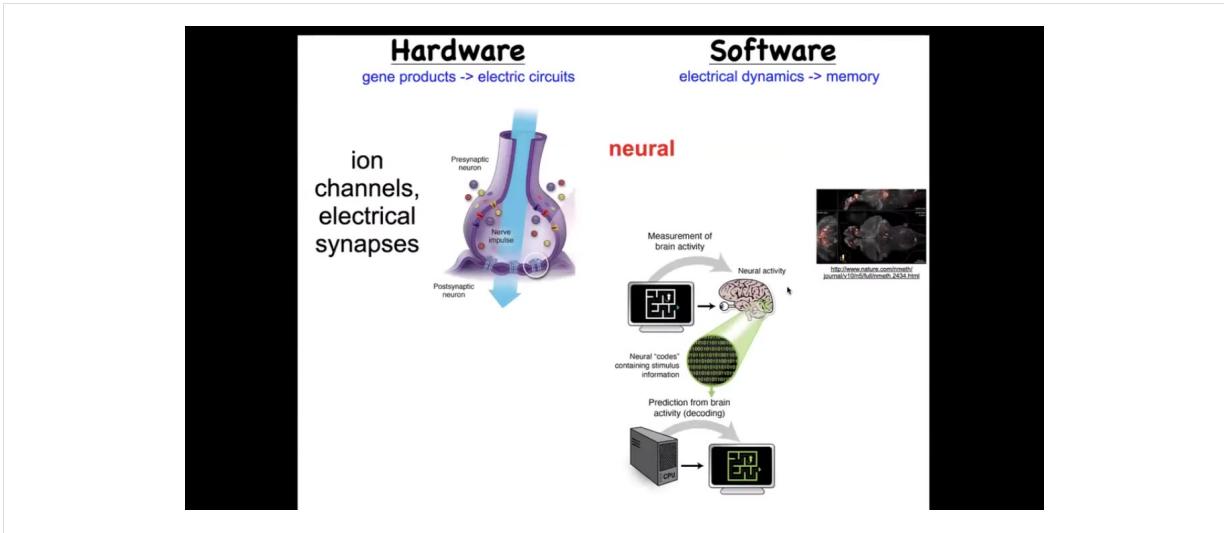
Here's another example. A creature like this is an axolotl. And these guys regenerate their legs, their eyes, their jaws, their portions of their heart and brain, their ovaries, spinal cord, and so on. If you amputate, and they bite each other's legs off all the time, so this is a natural experiment that keeps happening, you can amputate anywhere along this line, and the cells will very rapidly build exactly what's needed. And then they stop. Now this is the most amazing part of this whole process. Not only does it regenerate no matter where you cut it from, it makes exactly what's needed, no more, no less, but it ever stops. How does it know when to stop? It stops when a correct salamander arm has been completed. So this tells you that this is some sort of navigation policy in this anatomical space. It's an error minimization scheme where errors like this become corrected over time and when the error is small enough, they stop.

So we really need to understand this ability to navigate. All the tools of autonomous self-driving vehicles and control theory from engineering and AI become very relevant here.



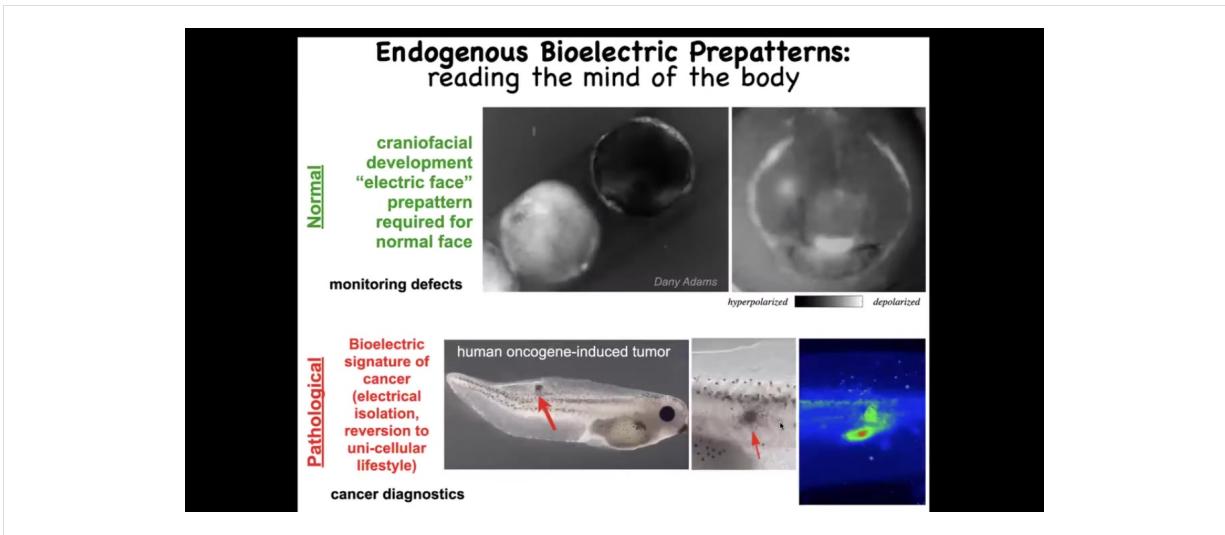
This is not just for salamanders and worms. Mammals do this too. Humans regenerate their livers. Human children regenerate their fingertips up until a certain age. If you leave it alone, an amputated fingertip becomes cosmetically a very nice finger.

Deer are large adult mammals that regrow meters of new bone, vasculature, and innervation. They do it every year at a rate of about a centimeter and a half of new bone per day. It is not the case that mammals are somehow unable to regenerate, but our capacities at present, and I do think this is changeable, are not as good as those of certain other organisms.



