

A LECTURE COMPANION

---

# **"Bioelectric networks: an interface to engineering with the agential material of life"**

---

Michael Levin

Recorded on March 1, 2025

## About this document

This document is a companion to the recorded lecture "*Bioelectric networks: an interface to engineering with the agential material of life*", recorded on March 1, 2025. You can watch the original lecture or listen in your favorite podcast feeds — all links are on the page [here](#).

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

---

## Lecture description

This is a ~1 hour talk (given at the Departmental Seminar series at the Virginia Tech-Wake Forest School of Biomedical Engineering and Sciences) on the use of a cognitive approach to bioengineering and regenerative medicine, in which morphogenesis is the behavior (in anatomical space) of a collective intelligence of cellular swarms. I describe our work that takes advantage of the bioelectric interface to re-specify the targets of morphogenesis.

---

## 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).

## Bioelectric Networks: an Interface to Engineering with the Agential Material of Life

Michael Levin  
Allen Discovery Center at Tufts

<http://www.drmichaellevin.org/>  
<http://allencenter.tufts.edu/>



ALLEN DISCOVERY CENTER at Tufts University

Computer-designed Organisms  
TUFTS UNIVERSITY | UNIVERSITY OF VERMONT

WYSS INSTITUTE

HARVARD

Thank you. I appreciate the opportunity to share some ideas with you all. What I'm going to talk about today is this notion of an agential material for bioengineering. We will talk about bioelectric mechanisms. But the important part of this talk is not so much the fact of this biophysical mechanism, but what it tells us about information processing and, in particular, diverse intelligence that goes all the way down to the biological material. If you want to download any of the papers, the datasets, the software — all the things I talk about are here at this website.

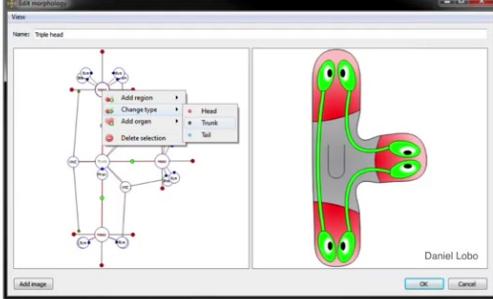
## Main Points:

- Most problems in biomedicine and bioengineering boil down to control of morphogenesis and collective decision-making among cells, and will not be solved by genomic (i.e., hardware-based) approaches alone.
- Biology uses a **multi-scale competency architecture** of nested problem-solvers in diverse problem spaces; we (and evolution) can exploit this architecture and the **bioelectrical interface that cells and tissues use to behavior-shape each other toward complex structure and function**
- Tools now exist to **read and write memories into this software layer between genome and anatomy**, with applications in birth defects, regeneration, cancer, and **synthetic bioengineering**
- **Engineering with agential materials: using tools from behavior and cognitive sciences to exploit collective intelligence of cells in transcriptional, physiological, and morphogenetic spaces**

Discussing the context that most problems in biomedicine and bioengineering really are fundamentally about the control of shape, and in particular about the collective decision-making among cells. I will argue that this is not going to be solved by hardware-based approaches such as genomics and molecular biology, but it's going to require thinking that comes from the cognitive sciences. Because biology uses a multi-scale competency architecture. We consist of nested problem-solving agents that work in all sorts of unconventional problem spaces. Engineers and evolution exploit this architecture, and it's this bioelectrical interface that allows cells to behaviorally shape each other towards creating and maintaining complex anatomical form.

We now have tools to read and write the pattern memories of the software layer of life, that layer that sits between the genome and the anatomy. There are numerous applications in birth defects, regeneration, cancer, and synthetic bioengineering. In particular, the main message of today is that when you are engineering with agential materials, you can use tools from the behavioral and cognitive sciences to exploit the collective intelligence of cells as it operates in these problem spaces.

## Biomedical Endgame: Anatomical Compiler



**Why we need it:**

- Birth defects
- Traumatic injury
- Cancer
- Aging

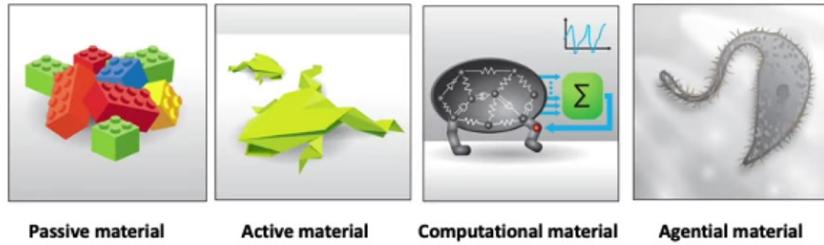
Problems of information processing

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

So let's work backwards from what I think is the end game of biomedicine. Someday, you will be able to sit down in front of a computer and draw the animal, plant, organ, biobot, any kind of living system that you want, draw it at the level of anatomy, not at the molecular biological level. And what the system will do then is to compile your goal, your description, into a set of stimuli that would have to be given to cells to get them to build exactly this, in this case, a nice three-headed flatworm. If we had this, all of these things would go away. Birth defects, traumatic injury, cancer, aging, degenerative disease. All of this would go away because we could tell cells to build and rebuild all the complex, healthy structures that we need. So how would we do this? And in particular, why don't we have anything like this? Molecular biology and genetics have been going very strong for decades.

## Engineering with Diverse Substrates

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



because we are made of competent parts with agendas and problem-solving capacities

I'm going to argue that here's what's happening: this anatomical compiler, this thing that's going to allow complete control over growth and form, is nothing like a 3D printer. It is not about micromanaging the details as you would when you were engineering with a passive material such as Legos, where you're in charge of every single piece. It's a communications device because in biology we are not dealing with active matter or even just computational matter. We are dealing with an agential material because we're made of competent parts with agendas and problem-solving capacities. The goal is to transmit the needs of the regenerative medicine worker or bioengineer to the collective intelligence of cells so that is what they build.

## We All Used to be Single Cells!

Lacrymaria = 1 cell  
no brain  
no nervous system

high competency  
at cell-level  
agendas

*Lacrymaria olor*      video by Charles Krebs

Normal Embryonic Development

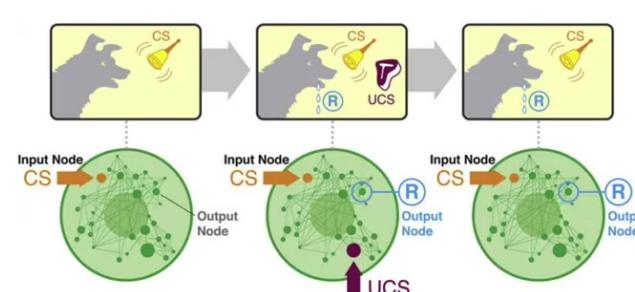
René Descartes

So here's what we are made of. This is a single-cell organism, a lacrimaria, a free-living organism, and this just illustrates what the individual cells can do. This thing has no brain, no nervous system, but you can see it's very competent at handling its tiny little agendas in physiological space and metabolic space.

All of us started life this way. We all began life as an unfertilized oocyte, a little BLOB of chemistry and physics. Then eventually through this incredible process of embryogenesis, we became one of these things or even something like this.

So this is René Descartes, a complex metacognitive human who will make comments about not being a machine and being something more than physics and chemistry. This process is slow and gradual. We all have this origin. That means we need to understand how this works. How did we get here?

## Collective Intelligence Below the Cell Level



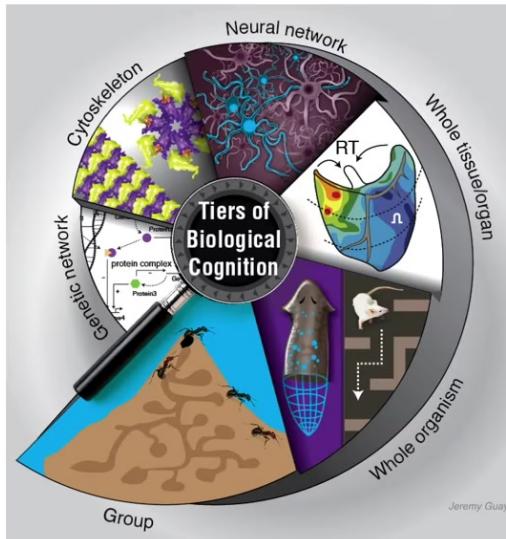
Biomedicine:  
- drug conditioning

**Article**  
Learning in Transcriptional Network Models: Computational Discovery of Pathway-Level Memory and Effective Interventions  
Surama Biswas <sup>1,2,\*</sup>, Wesley Clawson <sup>1,2</sup> and Michael Levin <sup>1,3,4,5</sup>  
**iScience**  
Gene regulatory networks exhibit several kinds of memory: quantification of memory in biological and random transcriptional networks  
MDPI | CellPress | Open Access  
Patrick Erickson



Now, I will point out that there is even intelligence below the cell level, meaning that it's not just that the individual cells are competent, but even below the cell level, there are pathways, gene regulatory networks and other biochemical pathways that you may think of as mechanical, as deterministic, hardwired kinds of things. But in fact, if you pay attention, it turns out that they are capable of learning. They show about six different kinds of memory, including Pavlovian conditioning. Not neurons, not brains. You don't even need the cell. It's the molecular networks themselves that can learn. And in our lab, we are currently building devices to take advantage of the learning capacity of cells and do things like drug conditioning and other ways to control the molecular pathways, memories, and the expectations that they have and how they deal with future stimuli and so on.

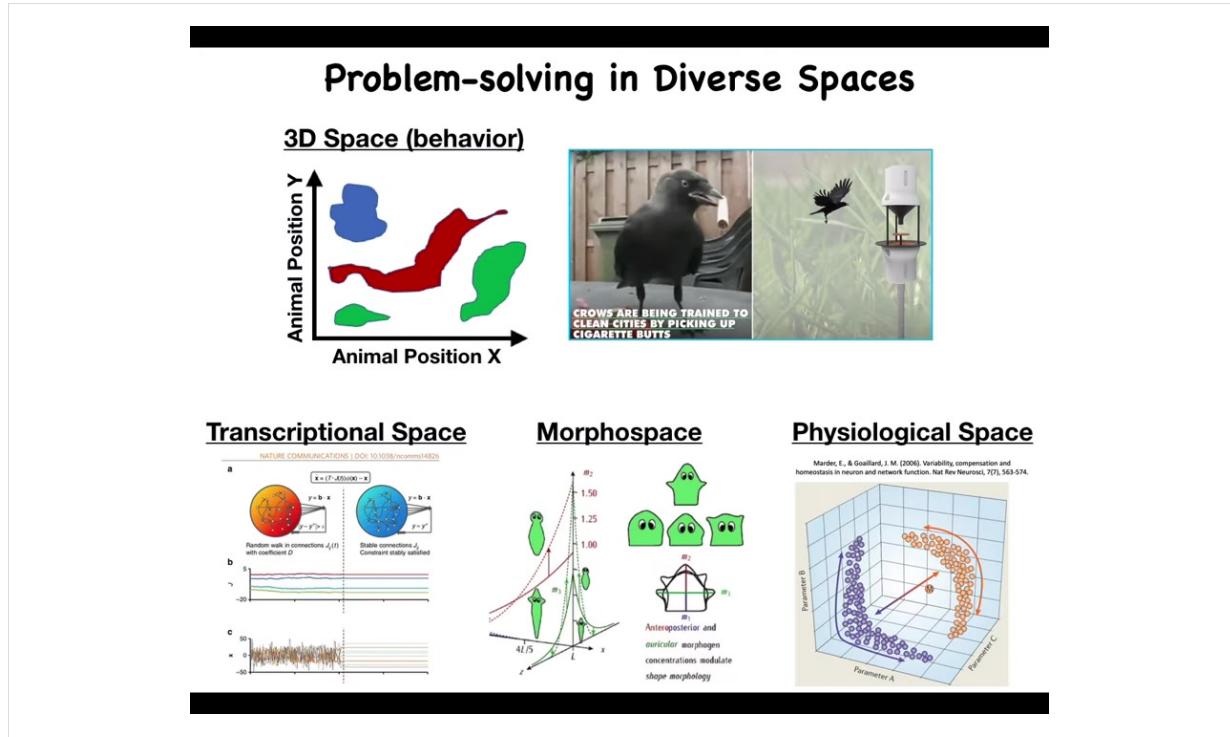
## Nested Cognition, not Merely Structure



### Multi-scale Competency Architecture

each level of organization solves problems in its own space (morphospace, transcriptional space, physiological space, 3D behavioral space, etc.) using some of the same tricks, at various levels of sophistication

So here's our architecture. We are not just nested in a structural sense. Everybody knows molecules, subcellular components, organs, organisms, swarms, and so on. But it's not just that. It's the fact that each one of these is itself a competent agent navigating things like anatomical space and transcriptional space and so on. And so there are these autonomous competencies at every stage.



We as humans are primed to recognize intelligence in three-dimensional space, specifically medium-sized objects moving at medium speeds, navigating three-dimensional space. We observe some kind of problem-solving capacity. We're actually not good at all at recognizing the same thing when it occurs in other spaces.

For example, cells navigate the space of all possible gene expressions. That's transcriptional space; that's a pretty high-dimensional space. They navigate physiological spaces, and they navigate anatomical morphospace, the space of all possible anatomical configurations. This is the one we're going to talk the most about today. I need to show you what I mean when I say they navigate the space, because you might have the idea that morphogenesis is the result of a kind of mechanical working out of the rules of biochemistry. While that's not exactly false, it's a very limiting picture that does not give access to the opportunities of regenerative medicine.

## Where do Anatomies Come from?

- DNA specifies proteins; whence Anatomy?
- how do cell groups know what to make and when to stop?
- How to convince cells to repair/rebuild?
- **how far can we push shape change?**  
Engineers ask: what's possible to build given default genome?

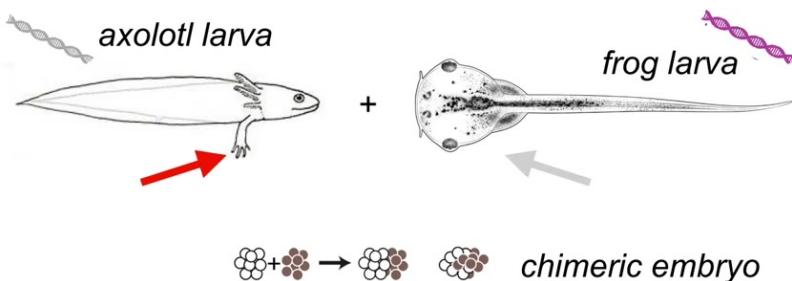
And I'll show you what I mean. Let's first ask ourselves where anatomies come from. This is a cross-section through a human torso. You can see all the different organs, the structures, everything is in the right place relative to each other. Where is this pattern coming from?

We all start life like this. This is a set of embryonic blastomeres that are going to build this. We need to understand where this pattern is specified. You might think it's in the DNA. That's the most common thing people say: it's in the genome.

But we can read genomes now. We know that this is not directly in the genome at all. What the genome specifies are proteins, the tiny molecular-level hardware that every cell gets to have. But the final outcome is not directly readable from the genome any more than the shape of a termite mound or a spider web is readable from the genomes of these creatures.

We need to understand what the software layer here is. How do the cells know what to make? How do they know when to stop? If something is missing or damaged, how do we convince them to repair and rebuild? As engineers, we would also like to know what else is possible. They build this, and what else can we get the exact same cells to build?

## Genetic Information is not Enough



We know both genomes;

Will **frogolots** have legs?

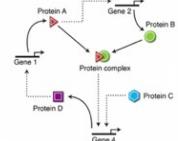
And we want to focus on this genetic information. Here's just one example. Here's a baby axolotl, and these juvenile axolotls have little legs. Here is a tadpole of a frog. They do not have legs at this stage.

Now, in my group, we make something called a frogolotl, which is an early chimeric embryo that has some frog and some axolotl cells. We have the genomes. We have the axolotl genome. We have the frog genome. They've been sequenced. They've been read.

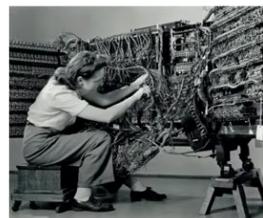
Now I ask you a simple question. Given this, can you tell me if a frogolotl is going to have legs or not? And if it is going to have legs, are those legs going to be made entirely of axolotl tissue or some of the components will be frog cells? And the answer is there's no way to tell. We have no way of knowing from this genetic information what's actually going to happen here.

## 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

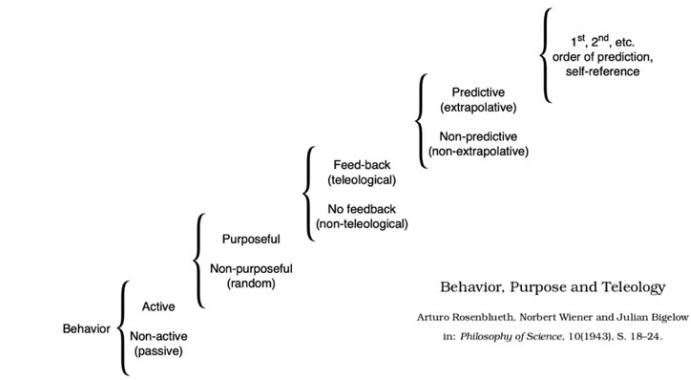


biomedicine has been largely at the hardware level of analysis and control, we do not understand higher-level information processing in biology (agency)

While we are, as a field, very good at manipulating molecules and cells, and making these pathways, who interacts with whom, we're quite good at this. We're really a very long way away from predictive control of form and function. If somebody's missing an arm, how do we convince the cells to regrow it or to explain the form to begin with? We have a lot of difficulty with that still. I'm going to argue that's because biomedicine has been largely focused on the hardware level of analysis. It's like computer science and information technology in the 1940s and 1950s. Here she is, she's reprogramming this computer. In those days, you had to interact with the hardware. To reprogram it, you had to move wires around. This is what's going on today: all the most exciting advances in molecular medicine, genome editing, pathway rewiring, protein engineering, all these things are down at the level of the hardware. If I told you that to switch from Photoshop to Microsoft Word on your laptop, you have to get out your soldering iron and start rewiring, you would laugh. Why is that funny? Because we know what's happened with the field of computer science. It's spawned the age of information technology by realizing that if your hardware is good enough, and I'm going to argue that biological hardware is definitely good enough, then you have access to something very powerful, which is reprogrammability and intelligence of the material. You don't have to rewire it to get new components.