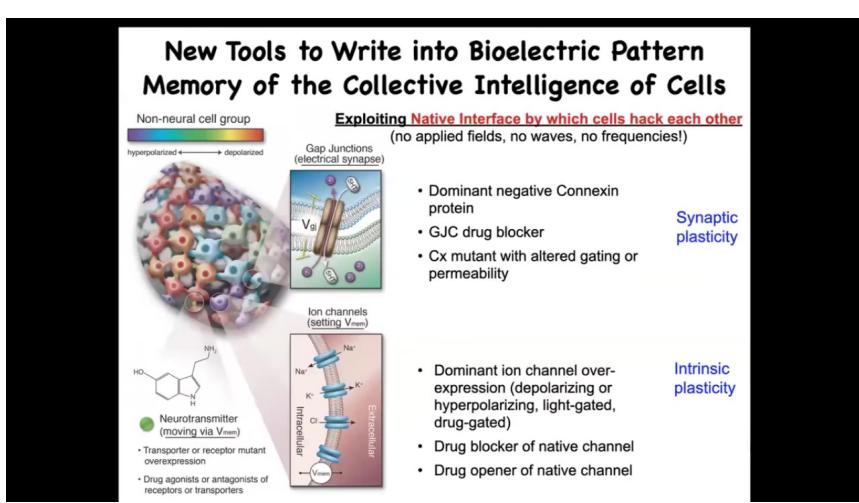
We've been studying for some years how it is that this happens. How does a collection of cells know what to build and when to stop? We took our inspiration from the nervous system, which is a collection of cells that guide your body through behavioral space using memories and goals. Exactly the same thing, it turns out, happens in the anatomical space. The way it happens in the brain is that the hardware looks like this. It's a bunch of cells we call neurons. They're electrically active. They have these little ion channels in their membrane that set a voltage. That voltage sometimes is communicated to their neighbors through these gap junctions or these electrical synapses. There's this network and there's this physiology.

Here it is: this group imaged all the electrical activity of a living zebrafish brain as zebrafish are thinking about whatever it is that zebrafish should normally think about. You can see all this activity. It's the commitment of neuroscience that all of the cognitive activity, the goals, the preferences, the memories, the behavioral repertoires of this animal are in this physiological activity, if only we could decode it. There's this notion of neural decoding where you can scan a brain, decode it somehow, and figure out what the animal is thinking about or looking at. There's been success in this both in humans and in model systems. The amazing thing is that brains and neurons didn't invent it from scratch. This is something that living systems have been doing since the time of bacterial biofilms. That far back, evolution discovered that electrical networks are amazing at storing and processing information. Every cell in your body has these ion channels, and figuring out what is a neuron and what isn't is not trivial at all. Most cells are connected to their neighbors with these electrical synapses. You can look at early embryos or regenerating limbs or anything else using the same techniques, electrical imaging, and try to do the same decoding program. You could ask, could we learn to decode what the cells are saying to each other, and then rewrite those patterns and get them to build something else?

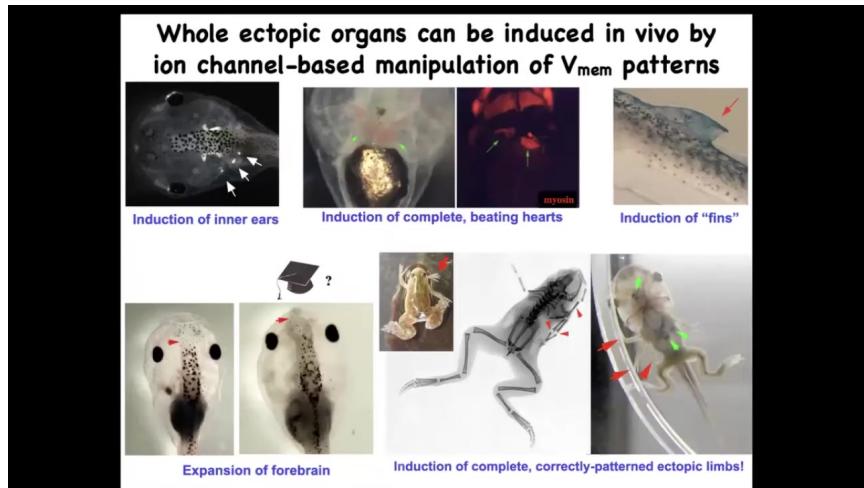


This is what we've been working on. I want to show you two examples of these bioelectrical patterns. This is using a voltage-sensitive fluorescent dye. Here the grayscale colors represent a voltage of the cells. This is a time-lapse movie of an early frog embryo putting its face together. This is a frame, one frame out of that movie. We call this the "electric face." Before the genes come on that are going to pattern this whole structure, you can see this is where the eye is going to be, this is where the mouth is going to be, the animal's left eye comes in later. This is the placodes out to the side that give rise to all kinds of tissues. You can read the electrical pattern memory that's going to guide what happens next. I'm going to show you how we know that this is an instructive set of interactions.

This is a normal pattern. This is required for normal development. This is a pathological pattern. Here we've injected a human oncogene into an embryo. Eventually it's going to make a tumor. Eventually the tumor will metastasize. Using this voltage dye, you can see what's going on here. You can see that these cells are electrically decoupling from their neighbors, they're acquiring a weird depolarized membrane potential, and they're just going to roll back to their amoeba lifestyle and migrate and treat the rest of the body as external environment.