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

- William James



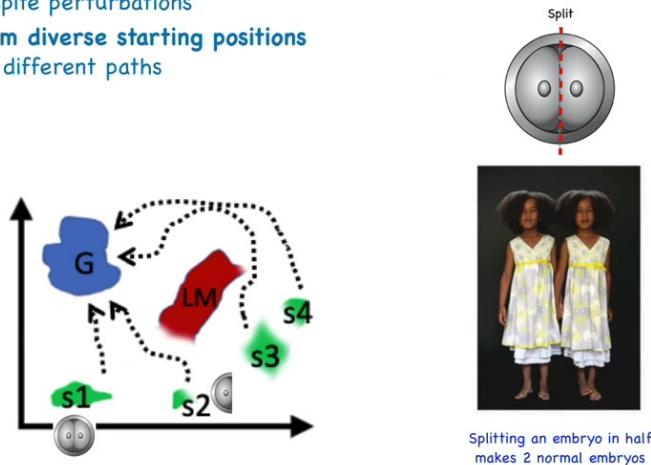
Here we have to have a definition. I've already used this word a couple of times. I'll use it many more times. Here's what I mean by intelligence. It's William James's definition, which is the ability to reach the same goal by different means. Notice it doesn't say anything about brains. It doesn't say anything about what kind of goals or what space, what problem space your goal is in. It's the ability to navigate some problem space to get your goal met despite various perturbations, barriers, novelty.

This is part of a much bigger program we call diverse intelligence, which seeks to understand the whole spectrum, from passive matter up through the different kinds of cybernetic control mechanisms and ultimately things like 2nd order human metacognition.

The next thing I'm going to do is show you some examples of intelligence in the cellular collective. What are your cells doing that is actually intelligent? The first thing we know about development is that it's reliable. A human embryo almost always results in exactly the right species-specific target morphology. The reliability of it and the complexity of the outcome are not why I call it intelligence. This is important. The intelligence is not merely complexity. It is not merely doing something reliably. All of that can be done mechanically. That is not what I'm talking about. What I'm talking about is creative problem solving. Here's an example.

## Same anatomy, from different starting states

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



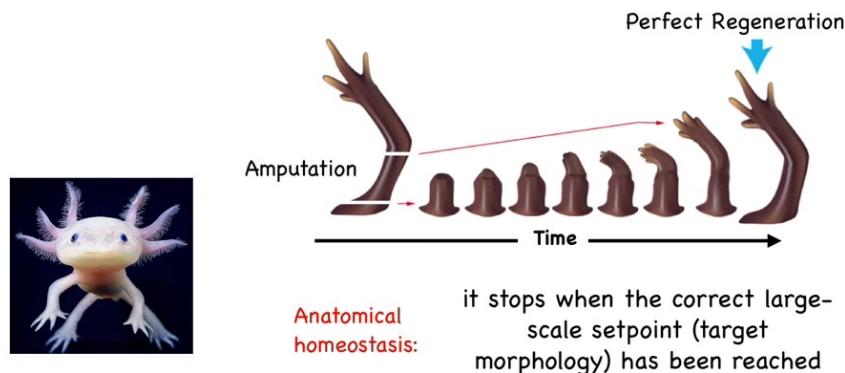
Splitting an embryo in half makes 2 normal embryos

If I take this same embryo and cut it in half or into quarters or even eighths or beyond, I don't get half bodies. I get perfectly normal monozygotic twins, triplets.

That's because this system doesn't mechanically go from its normal starting position to its goal in anatomical space, but it can get there from other positions. You can do all kinds of things with this embryo. It's not limitless. There are things you can do to mess it up. But there are many things you can do where the system recovers. Half of it is missing. It can tell what's missing, how much, create the missing pieces and get to where it's going, even avoiding these local maxima in this space.

## Same anatomy, from different starting states

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



Some animals are able to do it throughout the lifespan. Here's an axolotl. This amphibian can regenerate its limbs, its eyes, its jaws, portions of the heart and brain. When these guys inevitably bite their legs off, as they often do, for each other, whether that happens down here or whether that happens up here or up here, the cells will build exactly what's needed, and then they stop. The most amazing thing about regeneration is that it knows when to stop. When does it stop? When the correct limb has been completed. That's when it stops.

What you're seeing here is a homeostatic loop. You're seeing a system that can tell that it has been deviated from the correct position in the anatomical space. It works really hard to get back to where it's going, and then it stops when it's met its goal. That is what I mean by a goal-seeking system. I don't mean that it's metacognitive and it knows that it has a goal. I just mean that it is competent at pursuing a particular goal, and it will do so even when it has been deviated from it.

## Regeneration is not just for “lower” animals



The human liver is highly regenerative

Every year, deer regenerate meters of bone, innervation, and skin

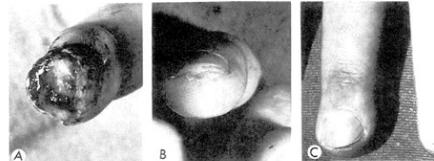


Fig. 2. (A) Amputation of finger tip in 5-yr-old girl. (B), (C) Twelve weeks after accident.

Human children below 7-11 years old regenerate fingertips

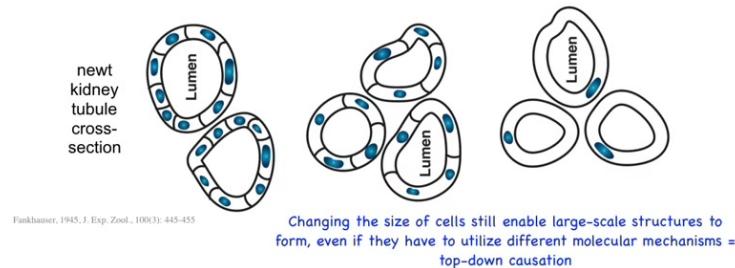
This kind of regeneration is not just for so-called lower animals. The human liver is highly regenerative. Every year, this animal, deer, will regenerate huge amounts of bone, vasculature, innervation as it regrows the antlers, up to a centimeter and a half of new bone per day. It's an amazing regenerative rate.

Even human children can regenerate their fingertips. Up until a certain age, they will regenerate them.

We have this problem of regenerating, which shows us that creatures can restore missing parts, but it's even much more impressive than that.

## Same anatomy, despite perturbations

- get to the same outcome
  - despite perturbations (external and internal)
  - from diverse starting positions
  - via different paths



Here's one of my favorite examples. This is a cross-section through a kidney tubule in a newt. And you'll see normally there's 8 to 10 cells that form a ring structure around a space in the middle. That's the lumen.

One thing you can do is create a situation where there's too many copies of the genetic material at the one cell stage in the egg. This was done back in the 1940s. When you do this, first of all, you still get a normal newt. So that's wild. It doesn't matter how many copies of your genome you have. You still get a normal animal.

The second thing that happens is the cells get very large. The cells actually adjust size to keep a certain ratio to the size of their nucleus. But there's a problem, because if you have these giant cells and you still try to do 8 to 10, you're going to get an enormous newt. That's not what happens. The newts are normal sized. How can that be? That's because these large cells can create the same structure by using way fewer cells. So now there's maybe only four or five cells that are making the same lumen. The most impressive thing is when these cells get truly gigantic. This is like 6N polyploid newts that have many more copies of their genome. The cells get really big. One cell wraps around itself to create the same structure.

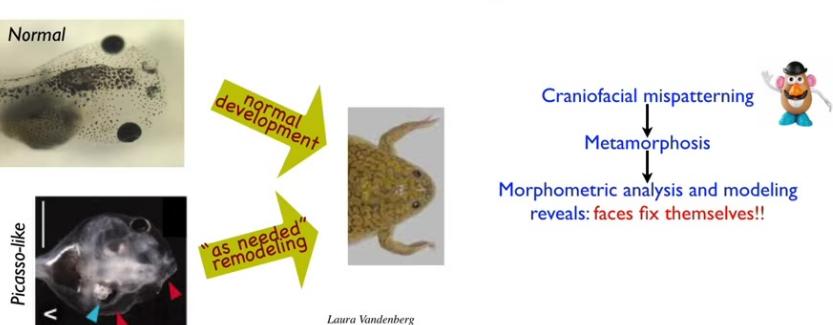
A couple of interesting things to note here. One is that this and this are completely different molecular mechanisms. This is cell-to-cell communication. This is cytoskeletal bending. So what you have here is a situation where the goal of making a particular anatomical structure—navigation in those spaces of anatomical shapes towards the correct path—can make use of different molecular components. That's

actually the schematic for many types of IQ tests where they give you a task and they say, find creative ways to use the components you already have to solve this task. In order to do this, the cells literally use a different molecular mechanism.

Now just think of what this means. If you're a newt coming into the world, what can you count on? We know you can't count on your environment, and we know you might get damaged, and so it makes sense that you regenerate. But it can't even count on its own parts. You don't know how many copies of your genome you're going to have. You don't know how big your cells are going to be. Your own parts are unreliable, and yet you have to have the goal met. This is what I mean by intelligence. It's literally creative problem solving. It's the ability to handle novel situations by using the tools that you have in new ways.

Slide 17 of 57 · Watch at [17:18](#)

## Intelligent Problem-solving in Morphospace



*Normal*

*Picasso-like*

*Laura Vandenberg*

normal development

"as needed" remodeling

Craniofacial mispatterning

Metamorphosis

Morphometric analysis and modeling reveals: faces fix themselves!!

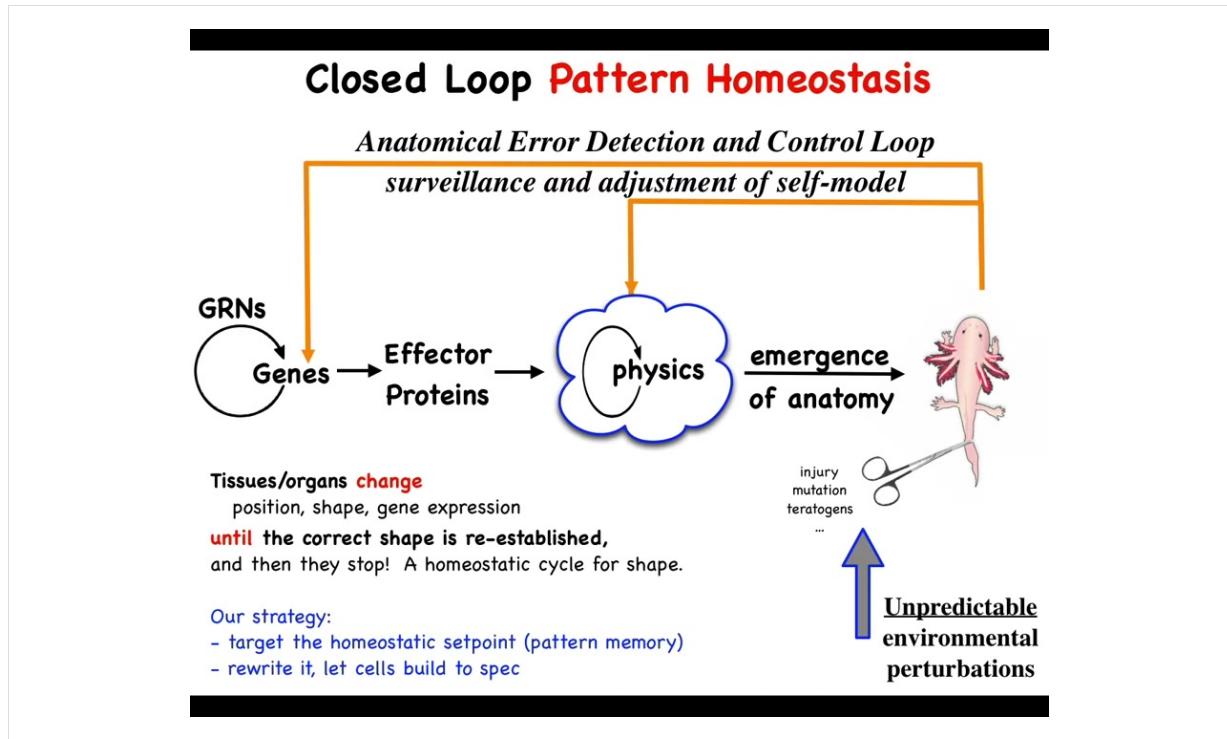
Genetics does not specify hardwired rearrangements: it specifies a system that executes a highly flexible program that can recognize unexpected states and take corrective action.

What is the cognitive glue that binds cells together toward complex goals and problem-solving capacities, and how do we exploit that interface?

Another example like this is a simpler example that sets up the next thing I'm going to talk about, which is that these tadpoles need to become frogs, and in order to do that, they rearrange their face. Their eyes move forward, their jaws move, everything moves. You might think that this is a hardwired process that basically, each of these organs just moves in the right direction the right amount, and then you get a normal frog. We

decided to test that, and this is very critical for estimating the intelligence of anything. You have to do perturbative experiments. You can't just guess. So what we did was we scrambled the position of all these organs, like a Mr. Potato Head doll. We put the eyes on the back of the head, the mouth is off to the side, everything is shifted. These animals, it turns out, make quite normal frogs, because all of these organs will move in novel paths. It is not hardwired. They will move in novel paths until they get to a correct frog face, and then they stop. That tells us that the genetics does not specify a bunch of hardwired rearrangements. It specifies a problem-solving agent. It specifies a system that executes a flexible program that recognizes unexpected starting positions and takes corrective action.

That brings up two questions with a common answer, which is, first, what are the mechanisms—you can call them cognitive glue—that allow individual cells to work towards large-scale outcomes? No individual cell knows what a finger is or how many fingers you're supposed to have or where the eye goes. Individual cells don't know this, but the collective absolutely does. I've shown you a number of examples where you can see they know exactly where they're going in that anatomical space. So what enables that to happen? And how do they know what the goal is? It's one thing to form a collective that knows more things than its parts, basic collective intelligence, but how does it know what the correct frog looks like?



This is what we've been working on for some years as a model system. Collective intelligence of groups of cells as a model system for scaling up the competencies of individual subunits into something greater that operates in a new space. This is fundamentally a problem of cognition and collective intelligence.

What we have to do is in addition to this standard scheme that is in all of your developmental biology textbooks, you have gene regulatory networks, they make proteins, then there's this process of complexity and emergence. And then out comes something like this. We have to add to this these homeostatic loops that I just talked about, which if there is damage that gets you away from this state, injury, mutation, teratogens, then this whole system has to kick in and do something different in order to return you to that state.

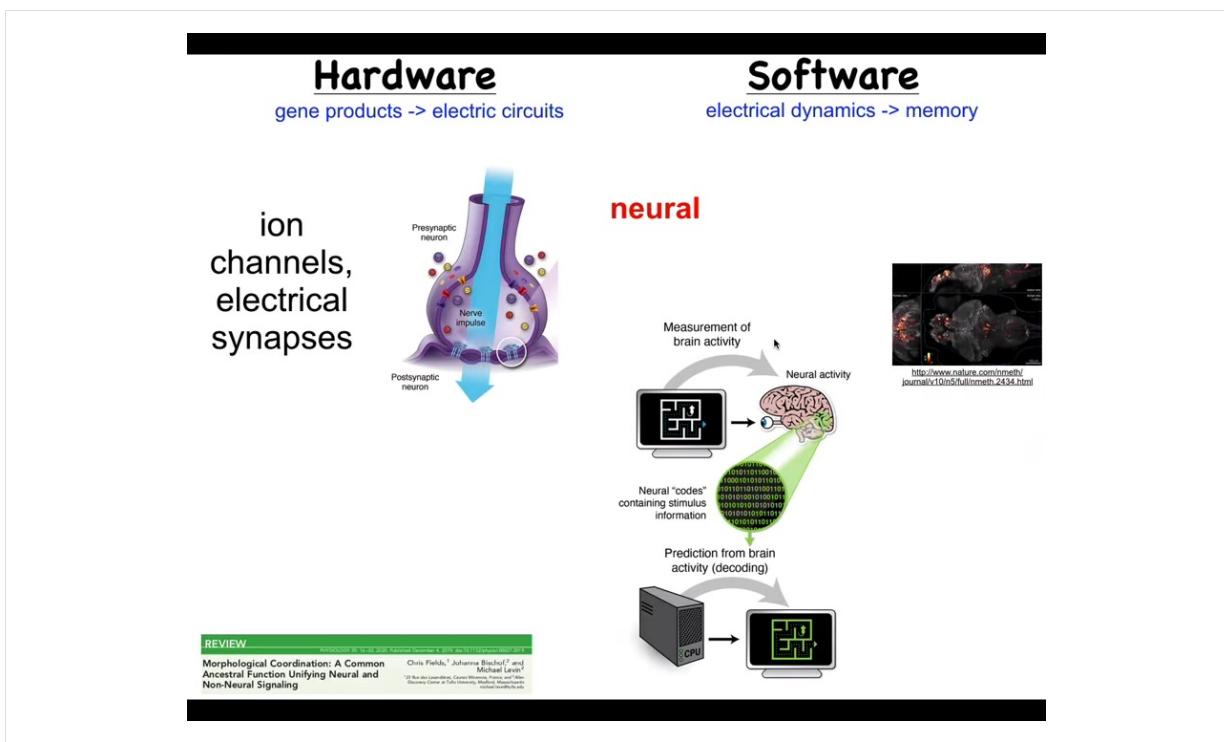
This unusual way of looking at it, the standard scheme is just open loop feed forward emergence. These things happen at the molecular level and then eventually something happens. But this error calculation and correction model is something different.

What it enables is the following implications, which we have been testing for some decades now: if you want to change the way the system works, you don't necessarily have to change the hardware. What if you just changed the set point? What if you changed the memory that tells these cells what it is that they're trying to build. In your thermostat, if you want it to be at a different temperature, you don't need to change the thermometer, you don't need to change the plant that heats or cools, you can just change the set point.

Could we do this? That means that if we could decode where that goal information was being stored, the memory of this collective intelligence, we could rewrite it and in fact leave the hardware entirely alone. We wouldn't have to change the genes. That's a very convenient feature because this doesn't work backwards. With a few exceptions, in the general case, if you want to make large-scale system-level changes here, we have no idea how to figure out what we would have to change down here.

This inverse problem is what's limited regenerative medicine until now. It would be nice to be able to avoid that, change the set point, and let the system do what it does best, which is to build to that set point. In order to think about how to do that, how could we decode what are the memories of a group of cells?

Slide 19 of 57 · Watch at [21:53](#)

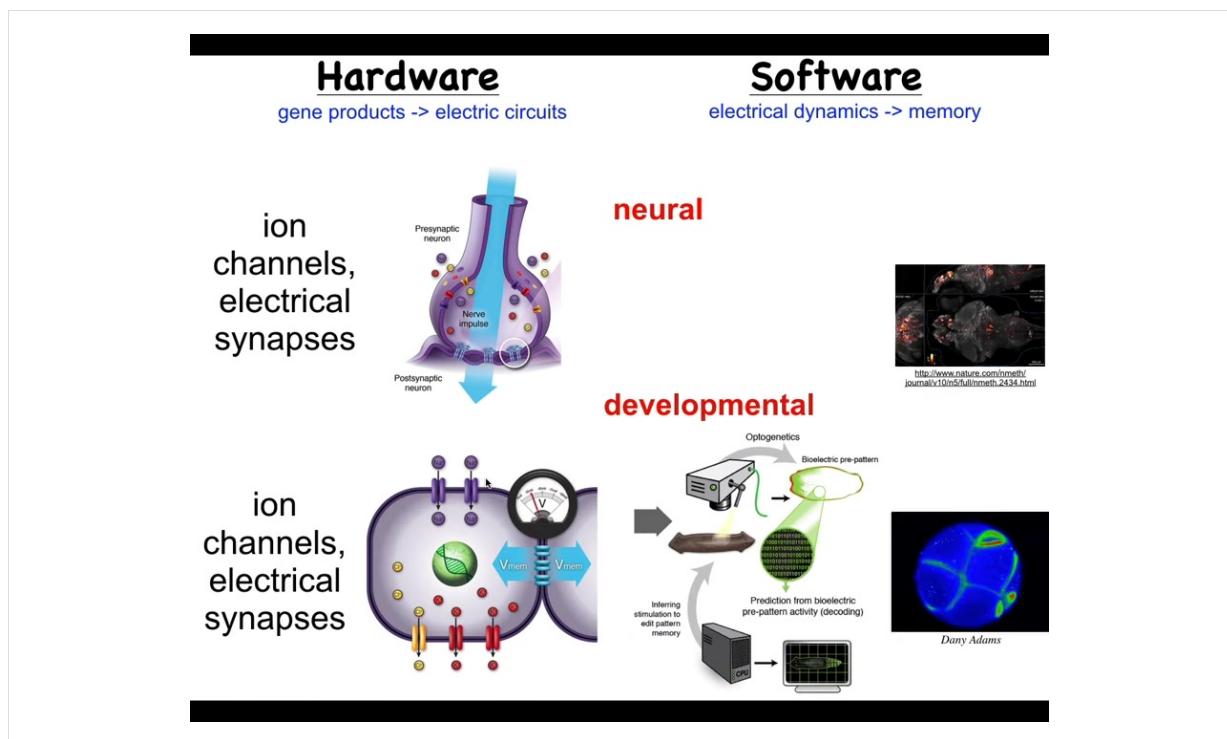


We looked at a non-controversial example where cells store memories and act towards them, and that's the nervous system. So we know that in the brain, you have hardware like this. You have cells expressing ion channels. They acquire certain electrical properties. They pass them on, or not, to their neighbors using these electrical synapses. And that's the hardware, that network runs a kind of physiological software.

Here's an example. You're seeing this group did this amazing video of a zebrafish brain as the fish thinks about whatever it is that fish think about. And that suggests this program of neural decoding. So neuroscientists try to read this electrophysiological information and decode it. Because the idea is that the animal's memories, goals, preferences, and everything else are in some way encoded in this electrical activity.

So that's the paradigm, that if you have a network like this of subunits communicating electrically, the collective, which is more than the sum of its parts, stores its memories in the dynamics of the electrical signaling. It turns out, and you can see the details in this paper, that this is an ancient system.

Slide 20 of 57 · Watch at [23:03](#)



It's not about brains at all. Evolution discovered this around the time of bacterial biofilms. Every cell in your body has these ion channels. Most cells have electrical synapses between their neighbors. Now we can run the same kind of program.

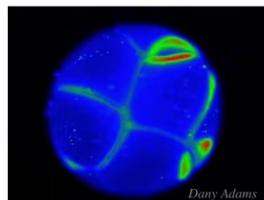
Neuroscience outside of the brain and nervous system: we can ask, what is it that the body was thinking about before nerve and muscle evolved? These networks have been around for a really long time. What they were thinking about is shape — anatomical

shape. Just as we can try to decode the electrical activity of the brain, we can try to decode the electrical activity of the body to understand how the cells know what to do in embryogenesis.

Slide 21 of 57 · Watch at [23:46](#)

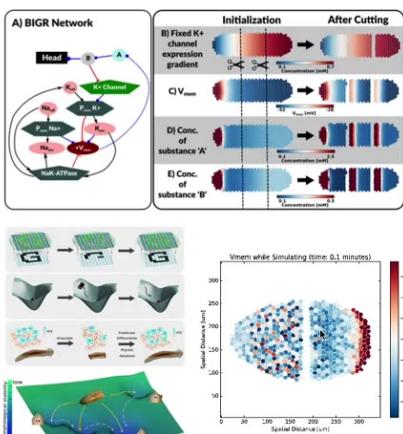
## How we detect and model bioelectric patterns:

**Characterization** of endogenous voltage gradients - direct measurement and correlation with morphogenetic events



A fluorescence micrograph showing a cluster of cells. The nuclei are stained blue, and there are green fluorescent puncta, likely representing viral particles or specific cellular markers. The background is black, making the fluorescent signals stand out.

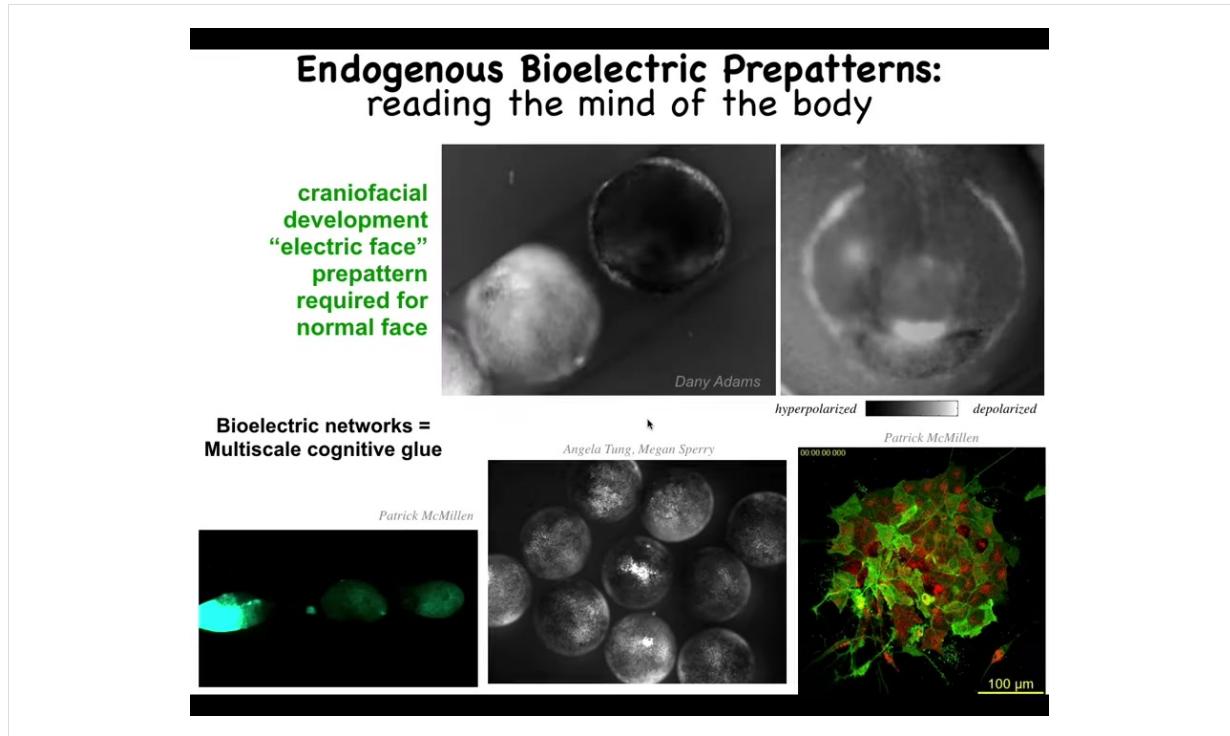
**Quantitative computer simulation:** synthesize biophysical and genetic data into predictive, quantitative, often non-linear models



Alexis Pietak

We developed some of the first tools to do this. It started out with voltage-sensitive fluorescent dyes. We used existing voltage-sensitive dyes to track the outcomes of embryogenesis, of regeneration, of cancer suppression here. These are individual cells. This is an early frog embryo. You can see all the cells having these electrical communications with each other. You can see what's happening. There are some fast voltage changes. There are some very slow voltage changes here. We do a lot of imaging like this.

We also do a lot of quantitative simulations. For example, we can try to understand how the gene regulatory networks that express channels result in excitable tissues that can form patterns and how these patterns repair after damage, as pattern completion in both neuroscience and machine learning. We have these architectures where if you're missing part of the input, the system can reconstruct what goes there. We do a lot of computational modeling like that.

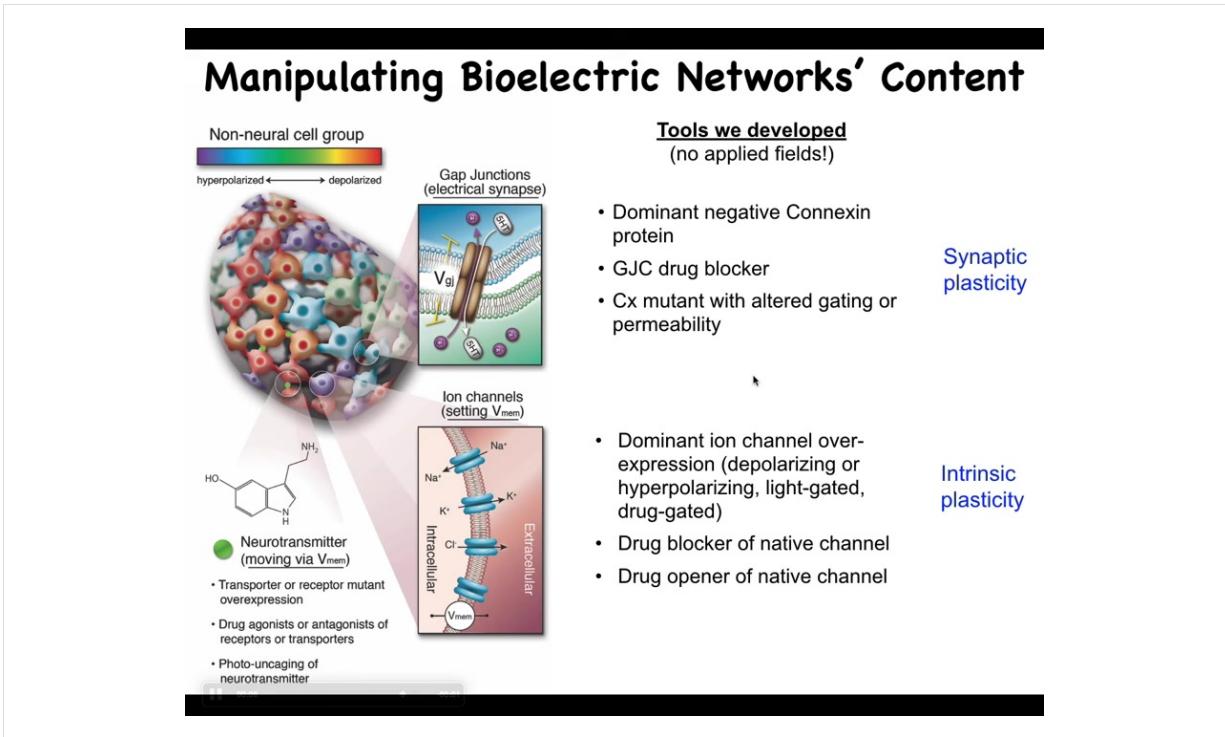


Let me show you a few examples of these bioelectric pattern memories. We call this one the electric face. This is one frame of this time-lapse video of an early frog embryo putting its face together. There's a lot going on. What you can see in this frozen frame here is that there's a bioelectric pre-pattern that already tells you what the face is supposed to look like. Here's where the eye is going to go, here's where the mouth is going to go, here are the placode structures out here. You can already see the pre-pattern long before the genes start to turn on to regionalize the face.

This bioelectricity allows the cells to cooperate to figure out what the large-scale structure is supposed to look like, meaning how many eyes it is supposed to have and where the eyes go. Not only is it a cognitive glue binding individual cells into an embryo, it actually binds individual embryos into a group. Here you can see if this guy is poked, these two find out about it in short order. If this one has damage, all the ones around it find out. It's a multi-scale communication mechanism.

Here it is on the single-cellular level. You can see some neurons, but you can also see a lot of activity out here that is developmental bioelectricity. This kind of electrical communication can produce wholes out of parts.

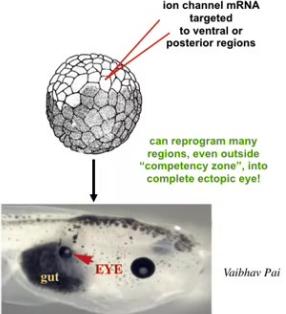
Now, observing and analyzing these patterns is all well and good, but even more important is being able to change them and to rewrite them.



We took all the tools from neuroscience. We do not use any applied fields, waves, radiation, or electromagnetics. There are no frequencies.

What we do is manipulate the actual interface, the bioelectrical interface that cells use to control each other. That means we can control the ion channels. We can open and close them using drugs and optogenetics. The same for the electrical synapses, the gap junctions. We can open and close them. That allows us to control what in neuroscience would be called synaptic plasticity and intrinsic plasticity. Those are the tools we have: molecular biology to control the bioelectric patterns inside groups of cells. So now, can we rewrite the information content of these networks? Can we put false memories into the collective intelligence of cells?

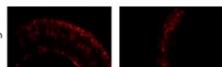
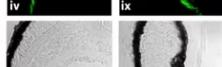
## Bioelectric Interface to Agential Material: Self-scaling morphogenetic subroutines



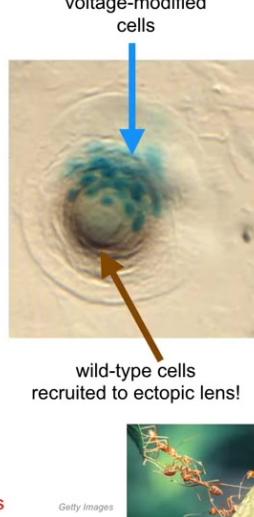
ion channel mRNA targeted to ventral or posterior regions

can reprogram many regions, even outside "competency zone", into complete ectopic eye!

Vaibhav Pai

	Native eye	Induced Eye
Muller glia		
Amacrine cells		
Rods		
Cones		
Phase contrast		

Developmental Modules because morphogenetic goals



voltage-modified cells

wild-type cells recruited to ectopic lens!

Getty Images

Here's a couple of examples. One example is that in an early frog embryo, we target some cells that are going to become gut. We target them with ion channel mRNA, encoding some potassium channels. And what they do is they reproduce that little eye spot that I showed you in the electric face, but they reproduce it somewhere else. If you put that eye spot, that little voltage spot here in the gut, then what you've told the cells is build an eye here, and that is exactly what they do. These eyes have all the same lens, retina, optic nerve, all the stuff they're supposed to have.