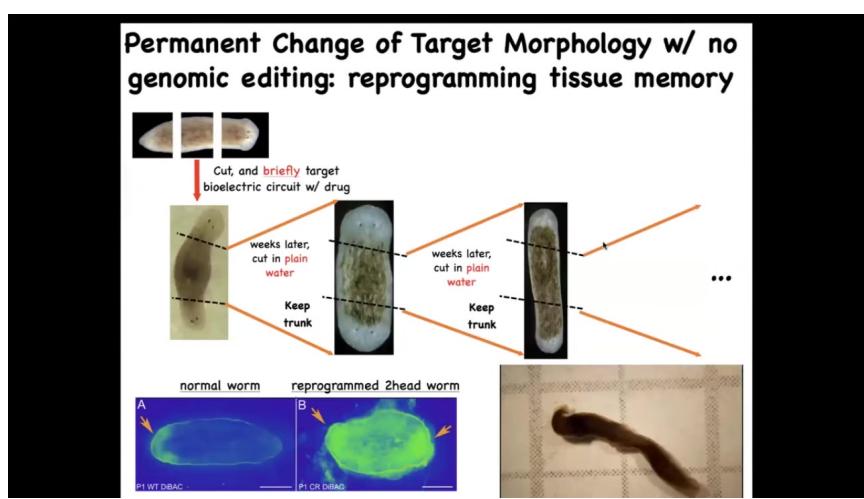


That's how we read these patterns, and we've developed tools to rewrite these patterns. We don't use electrodes, applied fields, waves, or frequencies. What we do is exploit the native interface by which the cells are normally hacking each other's behavior, which are the ion channels on their surface and the gap junctions that allow them to talk to each other. We have ways, mostly appropriated from neuroscience, using optogenetics and drugs to open and close these channels and to control the topology of the network.



When you do this, you can induce whole organs. We can induce inner ears. This is all in the frog model right now, showing inner ears. We can induce hearts, ectopic beating hearts. We can induce forebrain. Here's where the forebrain normally stops in this tadpole. Here we've made this extra-large forebrain.

We can make extra limbs. Here's our five- and six-legged frogs. We can make some structures that don't even belong on this animal, like fins. I don't have time to tell you today about changing the actual shape of an animal to a different species without changing the genome. We have several stories about that.



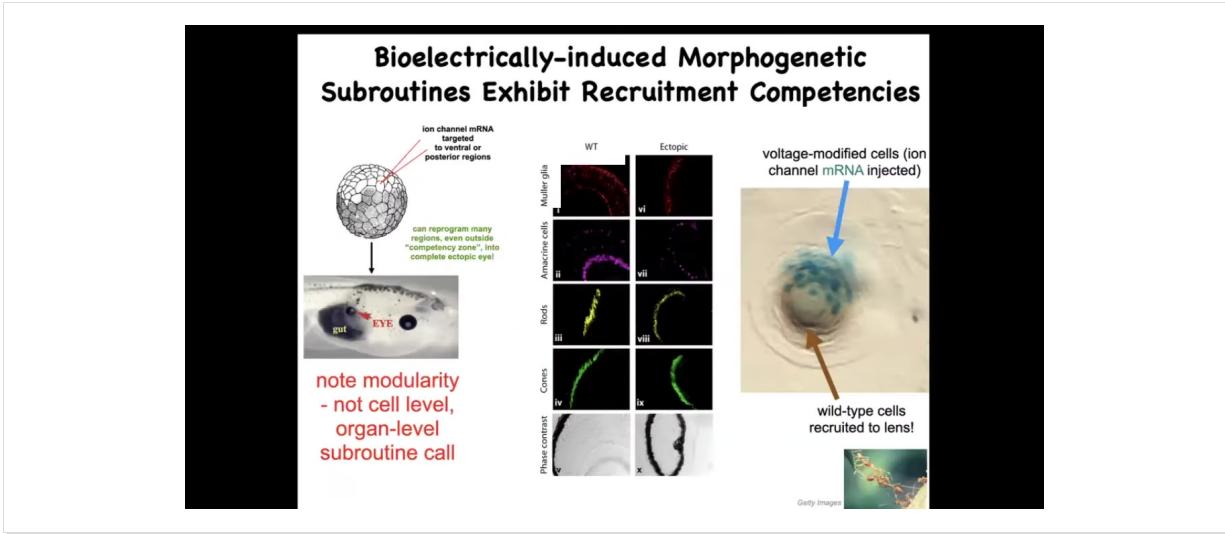
One of the amazing things about this is that these changes can be permanent because we're dealing with a tissue-level intelligence that has memory.

When we take this flatworm, it's got a head and a tail, and we amputate the tail, we amputate the head, we leave just the middle fragment. This middle fragment will normally regenerate 100% of the time to make a normal worm, one head, one tail. How do they know? How do they know how many heads they're supposed to have? We found this electrical gradient that says one head, one tail. We change that gradient to say 2 heads. What you get is this two-headed worm.

This is not Photoshop; you can see the video of these two-headed worms.

The amazing thing is that once you've got this two-headed worm, that's it. From then on, if you amputate the primary head, you amputate this ectopic secondary head; this middle fragment will continue to regrow 2 heads. The genetics are untouched. We never did any genomic editing here. There's no transgenes, no new circuits, no synthetic biology, nothing. All we did was a very brief physiological stimulus that rewrites a pattern memory. This is the way that people read the memories of a brain by looking at the electrical activity. We can now read the memories of the collective intelligence of the body.

What we can see is that once you've rewritten what the system encodes as the correct morphology, it will continue to build this new morphology, as far as we can tell, forever.

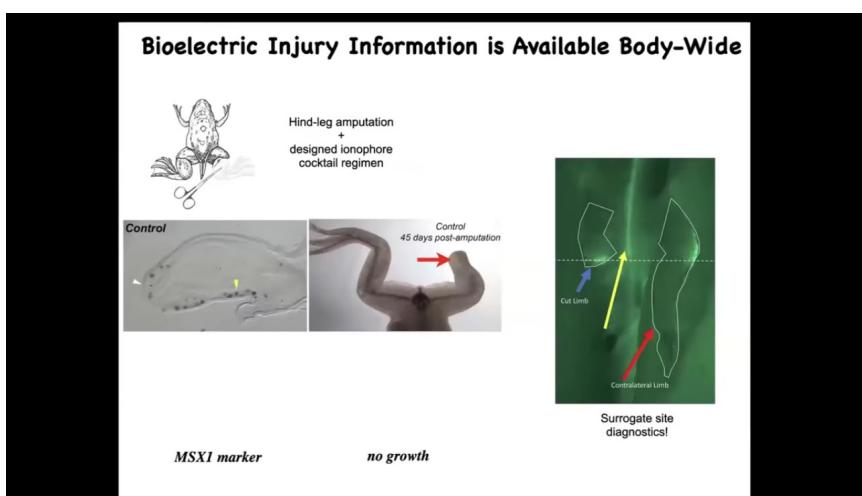


one-headed. Now, I'm going to show you one other example, and then we'll get to some of the biomedical stuff in the biorobotics.

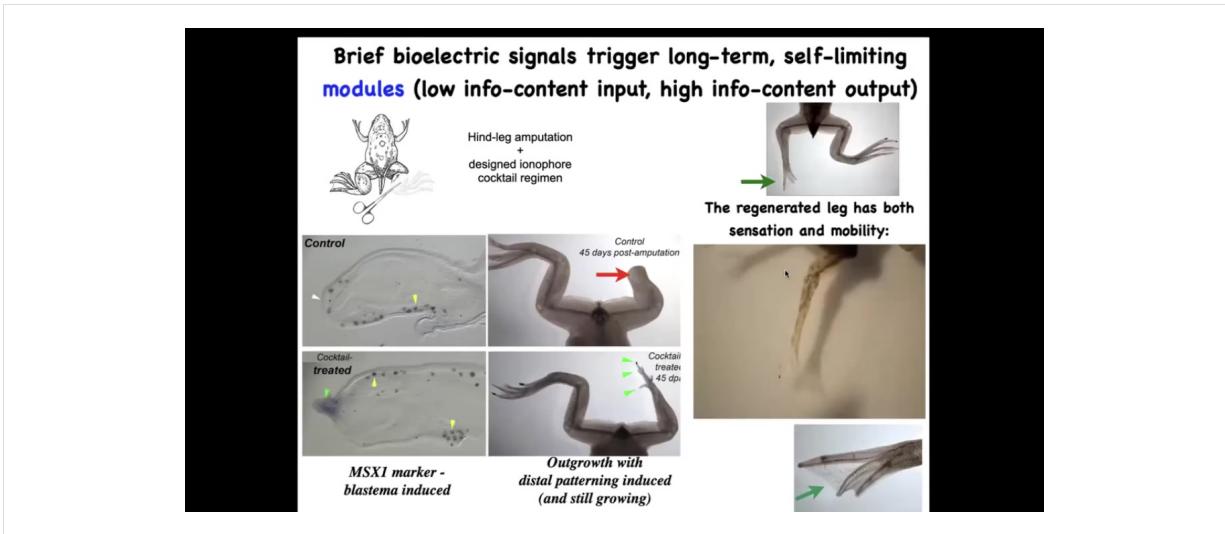
This has to do with organ-level induction. We can inject certain ion channel RNAs that will induce, for example, we can copy that eye spot that I showed you in the electric face. We can put it anywhere in the body. If we put it on the gut, these gut cells happily make an eye, and this eye has all the right lens and retina and optic nerve and all of that. But notice a couple of interesting things. One is that we didn't have to put in all the information about how to make an eye. In fact, we don't have any idea of how to make an eye. What we do know is that there's a high-level subroutine call that says to the local cells, make an eye. And then all the stuff downstream, all the morphogenesis, all the gene expression, all the gradients, everything else is taken care of.

That's a really attractive property, this modularity. But there are other hidden competencies here. For example, this is a lens sitting out in the flank of a tadpole somewhere. The blue cells are the ones that we injected with this potassium channel, and they go to make an eye. But there's not enough of them. There's only a few of them. What do they do? They recruit a bunch of their normal cells, which are not blue. That's how you know we have not manipulated them. These brown and clear cells participate because these cells hijack them. They recruit them towards this goal. This is the mark of a good collective intelligence like ants and termites. When there's a task that's too big for the individuals that have started on it, they will recruit their neighbors to all work together.

Again, we didn't have to orchestrate any of this. This competency is already in the tissue. Discovering these subroutines and discovering the various competencies and the abilities of these networks to restore the correct structure and function with minimal intervention and minimal tweaks are the name of the game here.



Moving to the regenerative part. In a frog, unlike a salamander, if it loses a leg 45 days later, there's nothing. Once you do the amputation, the opposite leg lights up bioelectrically within 30 seconds at exactly the location where the damage took place. The whole body knows about this, and this can serve as surrogate-site diagnostics. Already you can read out where the damage was.



Now what we did was we came up with a cocktail of ion channel drugs that induces regeneration. So here the pro-regenerative genes come on. Within 45 days, you've already got some toes, you've got a toenail, eventually you have a very nice leg. It's touch sensitive and mobile.