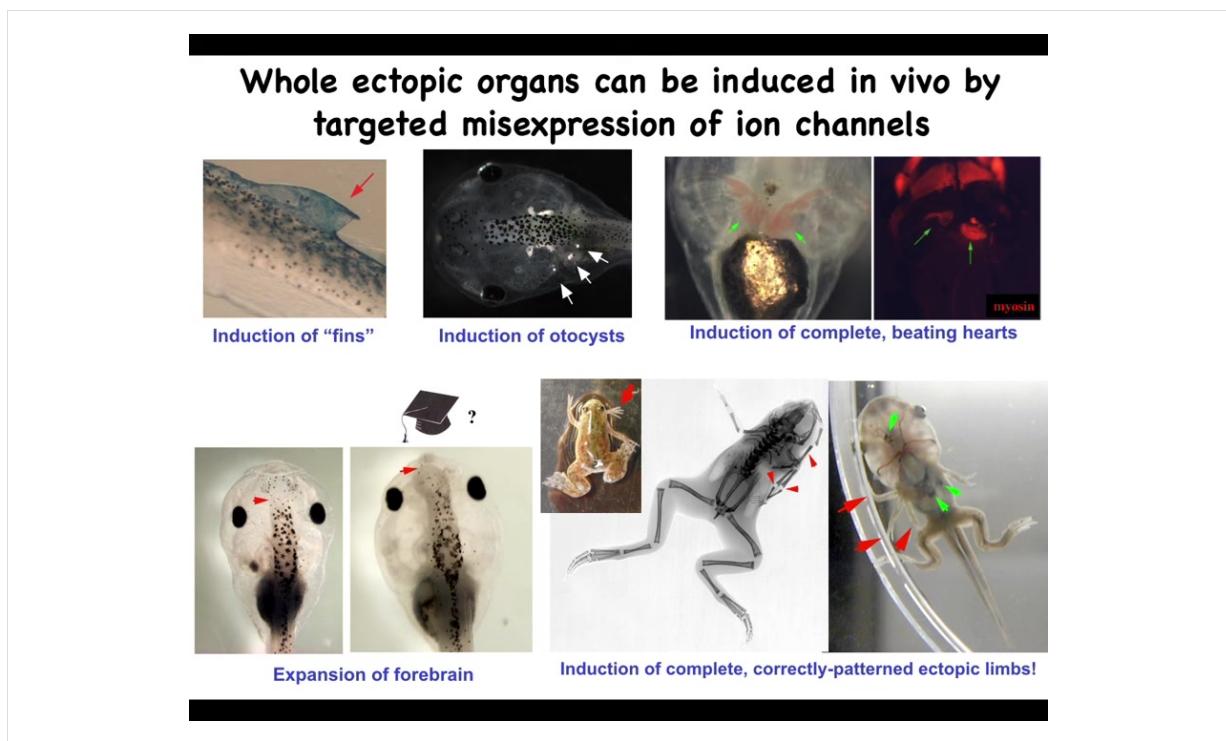
Notice a couple of interesting things. First of all, this tells you that the bioelectronics is instructive, meaning that it is actually telling the cells what to do. Second, notice that it's instructive at a very modular high level. We did not tell individual genes to turn on and off. We did not tell individual stem cells what to become. We told a large group of cells, be an eye. The signal encodes something a very large-scale state. It is not something that individual cells or molecules know. It is something that only the collective knows. And that works precisely because the system can trigger this very complicated cascade of molecular events that are needed to build an eye with a very simple signal. All we said is make an eye. We didn't say how big, what the structure of the eye is, and so on. That only works because the whole system as a whole is made up of modules that are good at following specific goals. They just need to know what those goals are.

We could spend a whole hour talking about just this experiment; a lot of interesting things here. For example, if you only inject a few cells like this, these blue cells, they

will actually go out and recruit all this other stuff, which we never touch. Those cells are completely untouched by us, but they will recruit them. It's a kind of secondary induction because this system knows how many cells it takes to build a proper eye, and it knows there's not enough of them, and so they will take over the task of recruiting their neighbors. We did not have to worry about that.

You see, when I say it's an agential material, this is what I mean. It has competencies we don't have to worry about. We don't need to know how to turn on which genes to make an eye. We don't need to know how many cells exactly have to be part of an eye. The material does all of this for us, and you don't find this in LEGOs or pretty much anything else that humans build with, at least to date. The material is incredibly competent at these kinds of things.

Slide 25 of 57 · Watch at [29:26](#)

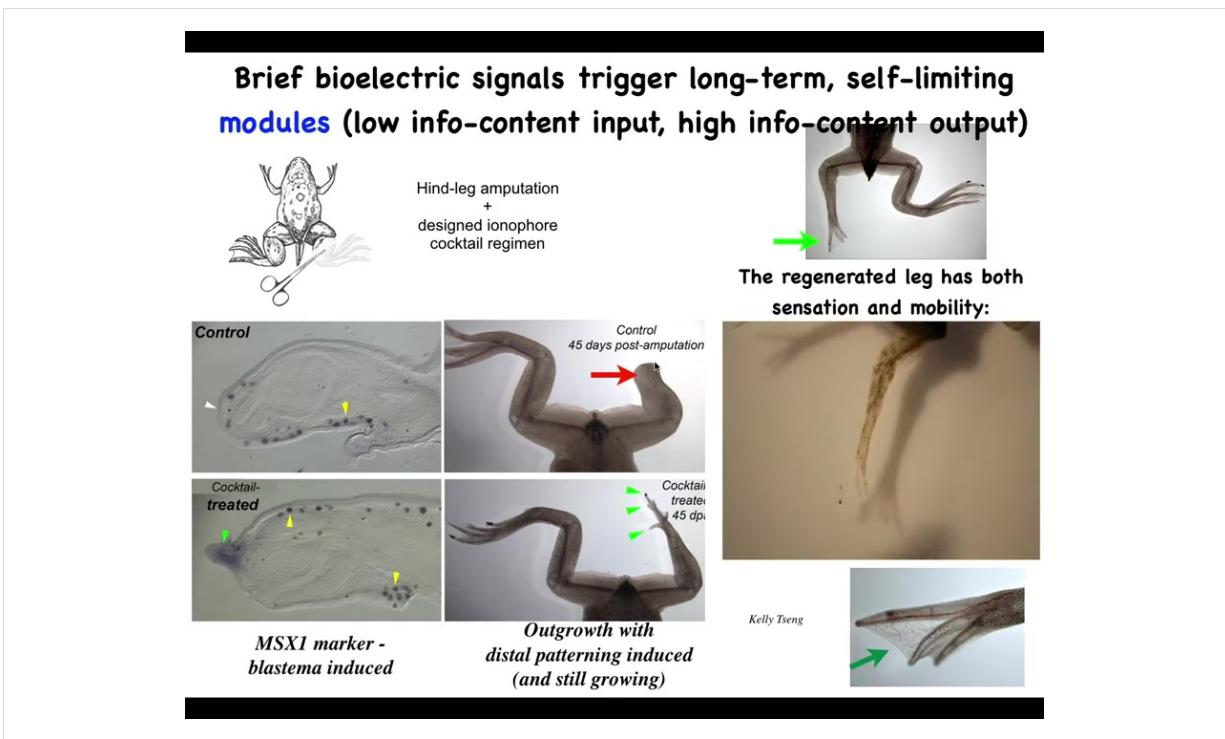


It's not just about eyes. We can make inner ear. We can make ectopic hearts. We can make ectopic forebrain, ectopic limbs. Here's our six-legged frog. We can also make fins, which is interesting. Tadpoles don't have fins. I'll get into that momentarily because you can make structures that don't actually belong on that species.

By showing that we can induce various organs in this way by providing bioelectrical states that encode meaning to the surrounding cells, we can use this for a regenerative medicine program.

For example, frogs normally do not regenerate their legs as adults. 45 days later, there's nothing.

Slide 26 of 57 · Watch at [30:08](#)



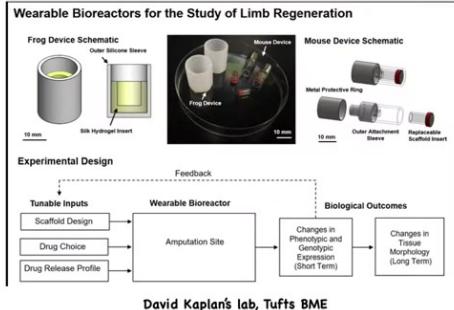
We created a cocktail that triggers a signal that causes the leg to regenerate. So within about 24 hours, you get this pro-regenerative blastema. You then eventually, by 45 days, you've got some toes, you've got a toenail, and eventually a pretty respectable leg that is both touch sensitive and motile. And our record at this point is 24 hours of exposure to the intervention, followed by a year and a half of leg growth, during which time we don't touch it at all. The reason is that we are not micromanaging it. This isn't scaffolds, it's not stem cell therapy, it's not 3D printing. We don't touch it at all at that time. The goal is during the first 24 hours to get across the idea that you should be moving towards a leg building region in that anatomical space, not the scarring region.

## Practical Applications for Regenerative Medicine

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

PLOS ONE

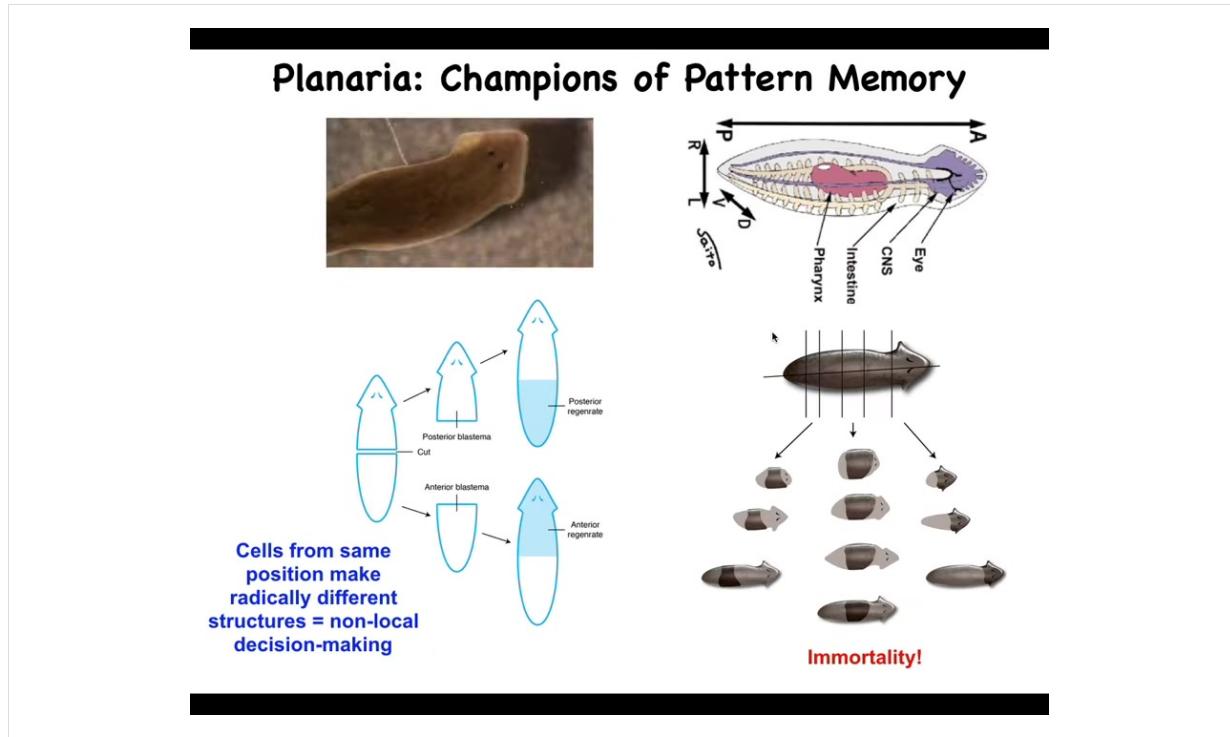
RESEARCH ARTICLE  
A Tunable Silk Hydrogel Device for Studying  
Limb Regeneration in Adult *Xenopus Laevis*  
Anne Golding<sup>1</sup>, Justin A. Guy<sup>2</sup>, Celia Herrera-Rico<sup>3</sup>, Michael Levin<sup>1</sup>, David  
L. Kaplan<sup>1\*</sup>



**Disclosure:**  
Morphoseuticals Inc.

And so here I have to do a disclosure. So Morphoseuticals is a company that Dave Kaplan and I spun off, which seeks to take that kind of technology into mammals and eventually human patients.

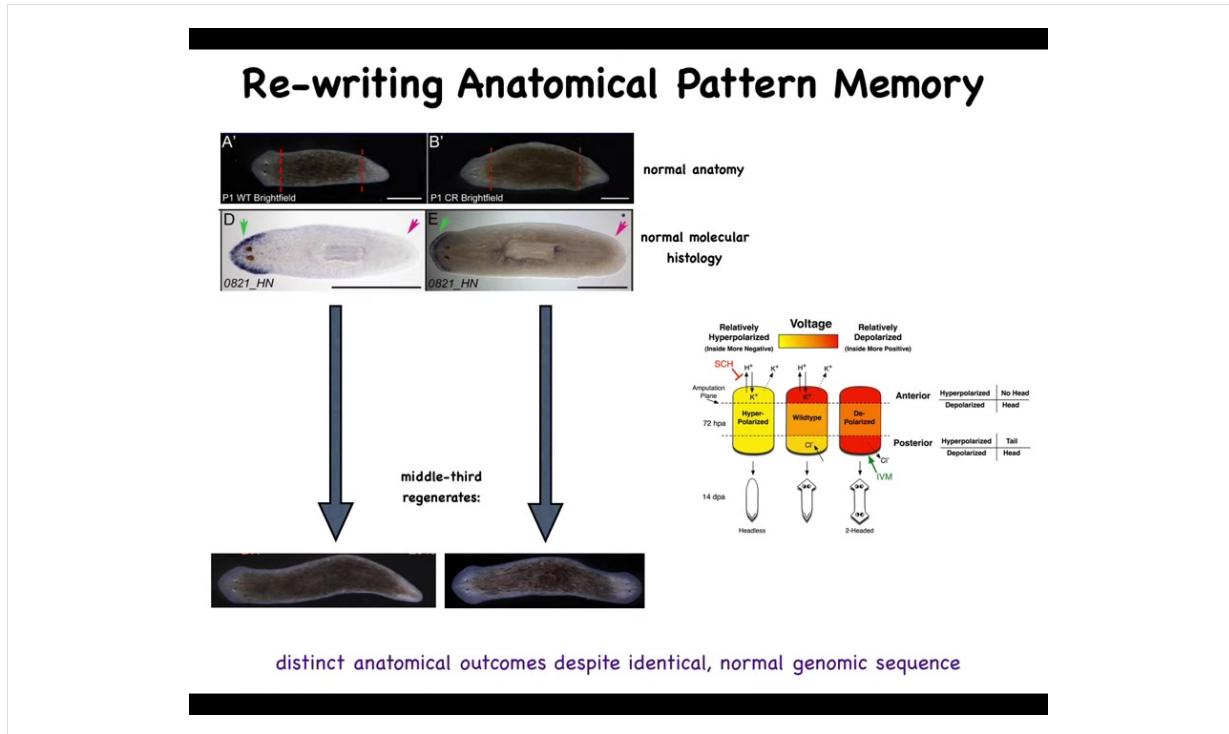
So there's a bioreactor that we make that is delivering the intervention and providing an environment for regeneration. And then there's the payload, which are ion channel modulators and some other things.



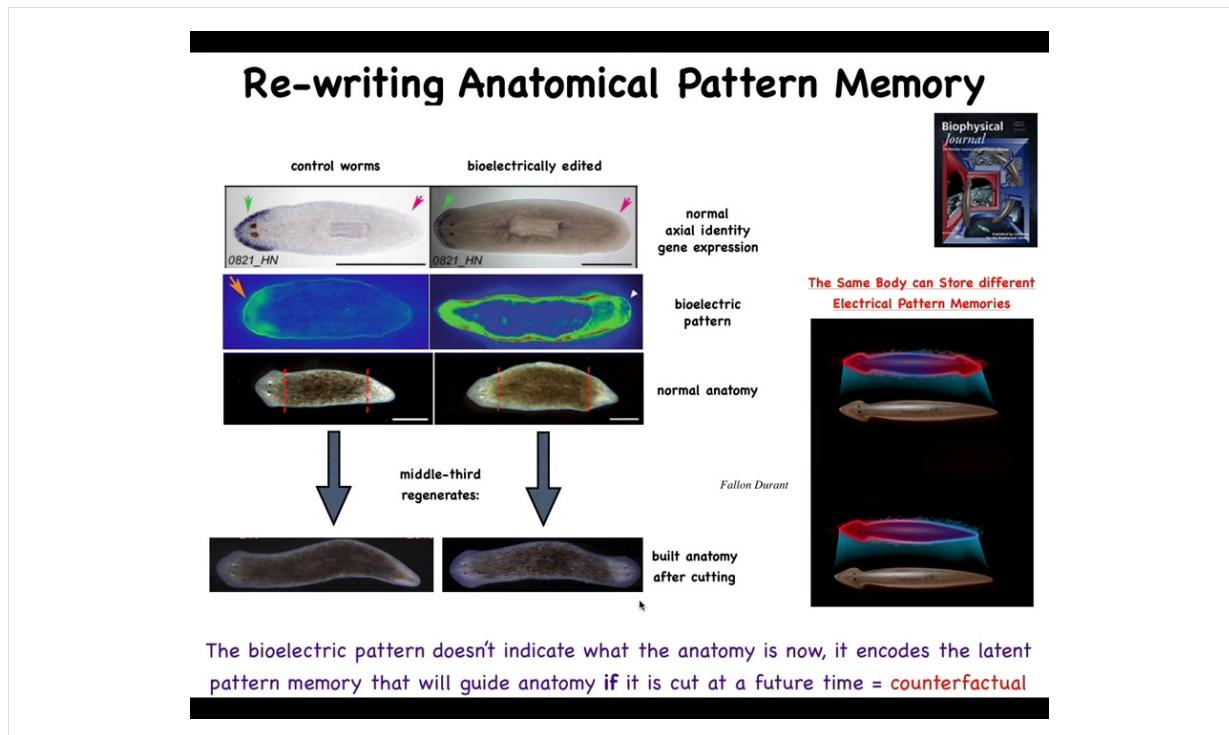
So the next story I want to tell you is to drill down on this notion of bioelectronics as memory, as pattern memory.

So these are planaria. These are flatworms that have a true brain, they have a central nervous system, bilateral symmetry, and so on. One of the amazing things about planaria is that you can cut them into many pieces. The record is 276. When you do, each piece regrows a perfect tiny little worm. There's all kinds of interesting problems. If you cut them in half, these cells up here have to make a new tail; these cells up here have to make a new head, but they were neighbors up until the time you cut them. How do they have these different anatomical positions, anatomical identities, rather? That tells you it's not local decision making. You can't, from the position because that's identical here, tell what you should be. How does each piece know how many heads it should have and which end should be a head, which end should be a tail?

We studied this and it turns out that there is an electrical...



There's an electrical component to this, which is that if you have a single worm and you cut off the head and the tail, this middle fragment will reliably give you a one-headed worm. But you can also make a two-headed worm. How do you do this?



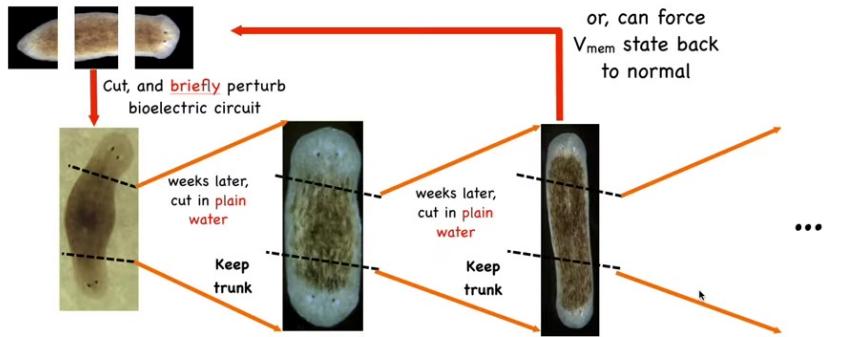
You do this by noticing that there's a bioelectrical gradient that says one head, one tail. That determines what the fragment is going to do. If you rewrite that information, put a false memory into this electrical network that says a normal worm should have two heads, then when you cut them, that's what you get. You get two heads. The cells obey the set point and they build the correct thing.

Notice that this is not a map of this two-headed worm. This bioelectrical image is a map of this anatomically normal one-headed worm. What we've done is change the voltage, but it's a latent memory. It doesn't do anything until we cut the animal. Then you get to find out what the memory was.

The molecular markers, the head marker and the head nod and the tail, and the anatomy are perfectly normal. It's the stored representation of what you should do if you get injured at a future time. That's what's been rewritten.

So this is a kind of counterfactual memory. It is possibly the basis of what we have in brains, which is this amazing time travel. It's this ability to remember things and envision things that are not happening right now. This electrical network is able to do that. It's able to store a pattern that is not the same as the current anatomy. The body of the flatworm can store at least two different representations, two different plans for what a correct planarian is supposed to look like. No doubt there's more, but this is what we've nailed down.

## Like any Good Memory, it is Stable and its content is not determined by the Hardware



### Basic properties of memory

- Long-term stability
- Lability (rewritable)
- Latency (conditional recall)
- Discrete possible outcomes (1H v. 2H)

Nestor Oviedo  
Junji Morokuma



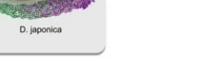
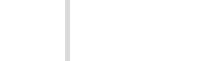
Now I keep calling it a memory because if you take these two-headed worms and you start cutting them in plain water, no more manipulations of any kind, what you'll find is that they continue to generate two-headed animals. Notice there's nothing wrong with them genetically. We haven't touched the genome. So the question of what determines how many heads a planaria has is subtle. You can't just say it's the genome because this two-headed line of animals has a perfectly normal genome. What the genome does is encode the hardware for a system that by default acquires a memory of one-headedness, but it can be changed. Like any good computational or cognitive system, it can be rewritten. It's long-term stable, like any good memory, but it's rewritable. And we can in fact take the two-headed guys and rewrite it back to being one-headed. Here you can see what these two-headed animals are doing.

Okay, we do a lot of computational modeling to try to understand how different memories and different representations become encoded in an electrical network. This is very parallel to what's going on in computational neuroscience and in AI and machine learning.

## A Single Genome Makes Hardware that can Access Bioelectric Memories of Other Species' Head Shapes

Tweaking of bioelectric network connectivity causes regeneration of head shapes appropriate to other species! (also includes brain shape and stem cell distribution pattern)

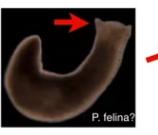
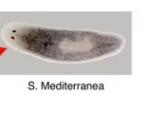
*D. dorotocephala*  7.5%

*P. felina*  *G. dorotocephala*  *D. japonica* 

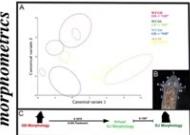
amputation  
free energy landscape  
altered cell network dynamics  
regeneration  
remodelling  
Alexis Pieltak

Maya Emmons-Bell

cut off head, perturb network topology

like: *P. felina?*  like: *S. Mediterranea* 

quantitative morphometrics



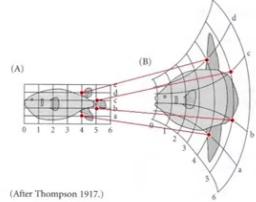
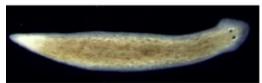
But now notice that it's not just about head number, it's far more than that.

For example, head shape. Here's a planarian; it's a different species with a nice triangular little head. If we amputate the head and we perturb the bioelectrical processes that are normally required for it to remember what its head shape is supposed to be, we can make flat heads like this *P. felina*; we can make round heads like this *S. mediterranean*; or it can make its own normal head.

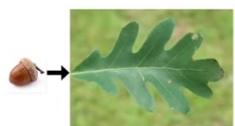
These animals are separated by about 150 million years, but again, no genetic change here. The same hardware—the genome made hardware that is able to navigate regions of that state space of the anatomical space and visit attractors that belong to other species. These other species normally live in these kinds of morphological attractors, but this one can visit them perfectly well if it wants to. It's not just about head shape; it's also the shape of the brain and the distribution of stem cells becomes like these other species. The plasticity is remarkable.

## Hacking Morphogenesis: same genome, explore its latent space via physiological stimuli

Normal



Bioelectric Circuit Altered After Bisection

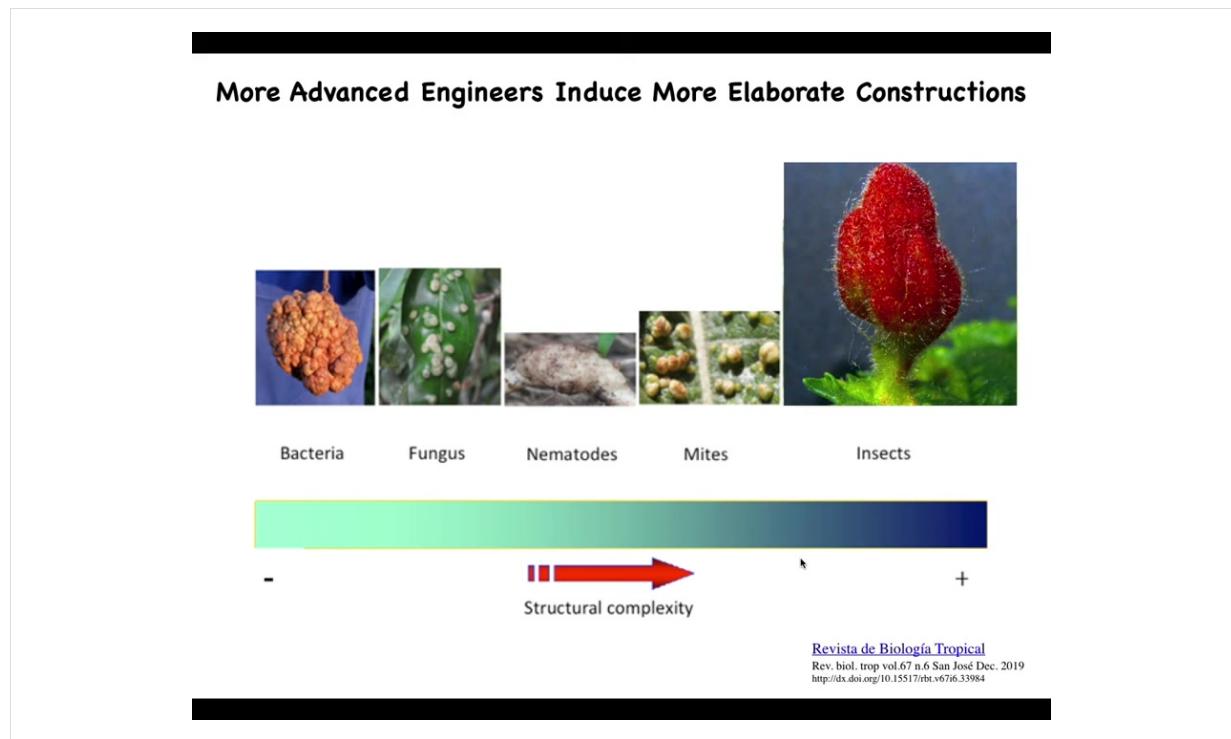


Plant galls –  
leaf morphogenesis  
hacked by signals  
from insect embryo



You can even go further than this, and you can make shapes that don't look like flat planaria at all. You can make these crazy spiky forms, you can make these cylindrical things, hybrids like that. The latent space of anatomical possibilities is huge, and the exact same hardware can be pushed into these directions.

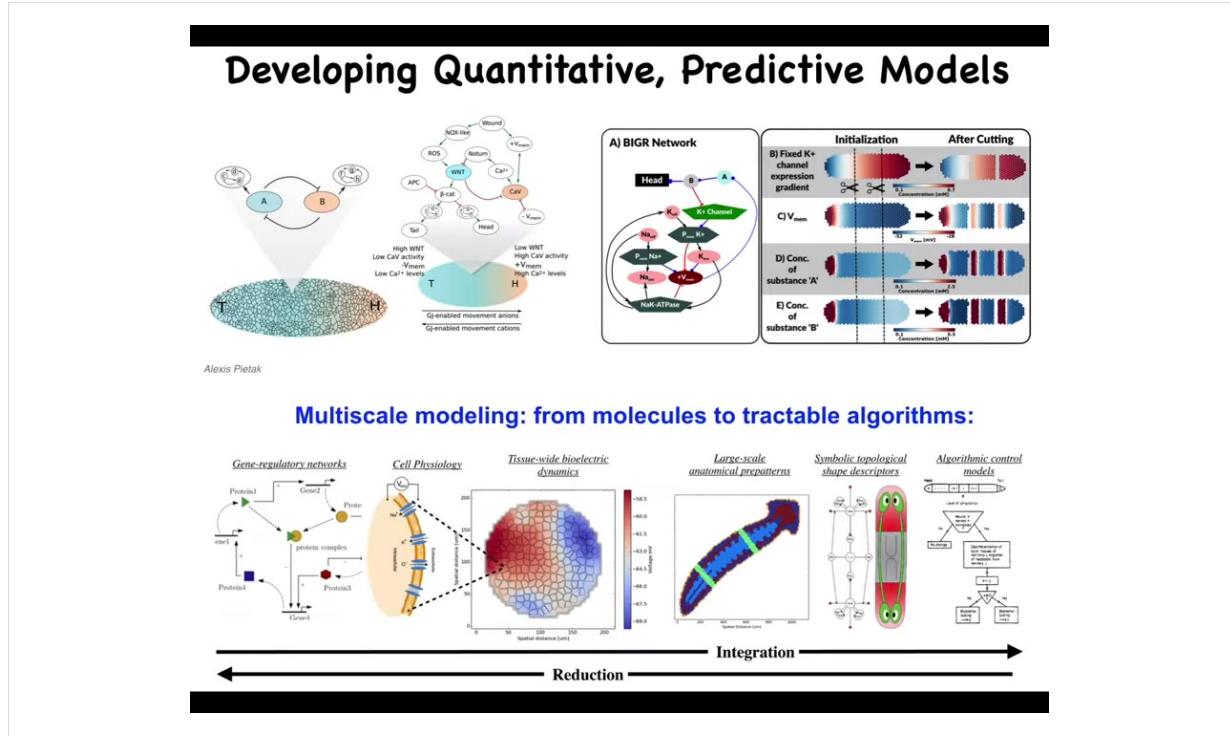
Now, this is a cool example that comes from the plant world. These are acorns and oak leaves, and you might think the oak genome knows how to build these flat green things, but along comes a non-human bioengineer, a wasp that has certain signals that it puts out to hijack the morphogenetic capabilities of the plant leaves and cause them to build something like this. Isn't that amazing? Would you ever think that something like this, that's happened so reliably all over the world, is actually capable? Instead of being this flat green thing, it can actually build these beautiful spiky red structures. And that's because it's been reprogrammed by signals from this non-human bioengineer.



In fact, if you look at the different types of structures that are formed, you see a cool relationship. More advanced bioengineers make more elaborate constructions.

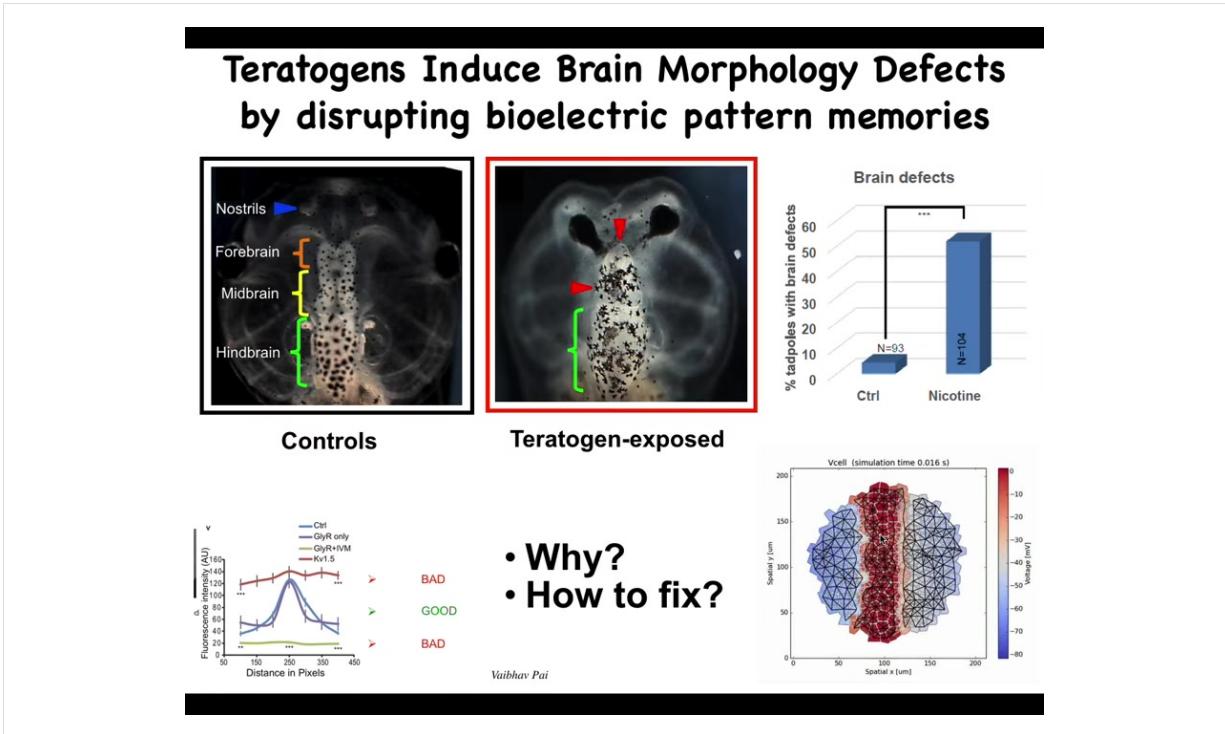
Bacteria and fungi make blobs like this, overgrowth. Nematodes make a shape and mites make little round things. But by the time you get to insects, you can make something like this. This is what they do with the agential material of the plant.

Again, this is presumably not a self-reflective metacognition. Nevertheless, this is what they're actually doing. They're reprogramming, they're taking advantage of the plasticity and the programmability of these cells to do something different from what the genome normally does.

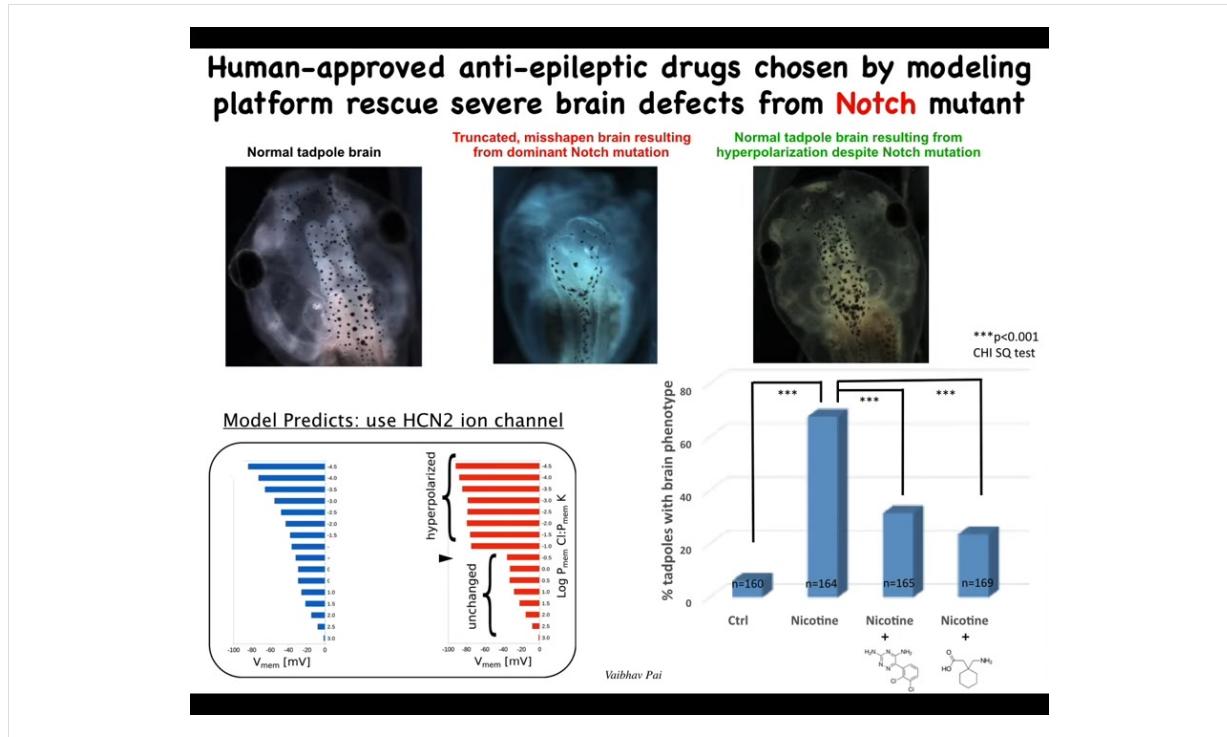


Our goal is to have a full computational understanding, not only of the molecular biology, but the electrophysiology and ultimately, the decision-making, the algorithmic decision-making of the collective, which has to read the electrical information, interpret it, and decide the mapping here that there's a particular state such as that electric face that maps onto a particular anatomical outcome.