And what we found is that one day, 24 hours of a wearable bioreactor made by David Kaplan's lab with the right payload gives you 18 months, a year and a half of leg growth, during which time we don't touch it at all. This is not about 3D printing. This is not about micromanaging what the stem cells do. It is about convincing the cells at the very beginning that they're going to go down the regenerative route and not the scarring route. And that's it. After that, you don't touch it again.

**Practical Applications for Regenerative Medicine**

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



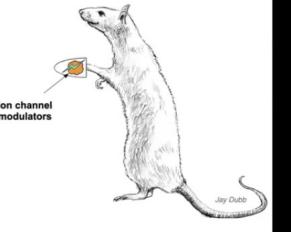
**Wearable Bioreactors for the Study of Limb Regeneration**

**Frog Device Schematic**  
Frog Device Schematic showing a porous hydrogel device with a central cavity and a porous shell. Dimensions: 10 mm height, 10 mm diameter.

**Mouse Device Schematic**  
Mouse Device Schematic showing a porous hydrogel device with a central cavity and a porous shell. Dimensions: 10 mm height, 10 mm diameter.

**Experimental Design**  
A flowchart showing the experimental design: Tunable Inputs (Surficial Design, Drug Choice, Drug Release Profile) lead to the Wearable Bioreactor, which is applied to the Amputation Site. The Wearable Bioreactor leads to Biological Outcomes: Change in Phenotype and Gene Expression (Short Term) and Changes in Tissue Morphology (Long Term). A Feedback loop connects the Biological Outcomes back to the Wearable Bioreactor.

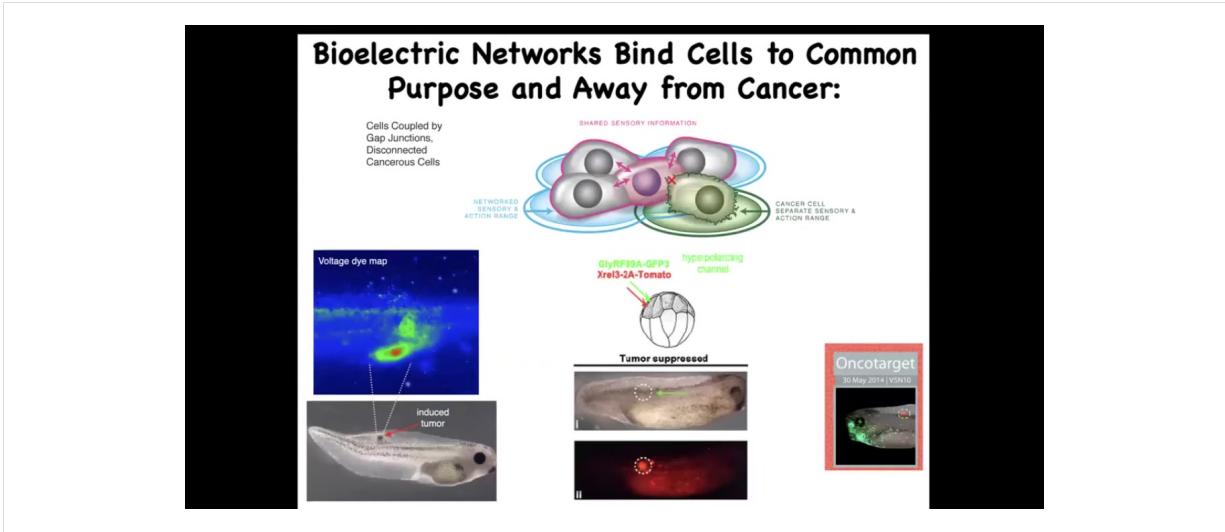
David Kaplan's lab, Tufts BME



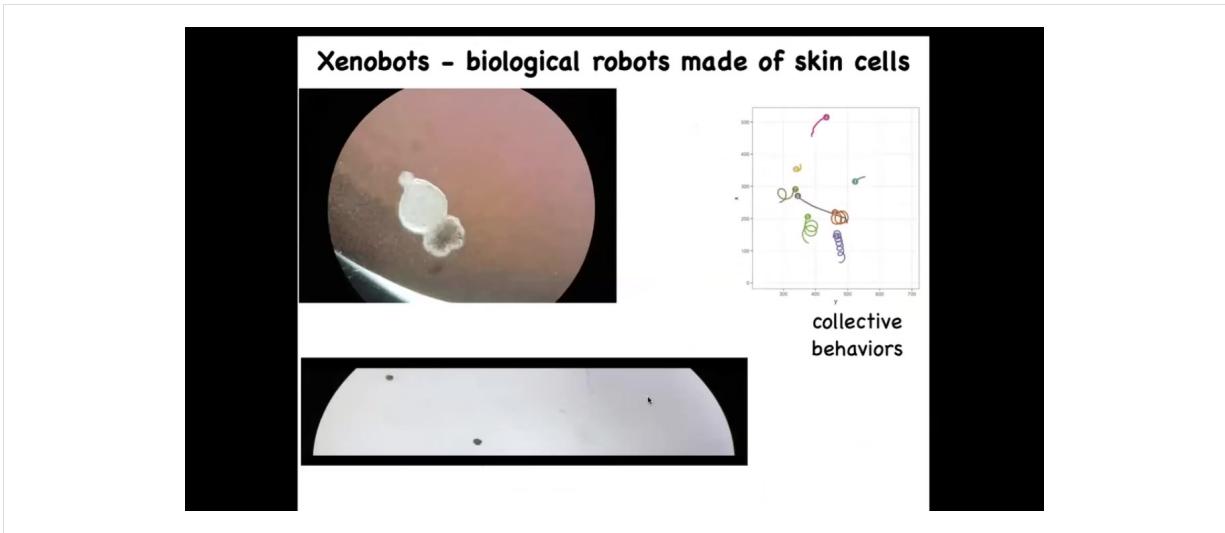
**Ion channel modulators**

Jay Dubb

These are our efforts through Morphoceuticals. David Kaplan and I are co-founders. We're now in mammals and attempting to do this in mice with a bioreactor that controls the wound environment. It has a payload of drugs that try to convince the cells that they need to regenerate.

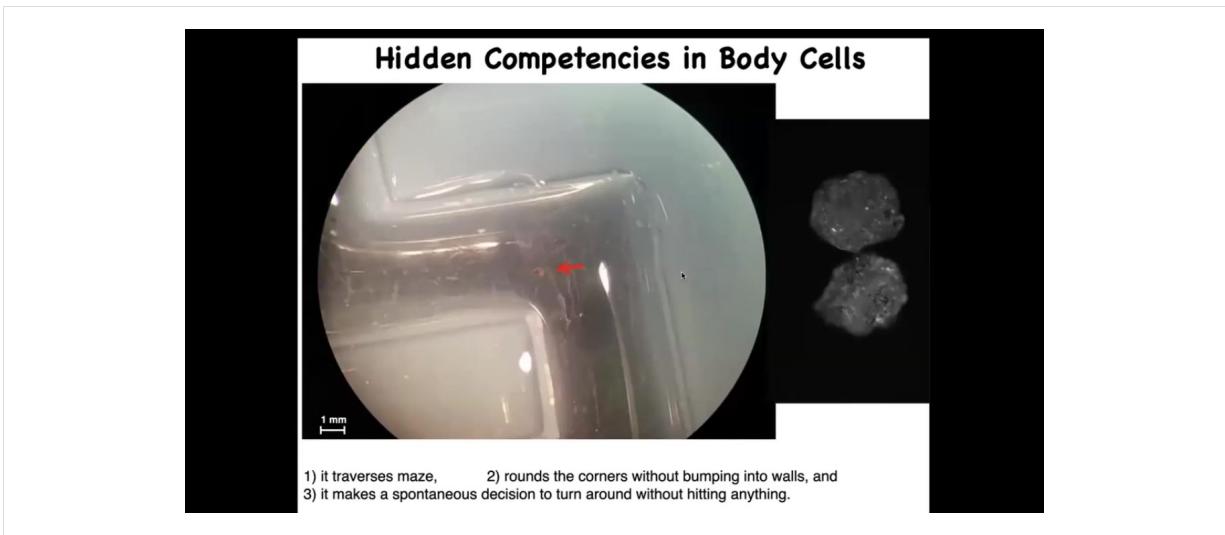


There are also applications to cancer. I showed you this picture a minute ago. Once you start thinking about this as these electrical networks as a way of creating a larger-scale system that can pursue large goals building nice organs instead of little tiny amoeba-level goals of proliferate and migrate. Then what we can do is we can artificially reconnect these cells. When the oncogene tells them to disconnect, we can force them to be in the right electrical state. You see here, the oncogene is labeled in red. This is the same animal. There's tons of it. If you were to sequence this, you would see that there's a nasty KRAS mutation. You would say, there has to be a tumor there. In fact, there is no tumor, because what we've also done is co-inject an ion channel that forces these cells into electrical communication with their neighbors and does not let them detach. There's a cancer therapeutics thing here.



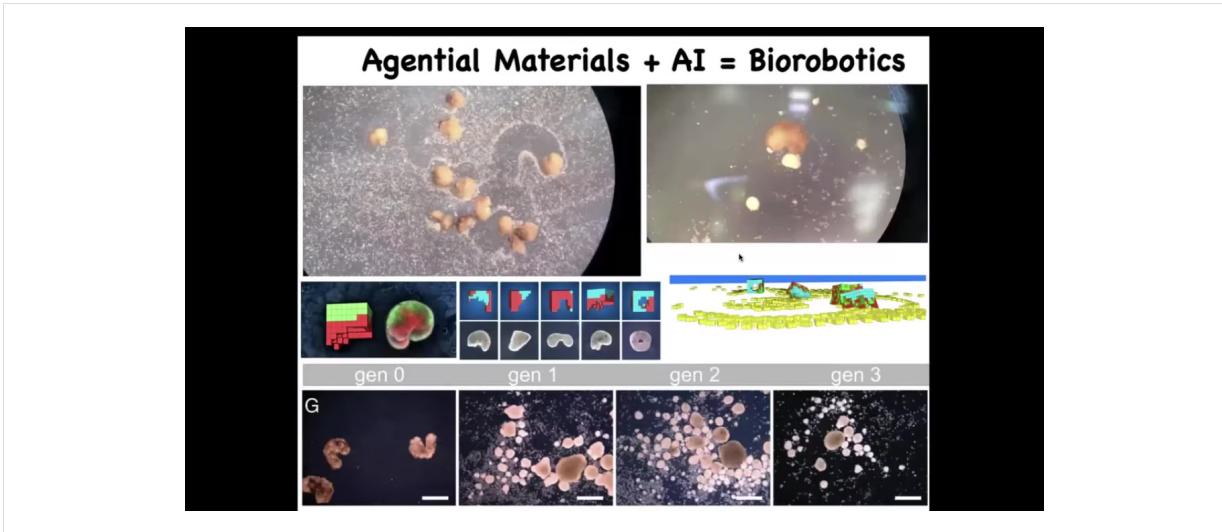
The last bit I want to show you in the last couple of minutes is some synthetic life. We call these Xenobots because they come from the skin cells of the frog called *Xenopus laevis*, and we think this is a bio-robotics platform. This is a collaboration between Josh Vongard at the University of Vermont and myself. Doug Blackiston is the biologist who does this work in my group. Sam Freeman was Josh's student who does the programming.