That whole process, much like in neuroscience, you can go all the way from the molecules up to the high-level executive decision-making and all of the cognitive processes underneath. The reason that these models are very helpful is because ultimately, what you want to do is make repairs, for example.



Here's an example of taking that approach to deal with birth defects. Here is a normal tadpole brain. You've got your forebrain, midbrain, and hindbrain. If that embryo is exposed to teratogens, you can see some defects here. There are profound defects here. The question is, how do we repair it? We created a model. This is Alexis Pytak in our center, who created this computational model where the voltage pattern is what tells the brain what size and shape to be. We were able to see that what happens in these teratogen-exposed embryos is that the bioelectric pattern is all wrong.



How do you fix it? The computational model can tell you that if you turn on or off a particular ion channel, you can get back to the shape that you want.

Again, this is a normal tadpole brain. Instead of a chemical teratogen, we used a mutation of a gene called Notch. You can see the forebrain is gone. The midbrain and hindbrain are a bubble. These animals are profoundly affected. They have no behavior. They just lay there doing nothing.

The computational model told us that we could get back to the correct bioelectrical pattern if we upregulated the activity of an HCN2 ion channel, because it would sharpen that bioelectrical pattern. When you do that, you get back to a normal embryo with a normal brain, despite the fact that a dominant Notch mutation is still there. It has normal anatomy. The animal's learning rates are indistinguishable from controls at this point. You rescue the anatomy and the function with this bioelectrical repair method, even though the mutation is there.

I'm not saying we can do this for all cases. There are many mutations that are not going to be solvable this way. But in some cases, you can make up for hardware defects with corrective patterning in software.

## Software Tools for Electroceutical Design

What channels to open and close, to get desired system-level bioelectrical state?

<http://bio-software-design-platform.appspot.com/eden>

What ion flows would repair the physiological state?

Which channels/ pumps can be targeted to achieve desired ion flows?

What drugs are available to implement the intervention?

w/ Jack Tuszyński and Philipp Winter

**EDEn platform**

The platform consists of three main components: 1. A flowchart on the left showing the process of identifying ion flows, channels/pumps, and available drugs. 2. A central image showing a tumor with a blue arrow pointing to a list of tissues. 3. A right panel titled 'EDEn - Electroceutical Design Environment' containing a tissue-specific ion channel database with a list of tissues and search filters.

**EDEn - Electroceutical Design Environment**

A tissue specific ion channel database

1. Select one or more tissues:

Normal tissues: bronchus, temporal lobe, testis, thalamus, thyroid, tongue, tonsil, trachea, urinary bladder, uterine cervix, uterine endometrium, uterus, vagus nerve, adipose tissue

Cancer tissues: colorectal adenocarcinoma cell, breast cancer cell, cervical carcinoma cell, thyroid carcinoma cell, testicular cancer cell, gastric cancer cell, glioma cell, lung cancer cell, pancreatic cancer cell, liver cancer cell, uterine adenocarcinoma cell, K-562 cell, colon cancer cell, carcinoma cell, HL-60 cell

2. Show ion channels with mRNA expression level greater than   Incl.

3. Select comprehensive to show all channels expressed in the selected tissues (ordered by mos). Select Unique to show only channels that are specific to the selected tissues.

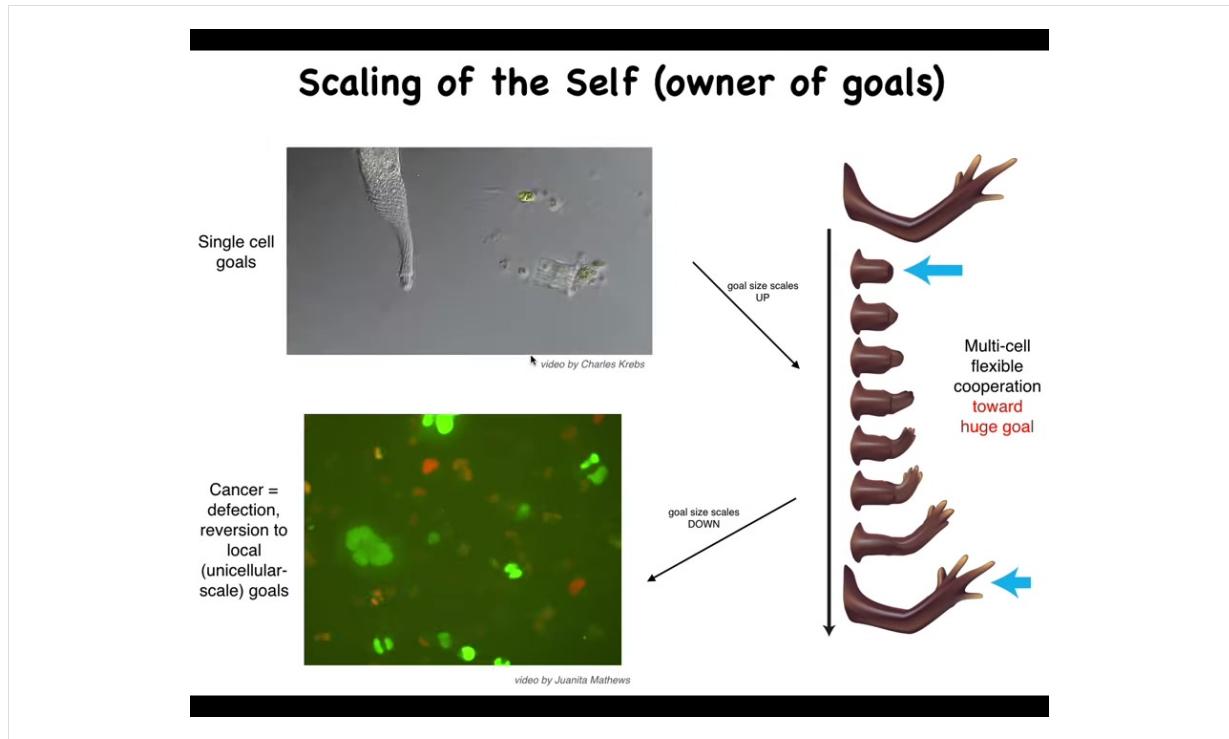
Comprehensive  Unique

4. Select one or more ion channels:

Sulfonurea receptors; K-ATP channels (Q15842|KCNQ8)  
 Transient receptor potential cation channel subfamily M member 7 (Q96QT4|TRPM7)  
 Transmembrane protein 109 (Q9BVCG|TMEM109)  
 Vesicle-associated membrane protein 2 (P63027|VAMP2)  
 Voltage-gated calcium channel subunit alpha 1C (Q92P9|CACNA1C)  
 Voltage-dependent calcium channel gamma-like subunit (Q8WXS4|TMEM37)  
 Voltage-dependent L-type calcium channel (Q16040|CACNA1F)  
 Voltage-gated L-type calcium channel (Q13698|CACNA1G)  
 Voltage-gated calcium channel alpha 1D subunit (Q13936|CACNA1D)  
 Voltage-gated calcium channel (Q05916|CACNA1A)  
 Voltage-gated calcium channel (Q15878|CACNA1E)  
 Voltage-gated potassium channel (P48547|KCNCl)

Where this is going ultimately is there will be a tool like this where the tool will have information about what the correct bioelectric pattern is, and what the incorrect pattern is in some disease state, for example tumor, which is what I'm going to talk about next.

And then there will be a computational component that will ask what channels and pumps could be turned on and off to repair the pattern. Once you know that, you can ask what drugs are available to do that. Something like 20% of all drugs are ion channel drugs. There's this incredible toolkit of our electroceuticals just sitting out there to be applied to all these problems with the appropriate computational models. You can play with an early version of it here. We already have a little bit of this going. I've shown you an example of doing this for birth defects. I've shown you an example of doing this for regeneration.

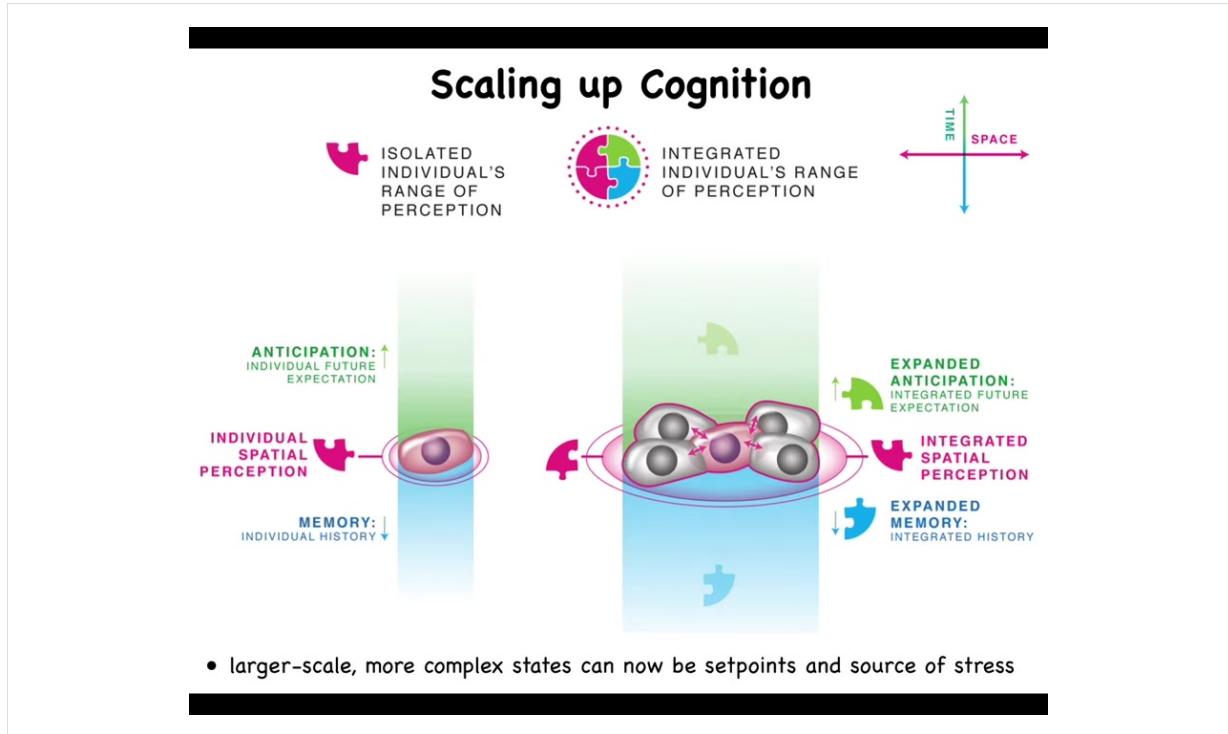


I want to get to cancer because cancer is an important component of all of this. So let's remind ourselves that what happens with this agential material during evolution and during development is that its goals scale. The cognitive light cone of these systems, meaning the size of the goals that they can pursue, scales up. Individual cells have tiny goals. The only thing the cell is going to work towards is a specific pH, metabolic, and other levels right around its own surface. It doesn't care about anything else here.

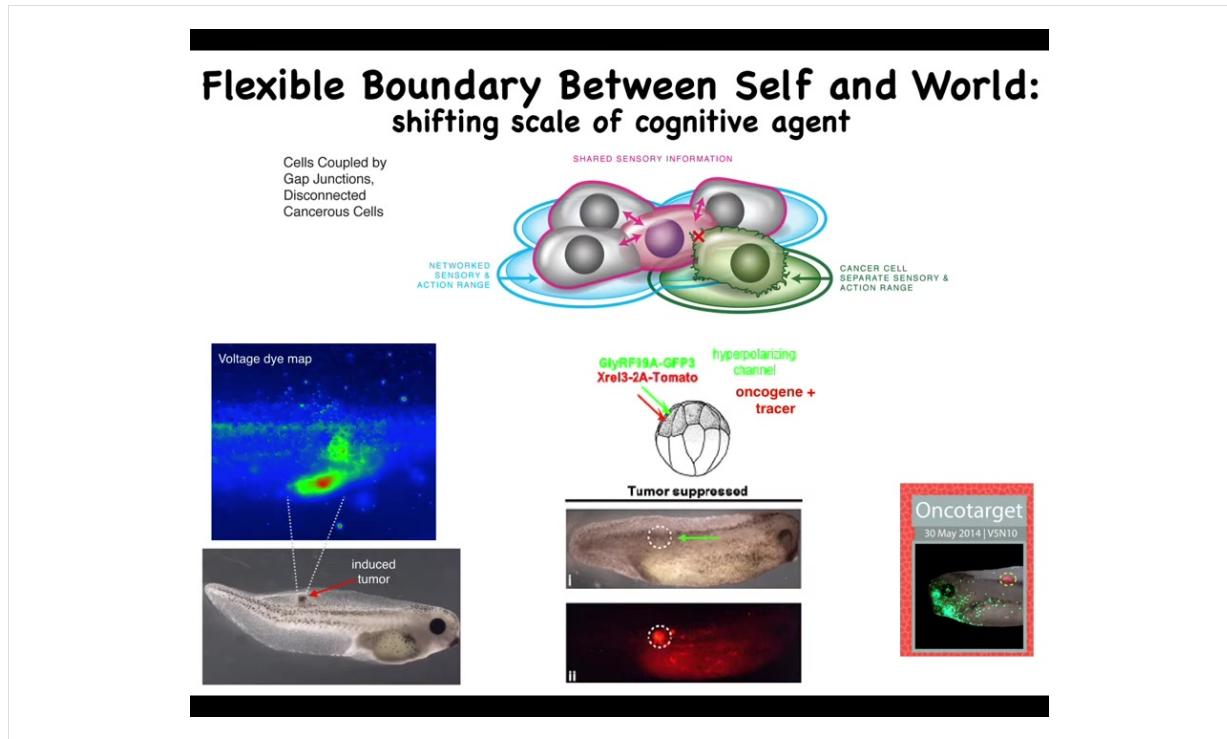
But this system has a huge goal. It knows it needs to have 4 fingers and be a certain size. No individual cell knows this, but the collective absolutely does. Because if you cut off fingers or you cut off the whole thing, it will do exactly what it needs to do to get its goal met. So the collective has a very large set point. These guys have little, tiny set points. And I've just shown you that the cognitive glue that binds individual cells towards systems that can store and pursue large goals is bioelectricity. It's the electrical networks that allow these cells to remember what a limb is.

That whole cognitive glue system has a failure mode, and that's called cancer. What happens in cancer is that cells physically and electrically disconnect from their neighbors. As soon as they've disconnected from the electrical network, they can no longer remember any of these large-scale goals. They just go back to their ancient, primitive, tiny goals. That is: go wherever life is good and proliferate. And that's metastasis.

So that's what you're seeing here is individual cells, which are glioblastoma, rolling back to their unicellular past.



That's what we have here. We have individual goals tying together via bioelectronics and other signals, biochemical and biomechanical signals, into systems with a much larger degree of both temporal and spatial goal directedness. That very weird view of cancer and cancer suppression suggests an alternative medical modality.



What if we tried to reconnect them to the electrical network despite whatever problems they had?

We already know from both the brain repair case and from the two-headed planaria that genetics does not tell the whole story, that what drives is the physiology that's operating in these cells.

What if we tried this for cancer? Here it is, we inject a human oncogene, KRAS mutations, P53, and so on. They will eventually make a tumor. If you watch that spot, you can already tell where the tumor is going to be. This is again a voltage DIMAP. You can already tell that there's some aberrant voltage going on here. These are the cells that are going to disconnect and treat the rest of the body as just external environment. They don't consider themselves a part of the whole anymore. It's just outside to them. They're going to go wherever they want, do whatever they want.

If we co-inject with that oncogene an ion channel that forces these cells to have a particular voltage pattern and to stay connected to their neighbors, then this happens, while this is the same animal, while the oncoprotein is still blazingly expressed. There's mutation all over the place here. If you sequence it, you would make the wrong prediction. You would say that there's definitely going to be a tumor here. Look at how much of this oncoprotein there is. There is no tumor, because that's not what drives. What drives is not the genetics. What drives is the physiology: once you've connected the cell to that neighbor, to their neighbors, the collective is saying, "We are building nice skin, nice muscle, and so on." The collective does not have thoughts about

migrating away and doing something else. What it knows how to do is navigate anatomical space to keep this tissue healthy.

This is both a diagnostic and a therapeutic modality that we're pursuing now in human tissues for cancer via this kind of boundary of a self approach where we try to inflate the boundary between self and world. The idea is that these cancer cells are not any more selfish than normal cells; they just have smaller selves.

Now for the last little bit of this, I want to shift from, so far we've had a focus on native outcomes. In other words, let's repair the birth defects to the correct shape. Let's regenerate organs that were there before. Let's prevent or normalize cancer so that you can have normal tissues. Now we're going to go beyond this. Here's the bioengineering part of this talk.

Slide 42 of 57 · Watch at [47:29](#)

## Biorobotics: novel anatomical goals

*What I cannot create, I do not understand.* R. Feynman

can cells, liberated from their normal boundary condition, reboot multicellularity? Do they like to cooperate, and what will this swarm agent decide to build?



The Institute for  
Computationally  
Designed Organisms  
Tufts University | University of Vermont



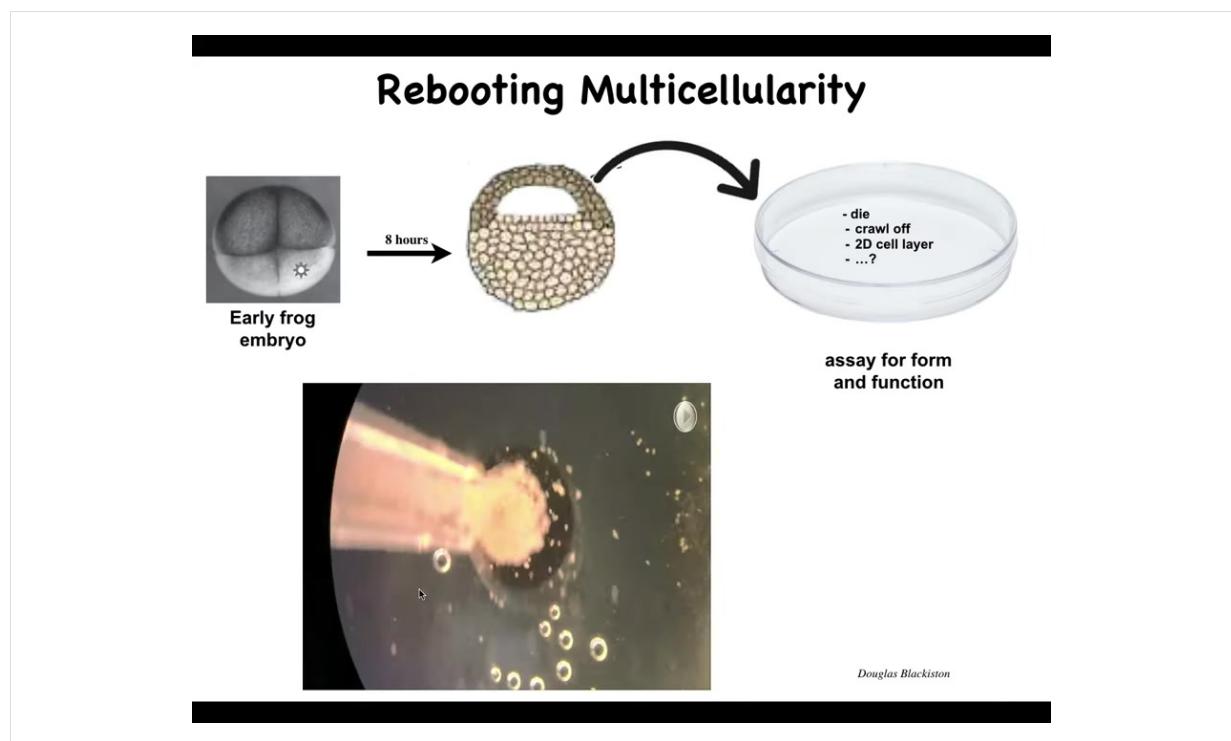
Douglas  
Blackiston

We're going to talk about bio-robotics and specifically this idea of novel anatomical goals. I've shown you systems that pursue evolved anatomical goals, such as having a proper face, the proper number of heads, the proper structure of the brain. But what about novel systems? Where do they get their goals, and what are their goals?

Famously, Feynman wrote that "what I cannot create, I do not understand." The idea is by synthetically creating novel beings that have never been here before through evolution, we gain some understanding of where goals come from in the first place.

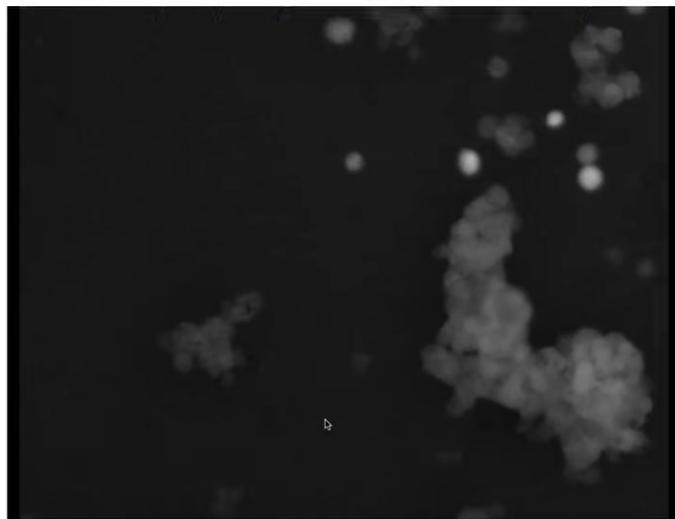
This is work largely spearheaded by Doug Blackiston in my group, in our collaboration with Josh Bongard's lab in our Institute for Computationally Designed Organisms, where we wanted to know, could we use cells liberated from their normal context, the embryo? Could they reboot multicellularity? What would they want to build?

Slide 43 of 57 · Watch at [48:24](#)



What we did here was, here's an early frog embryo, here is the cross-section, these animal capsules normally make ectoderm, so think prospective skin layer. We take these cells, we dissociate them and put them in a petri dish. Once you do this, they could do many things. They could die, they could crawl off and get away from each other, they could make a two-dimensional monolayer like cell culture. But instead they compact overnight and become this little thing. The flashes you're seeing are a calcium reporter or calcium signaling. They become something we call a Zenobot.

## Microbots



*Douglas Blackiston*

The first thing that happens is we call these microbots. Each one of these is an individual cell. And this little thing that remarkably looks like a little horse is one of many examples where they get together. They have these interesting movements. They start interacting with each other. There's that calcium flash. This thing is skipping along in the Petri dish. And there's all kinds of little structures.

### Xenobot behaviors - repurposing cilia for motion

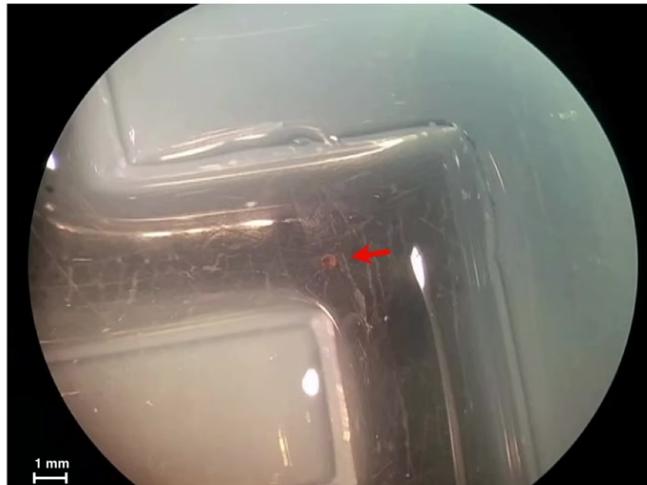
collective behaviors

Douglas Blackiston

And eventually they all come together and they make this. This is our Xenobot. We call it a Xenobot because *Xenopus laevis* is the name of the frog and we think it's a biorobotics platform among other things. So Xenobot.

They're swimming along because they have little hairs on their surface that row against the water. These are cilia that frogs normally use to push mucus around on their skin. They can go in circles. They can patrol back and forth. They can have group or collective behaviors. This is some tracking data. You can see this one's going on a long journey. These two are interacting with each other. Some of them are hanging out doing nothing. We can make them into weird shapes, these donuts that then swim around.

### Xenobot in a maze (still water, no flow):

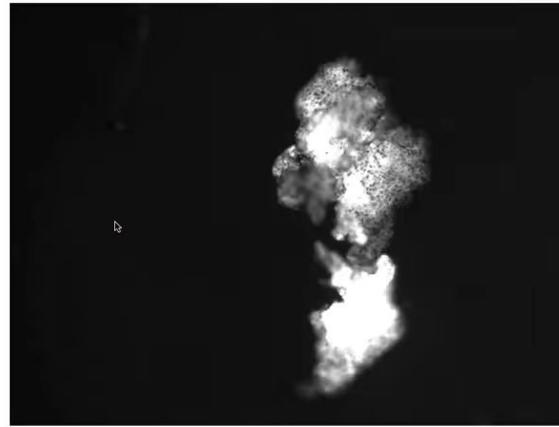


- 1) it traverses maze,
- 2) rounds the corners without bumping into walls, and
- 3) it makes a spontaneous decision to turn around without hitting anything.

*Douglas Blackiston*

Here's one traversing a maze. There's no fluid flow. It swims down here by itself, takes the corner without bumping into the opposite wall, and then decides to turn around and come back. These are completely autonomous. We're not pacing them; we're not making the move.

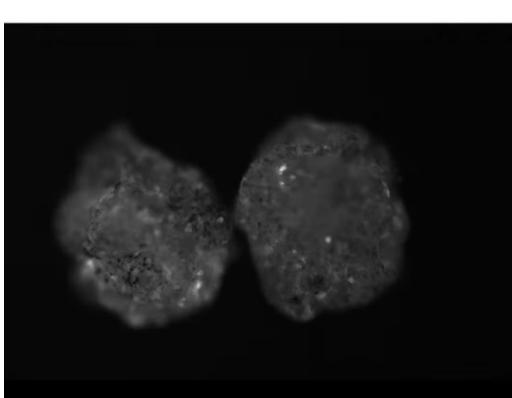
## Xenobots Regenerate after Massive Damage Back to their New Xenobots Shape



*Douglas Blackiston*

They also have the ability to self-repair. If you chop them almost in half, they will fold back up like this. It's like doing a bicep curl from 180 degrees; that hinge is generating a ton of force to fold this thing back up.

**Reading the Xenobot mind: calcium spiking  
in skin cells – there are no neurons here**



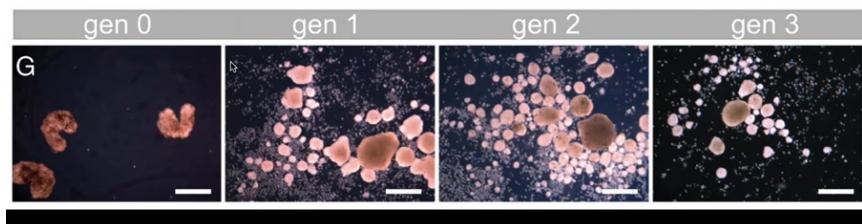
The image shows a screenshot of a bioRxiv preprint page. The title 'Reading the Xenobot mind: calcium spiking in skin cells – there are no neurons here' is displayed in large, bold, black font. Below the title is a grayscale image of two biological cells. The bioRxiv logo is in the top right corner, along with 'HOME | SUB'. Below the logo, the text 'New Results' and 'Revealing non-trivial information structures in a neural biological tissue via functional connectivity' is visible. The authors listed are Douglas Blackiston, Hanul Shimminck, Cecile Grasso, Thomas F. Tolley, Douglas G. Moore, Kristen S. Sosman, Olaf Sporns, Johnnie Roegner, Michael Lewis, and Sara I. Walker. The doi is https://doi.org/10.1101/2024.01.09.593467. The author Douglas Blackiston is mentioned at the bottom.

They do have these interesting flashes. This calcium signaling is how people profile activity in the brain. One thing you might ask, and we have a pre-print of this, even though there's no neurons here, how similar is this kind of calcium activity? What's the computation that might be done here? How similar are they to brains? Even though this is just skin.

## Kinematic Replication in Xenobots: evolution, like engineers, works in agential materials

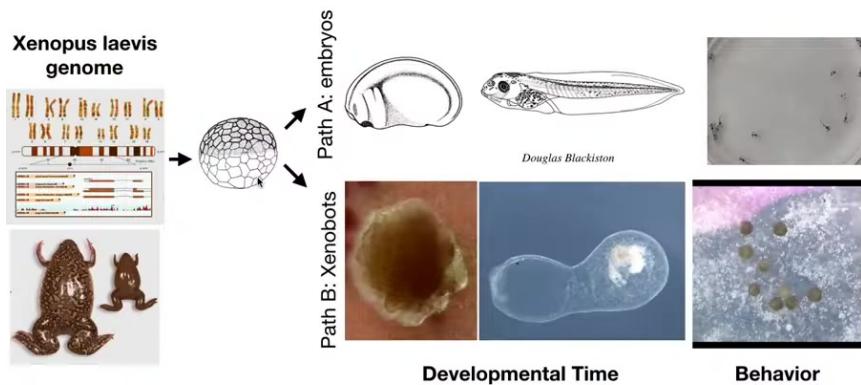


Douglas Blackiston



One of the most remarkable things they do is something called kinematic replication. We've made it impossible for these guys to reproduce in the normal \*\*\*\*\* fashion, but if you give them a bunch of loose skin cells, they gather these cells up into little balls and they polish them, both collectively and individually. These little balls mature into the next generation of xenobots. They make the next generation. It's a kind of kinematic self-replication. It works because evolution, like us, is working in an agential material. These cells are not passive particles. They are cells themselves that will become xenobots.

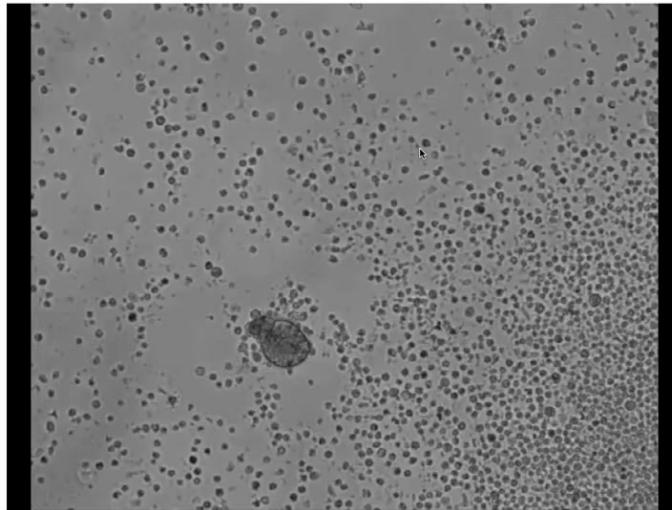
## Xenobots have a Standard Frog Genome Xenobots have a Unique Developmental Sequence



The amazing thing about this is that there haven't been any Zenobots before. There's never been any selection to be a good Zenobot. Where do their properties come from? What did the frog genome learn during its many years on Earth? It certainly learned to do this. It's hardware that can make this set of stages with behavior. I hear some tadpoles running around. Apparently, that's not the only thing it can do. It can do this. While this was under strong evolutionary selection for specific environments, these never directly existed until we made them. This is an 83-day-old Zenobot. What is it becoming? What shape is this? Where does the kinematic replication come from? It's a remarkable thing to be thinking about the goals and competencies of creatures that have never been here. You can't say that it's evolution that gave it to them. They have no straightforward evolutionary backstory like all the other animals that everybody studies. Soon, we will have some actual behavioral studies out, so stay tuned for that. We've been studying their ability to sense, to remember, to behave, and that's all coming.

One thing you might think of is that this is some weird, unique fluke of amphibian embryonic development. Amphibians are pretty plastic. Embryos are definitely plastic. Maybe this is just applicable to frog embryos.

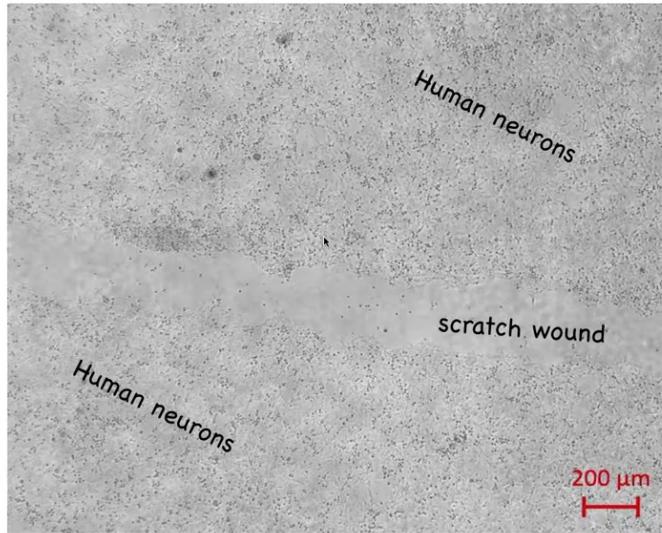
## Are Xenobots a fluke of frog development?



Gizem Gümüşkaya

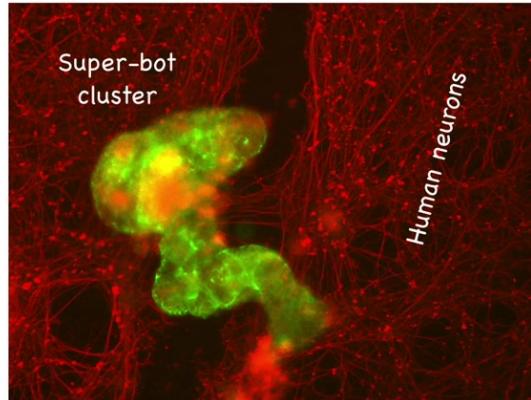
So I'm going to show you this. If you look at this thing and I ask you, what do you think this is? You might think that this is, we got this out of the bottom of a pond somewhere, some kind of primitive organism. And if I ask you what the genome would look like, you might guess that it would look like the genome of primitive organisms. If you did sequence it, you would find that it's 100% *Homo sapiens*. These are 100% human cells. They are not embryonic cells. They are taken from the tracheal epithelium of adult human donors. And much like the xenobots, when given the opportunity, these cells will take that chance to have another life. They reboot themselves into a novel creature that doesn't look like any stage of human development. And they have a completely different structure. And some very different properties. And again, they are self-motile.