What you see here is this thing swimming along. This proto-organism self-assembled from cells taken or scraped off of an early embryo. These are just skin. There's no neurons here. There's no brain. This is what the skin will do when left to its own devices, when freed from the instructive influence of the other cells. You can see what happens here. They can move in circles. They can patrol back and forth like this. These are all spontaneous behaviors. They have collective behaviors. They can interact with each other. They can go on these longer journeys. There's great individuality among them. What you're looking at is the answer to the question: what would skin cells do if given a chance to reboot their multicellularity, taken away from these other cells?



Here's one navigating a maze. It's swimming here. It takes the corner without bumping into the opposite wall. For some internal reason, it decides to turn around and go back where it came from. Again, completely spontaneous.

These xenobots have all kinds of calcium activity, which you would normally see in the brain, but again, no neurons, just skin.



One of the most remarkable things about this is that we can now use AI to control their behavior.

Here, what you're seeing is this white stuff, these little particles; these are other skin cells, just loose skin cells that we provided them with.

What Josh and Sam did was do a bunch of simulations that showed us that if we make this little Pac-Man shape, they become really, really good at this task.

What they do is they collect the cells, they push them together into a little ball and they polish that little ball.

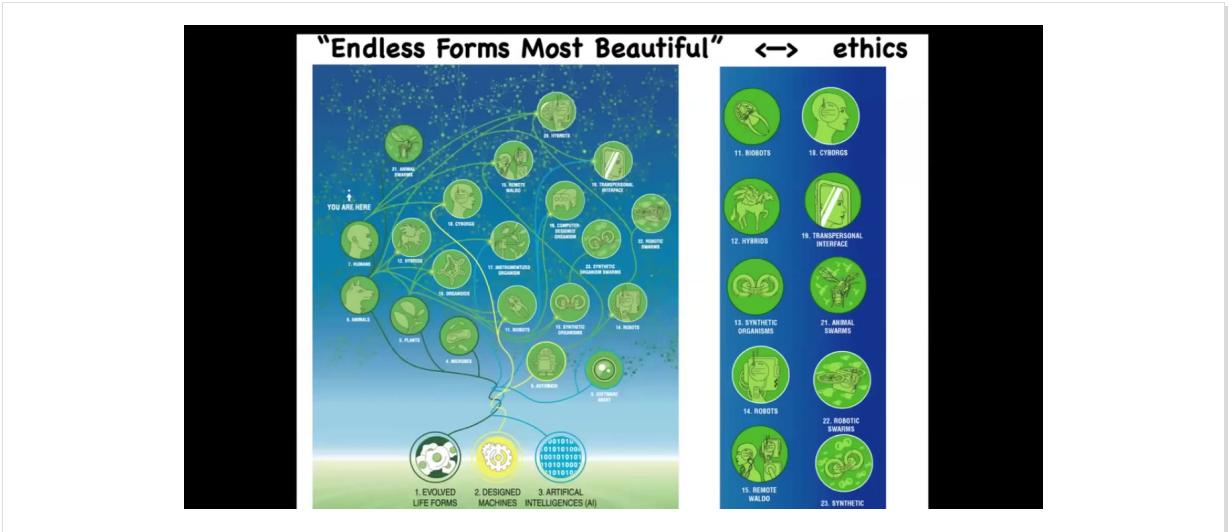
And because they're dealing with an agential material, meaning the cells are not passive particles, they have agendas, they like to do certain things.

Once they're pushed into this ball, that thing matures into the next generation of Xenobots.

Guess what they do? They run around and do the same thing, although not quite as well as the first generation.

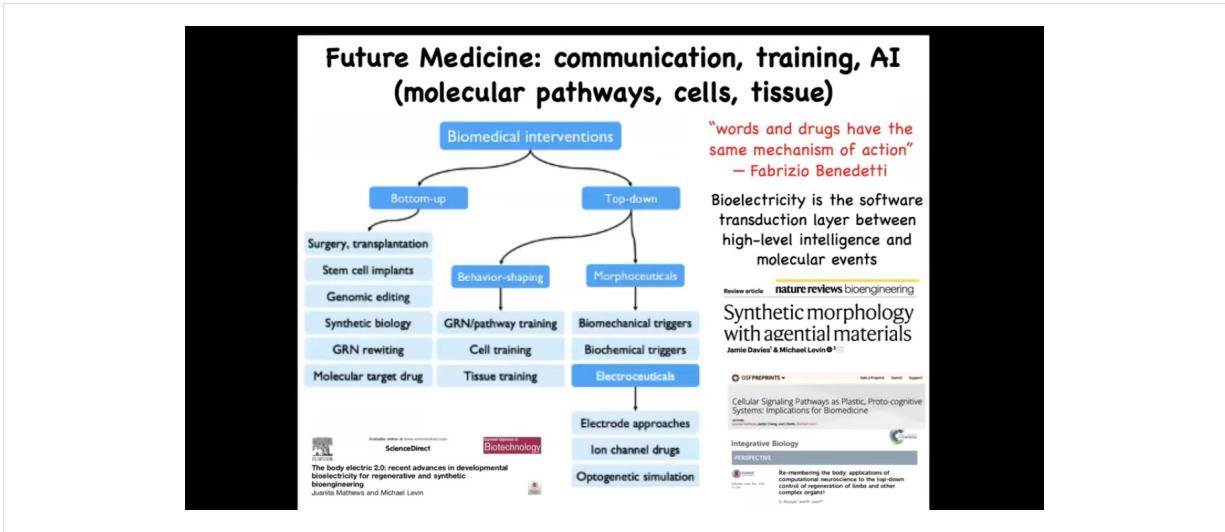
There are a few things we learned from this. One is that we can now use AI to begin to control the native behaviors, the very rich native repertoire, a very surprising repertoire that nobody had any idea that skin cells can implement von Neumann's dream, a robot that goes around and builds copies of itself from loose material it finds in the environment.

There are these amazing competencies and the research program now and the roadmap is to learn to recognize them and then to learn to control them and reprogram them for applications in regenerative medicine, robotics, and ultimately to take what we learn from the intelligence of these cells and use them to create new artificial intelligence.



Because biology is so incredibly interoperable, biology finds a way to live coherently under so many circumstances.

In the future, almost any combination of evolved natural material, designed material, and software at whatever scale is some viable feature. I think in the future, we are going to be surrounded by all kinds of hybrids, cyborgs, and all kinds of things, which are going to require a new kind of ethics for relating to beings that are nowhere on the normal tree of life with us. In other words, we cannot use familiar criteria of where did you come from, meaning evolved or designed, and what are you made of, meaning do you have a brain, does it look human. Those criteria are going to be out the window. Everything that Darwin meant when he said "endless forms, most beautiful" is a tiny corner in this enormous option space of possible bodies and possible minds.



There's some very deep new ideas coming in the space of biomedical interventions rooted in this idea by Fabrizio Benedetti, who studies placebo effects, which is that what you have here is the role of top-down cognitive information on molecular events.

You do that every time you get up out of bed. Your intention of getting up feeds down into the molecular properties of your muscle cell membranes that make you get up. The idea, especially using bioelectricity as this transduction layer, that software layer between the intelligence and the actual execution machinery and AI to manage this gives us a whole range of strategies that are in complement to all the successful mainstream strategies.

## Overview

- Rate limiting step in transformative regenerative medicine is **communication of goals to cellular swarms**. Genetics and big data are not enough.
- Indirect mapping between genes/pathways and outcomes means that bottom-up control is limiting. Fortunately, we have a roadmap for **exploiting the native computation and competencies of the software of life**.
- Bioelectric signaling between cells is the **proto-cognitive medium of the collective intelligence of cells**. Neuroscience and computer science give us tools for top-down control -> applications in birth defects, regeneration, cancer.
- At stake are not only **repair** to standard form and function, but also **synthetic bio-robotics, new GAI, and human augmentation**.

To summarize what I've said, the rate-limiting step in truly transformative technologies is the communication of goals to cellular swarms, not the molecular details, but the algorithms of life. We have a roadmap for exploiting the native computation and the competencies of living networks. Bioelectrical signaling is a crucial proto-cognitive medium of that collective intelligence. We now have tools that appropriate all kinds of deep concepts from neuroscience to help us understand this in other contexts. At stake are a huge number of transformative applications.

**Thank you to:**

**Post-docs:**  
Léo Pio-Lopez - heterochromatin, interface scaling, aging, and bioelectricity  
Douglas Blackiston - brain/ECM interface, bioelectricity, synthetic living bodies  
Samira Samani - computational modeling of heart development  
Devon Davidian, Evgenia Shmulev, Nirusha Murugan, Celia Herrera-Rincon, and AlSun Tseng - limb regeneration  
Vishav Pai - voice growth, gene expression induction and repair  
Jiaming Wang - bioelectricity of cancer  
Santosh Manicka - cognitive and dynamical systems models of morphogenesis  
Tilak Doshi - bioelectric imaging  
Nestor Oviedo, Wendy Beane, Johanna Bischof - bioelectricity of planarian regeneration  
Patrick McMillen - bioelectric imaging and transitions to multicellularity

**Ph.D. Students:**  
Alessio Nino - chimerization in morphogenesis and cognition  
Franz Kuehling - cognitive neuroscience modeling applied to pattern homeostasis  
Fallon Durant -  $V_{mem}$  and pattern memory in planarian regeneration

**Undergraduate Students:**  
Maya Emmons-Bell, Kelly G. Sullivan - non-genetic cross-species morphogenesis  
+ many other undergraduate students working in our lab over the years

**Technical support:**  
Rakela Colon, Jayati Mandal - lab management  
Eric Swanson - vermetide animal husbandry  
Emma Lederer - Kenyan behavior  
Joan Lemire - molecular biology

**Collaborators:** Allen Center members +  
Alexis Pietak - computational modeling of bioelectricity  
Jordan Miller - bioelectric simulations and AI applications for biology  
David Kaplan -  $V_{mem}$  and human MSC differentiation, regenerative sleeves  
Florencio Ornatnetto - optical approaches to regeneration  
Giovanni Cimatti - mathematical models of pattern regulation  
Vitaly Volpert, Chris Fields - mathematical models of pattern regulation  
Hannah Ceder - gene expression, regeneration and evolutionary learning  
Sara Walker, Karl Fritton - top-down causative models  
Don Ingber, Richard Novak, V. J. Koosmon, J. H. Dungan - mammalian bioengineering  
Jack Tuszynski - cell biophysics and drug discovery

**Model systems:** tadpoles, planaria, zebrafish, slime molds, human cells, and chick embryos

**Funding support:** TWCF, JTF, DARPA, Paul G. Allen Frontiers Group, Sloan Foundation, NIH

**Illustrations:** Jeremy Guay @ Peregrine Creative

**Disclosures:** Morphochemicals, Fauna Systems, Astonishing Labs



I'll end here by thanking the postdocs and the students who did all the work, our funders, and most of all, the animal model systems, because they do all the heavy lifting. There are some disclosures I need to make. Thank you very much.

**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).