## Anthrobots Navigate Wounds in vitro:

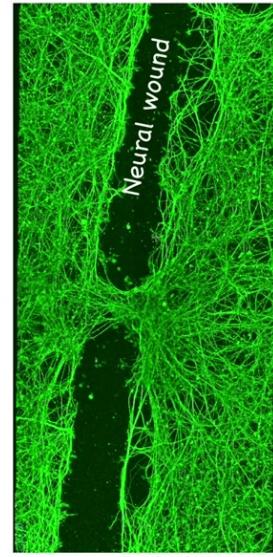


One thing that they can do, for example, is traverse wounds. This is a big scratch wound placed in some human neurons. Here, this is a lawn of human neurons. There's a scratch that we made, and they can traverse it.

## Anthrobots Exert Neural Repair



Toward personalized autonomous therapeutics



Gizem Gümüşkaya

When they do traverse it, they can collect in a pile. We call this a super bot cluster. It's probably around 10 bots here. You can see this clearly when you lift them up four days later: they're sitting there knitting the two sides of the gap together.

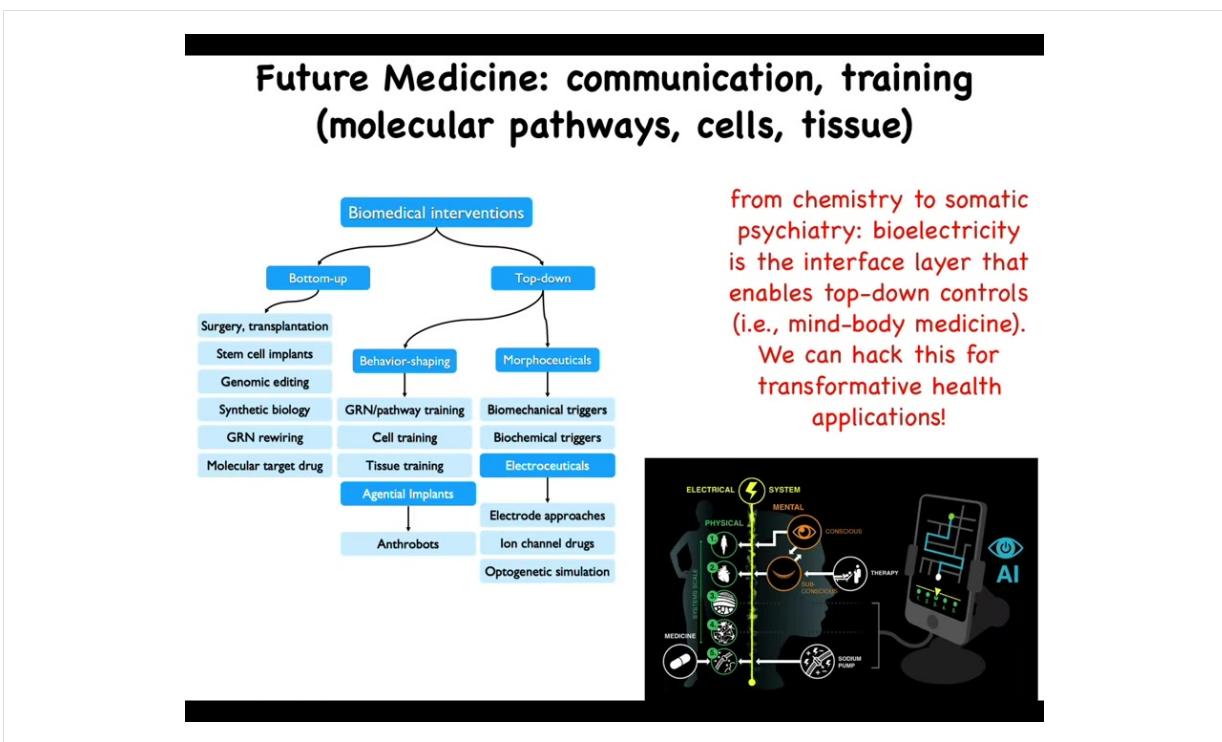
So who would have thought that your tracheal epithelial cells that sit there quietly in your airway, doing nothing other than moving mucus and dust particles, have the ability to reassemble into a completely novel little organism with a completely novel transcriptome. I didn't show you this, but they transcribed their genes: about half of the human genome was transcribed completely differently in these guys because of their new lifestyle. There are no transgenes, no synthetic circuits, no synbio, no scaffolds, no weird drugs. We didn't edit the genome at all. These are completely normal cells doing something that nobody had expected them to do. They have the ability to exert repair.

We think this is the beginnings of a technology for personalized autonomous therapeutics. We could take cells out of your airway, use them to create Anthrobots, and put them back for a range of applications, for example neural repair, but in a way that does not require immune suppression. It doesn't — they're your own cells. You don't have to worry about protecting them from the immune system. You wouldn't be worried about nanomaterials accumulating in the body. They're completely natural, made of your own cells, and they degrade after a few weeks.

We think this is the beginnings of these kind of autonomous therapeutics, which are way smarter than anything we can build today. They have a billion years of living

inside a body. They already know what inflammation is, what cancer is, what stress is like. I think the sky is the limit in learning to use these things for therapies.

Slide 54 of 57 · Watch at [56:38](#)

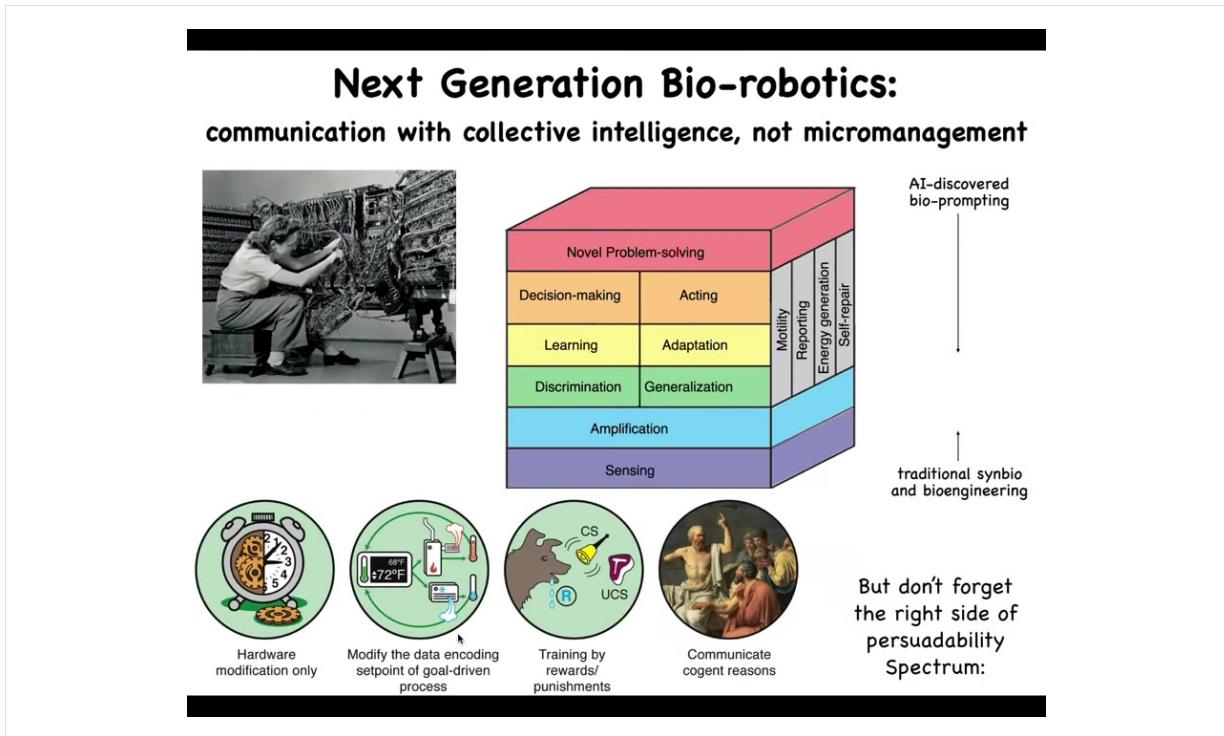


To wrap up here, a couple of things. First, I want to point out that the whole field of biomedical interventions can be categorized as bottom-up or top-down. This is everything that's been going on until now. Surgery, stem cells, genomic editing, synbio, and targeting drugs with molecular, targeting receptors with molecular drugs are all bottom-up, hardware-focused approaches. The future of biomedicine, I think, is complementing that with powerful top-down approaches. Behavior shaping — I didn't talk about the cell and tissue training that we're doing. Agential implants such as anthrobots and the morphoceutical, and in particular electroceutical, signals that are available for getting cells to do very complicated things that we do not want to micromanage. The name of the game here is going to be cooperation and collaboration with the material, not micromanaging.

We think that AI is going to be an important component of this to be able to crack the communication interfaces so that we know how to talk to all the different levels. The way we talk to each other at the highest cognitive level, we should be able to talk to all

of the somatic intelligences within the body. I think that future medicine is going to look a lot more like somatic psychiatry than it's going to look like chemistry. It's going to be quite different.

Slide 55 of 57 · Watch at [58:04](#)



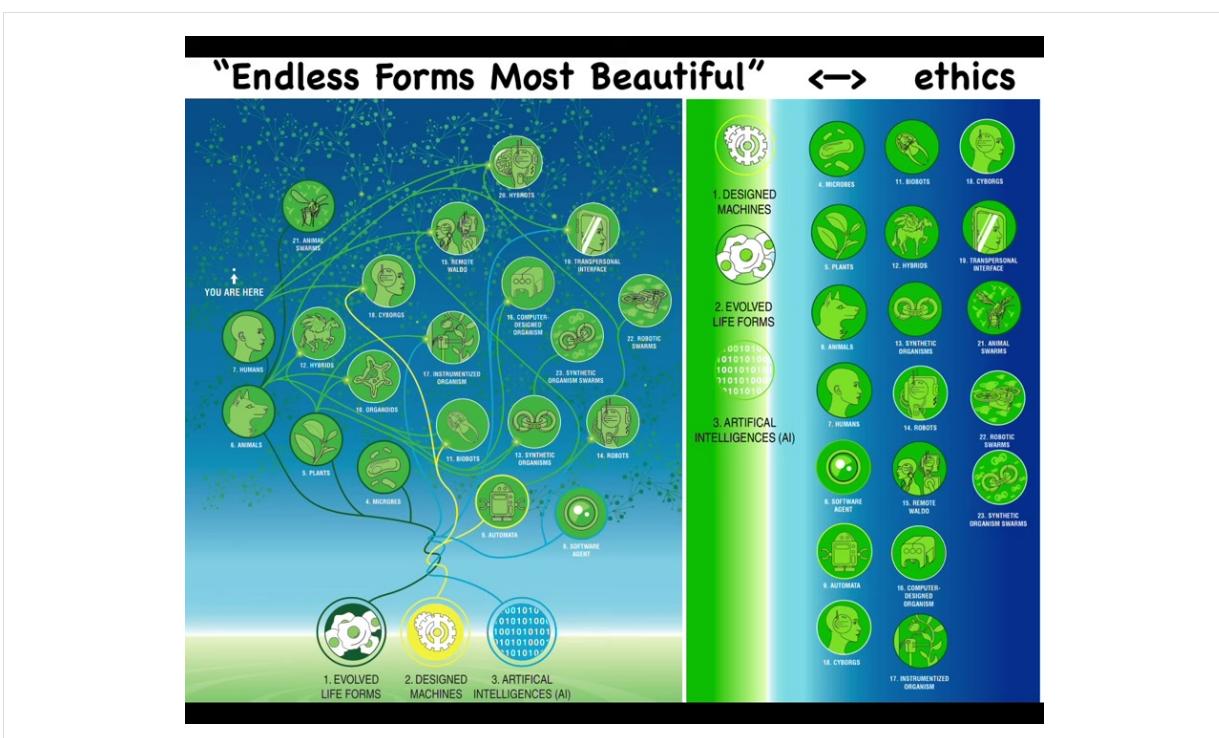
That's for biomedicine. For bioengineering and biorobotics, in many cases you don't want to build from scratch where you have to micromanage all the hardware. The material that we're dealing with already has built-in sensors, amplification machinery, the ability to discriminate signals, to generalize, to learn, to adapt, to make decisions. I showed you many examples. Taking advantage of this is the frontier of bio-robotics; it's not trying to force specific molecular states, but actually communicating and collaborating and reprogramming these existing capabilities. The AI discovered bio-prompting of the material.

The very last thing I want to point out is to step beyond medicine and engineering to the ethics component of all this, which is the following. This is what I call the spectrum of persuadability. All systems are somewhere along a spectrum like this related to what kind of approach you're going to take to relate to them.

Very simple machines are only amenable to hardware rewiring. Then you get these cybernetic things like thermostats, where there's a set point and you can modify them by controlling it. Then there are some systems that learn. You can use behavior shaping and rewards and punishments. I've argued that cells are somewhere in this neighborhood. It's not just brainy animals; cells and tissues are in this neighborhood too.

On the left it's all about control. As you move to the right of the spectrum, it's less and less exerting control and more communication and collaboration. By the time you get to something that's up here, it's a very bi-directional relationship because the level of agency rises. By the time you're having this kind of interaction, which is not about control or manipulation, it's about bi-directional, mutually beneficial interactions. It's about communication and thriving together as high-agency systems.

Slide 56 of 57 · Watch at [1:00:30](#)



The reason I bring this up is because of life's incredible interoperability and the future of bioengineering. We, and certainly you as students in this audience in your lifetime, are going to be living with a very wide range of beings.

Everything that Darwin meant when he said "the natural world and its endless forms, most beautiful" is like a tiny little region in this whole space of bodies and minds. Because every combination of evolved material, engineered material, and software is some kind of possible embodied being: hybrids and cyborgs, augmented humans, and weird combinations of material extended both biologically and technologically. All of these things we are going to live with, and most of them are going to be nowhere on the tree of life with us. They're quite different. And this means that we're going to have to develop novel forms of ethics that go beyond what do you look like and how did you get here, meaning evolved or engineered, in terms of knowing how we deal with others that are not quite like us. So that achieving some sort of ethical synthbosis with a very diverse range of beings goes way beyond the kind of control goals that we have in engineering and more like a bi-directional vulnerability and relationship.

I'm going to stop here. I'm way out of time. Lots of this is discussed in detail in these papers, both on engineering and medicine.

Slide 57 of 57 · Watch at [1:02:01](#)

## Thank you to:

### Post-docs:

Junji Morokuma - bioelectronics of melanocyte transformation  
 Douglas Blackiston - brain-body interface plasticity, synthetic living bodies  
 Sam Kriegman - computational modeling of biorobotics  
 Devon Davidian, Evgenia Shmelyova, Nirosha Mungan, Celia Herrera-Rincon, and AiSun Tseng - limb regeneration  
 Vaibhav Pai - voltage gradients in eye/brain induction and repair  
 Juanita Mathews - bioelectricity of cancer  
 Nestor Oviedo, Wendy Beane, Johanna Bischof - bioelectronics of planarian regeneration  
 Patrick McMillen - bioelectric imaging and transitions to multicellularity  
 Patrick Erickson - cellular learning



### Ph.D. Students:

Franz Kuchling - cognitive neuroscience modeling applied to pattern homeostasis  
 Fallon Duran -  $V_{mem}$  and pattern memory in planarian regeneration  
 Adam Goldstein - cancer and bio-inspired AI



### Undergraduate Students:

Maya Emmons-Bell, Kelly G. Sullivan - non-genetic cross-species morphogenesis  
 Pranjal Srivastava, Ben G. Cooper, Hannah Lesser, Ben Semegran - Anthrobots



### Technical support:

Rakela Colon, Jayati Mandal - lab management  
 Erin Switzer - vertebrate animal husbandry  
 Joan Lemire - molecular biology



### Collaborators:

Allen Center members +  
 Dame Adams - bioelectric imaging, electric face  
 Alexis Pietak - computational modeling of bioelectronics  
 Joshua Beckman - X-ray simulation and AI  
 Selvam Raja - open-ended evolution  
 Sara Walker, Karl Friston - top-down causation models  
 Simon Garnier - computational analysis of Anthrobot form and function  
 David Kaplan -  $V_{mem}$ , am, human MSC differentiation, biomechanics  
 Giovanni Pezzulo - cognitive science models of pattern regulation  
 Vitaly Volpert, Chris Fields - mathematical models of pattern regulation  
 Richard Watson - computational models of cognitive scaling and evolutionary learning  
 Don Ingber, Richard Novak, V. J. Koomson, J. H. Dungan - mammalian bioengineering  
 Jack Tuszyński - cell biophysics and drug discovery



### Model systems:

tadpoles, planaria, zebrafish, slime molds, human cells, and chick embryos



### Funding support:

JTF, TWCF, DARPA, Paul G. Allen Frontiers Group, Sloan Foundation, NIH, NSF

### Illustrations:

Jeremy Guay @ Peregrine Creative

Disclosures: Morphoceuticals, Fauna Systems, Astonishing Labs, Unravel

Really critical. I want to thank all the people who did the work. Here are the amazing postdocs and PhD students who did all of this work. We have lots of incredible

collaborators, funders, several companies that support us. Most of all, the model systems, the animals that teach us everything. We know because they do all the heavy lifting. I will stop here and take questions.

**Thank you for reading.**

---

**More lectures**

You can find more of my lectures [here](#).

---

**Follow my work**

[Twitter](#) • [Blog](#) • [The Levin Lab](#)

---

**Want one for your lecture?**

Want something like this for your own talk? Reach out to Adi at [adi@aipodcast.ing](mailto:adi@aipodcast.ing).