

A LECTURE COMPANION

"Biorobotics: engineering with agential materials"

Michael Levin

Recorded on May 5, 2024

About this document

This document is a companion to the recorded lecture "*Biorobotics: engineering with agential materials*", recorded on May 5, 2024. 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 an ~1 hour talk by Michael Levin on the topic of engineering with agential (living) materials, given for the students and other members of the ETHZ Zurich Soft Robotics Lab led by Prof. Robert Katzschmann).

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.

Biorobotics:
Engineering with Agential Materials

Michael Levin
Allen Discovery Center at Tufts

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

ALLEN
DISCOVERY CENTER
at Tufts University

Computer-designed Organisms
TUFTS UNIVERSITY | UNIVERSITY OF VERMONT

WYSS
INSTITUTE
HARVARD

Thank you so much. It's a great pleasure to be there. I love talking to students and I hope to tell you some interesting stories today. I'm going to talk about bio-robotics and in particular engineering with agential materials.

Afterwards, if you want to follow up any of the details, all of the primary publications, the data, the software, all of that is at this site. If you want to look at some broader thoughts that I have about this work, you can find my blog here.

Main Points:



Michael Levin

- 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**

Here are some of the main points that I'm going to transmit today. Most problems in biomedicine and many problems of bioengineering boil down to the control of morphogenesis. The idea of controlling the collective decision-making among cells is critical for advances in many areas. I'm going to argue that this is not going to be solved by hardware-based approaches such as genomics.

In fact, biology uses a kind of multi-scale competency architecture of nested problem solvers in very different problem spaces. We, in evolution, exploit this architecture, and in particular this bioelectrical interface that cells and tissues use to shape each other's behavior to create and maintain complex functional structures. We now have tools that allow us to read and write memories into this software layer, this physiological layer of control that sits between the genome and the anatomy, and this has many applications in birth defects, regeneration, cancer, and synthetic bioengineering. I'm specifically going to emphasize engineering with agential materials, that is, materials with an agenda, where we can use tools from behavioral and cognitive sciences to exploit the intelligence and the competency of cells and groups of cells in these unusual problem spaces.

Biomedical Endgame: Anatomical Compiler

Name: Triple head

• Add region
• Change type
• Add cell
• Add gene
• Delete selection

Head
Tail
Sai

Daniel Lobo

Add image

OK Cancel

Why we need it:

- Birth defects
- Traumatic injury
- Cancer
- Aging

Problems of information processing

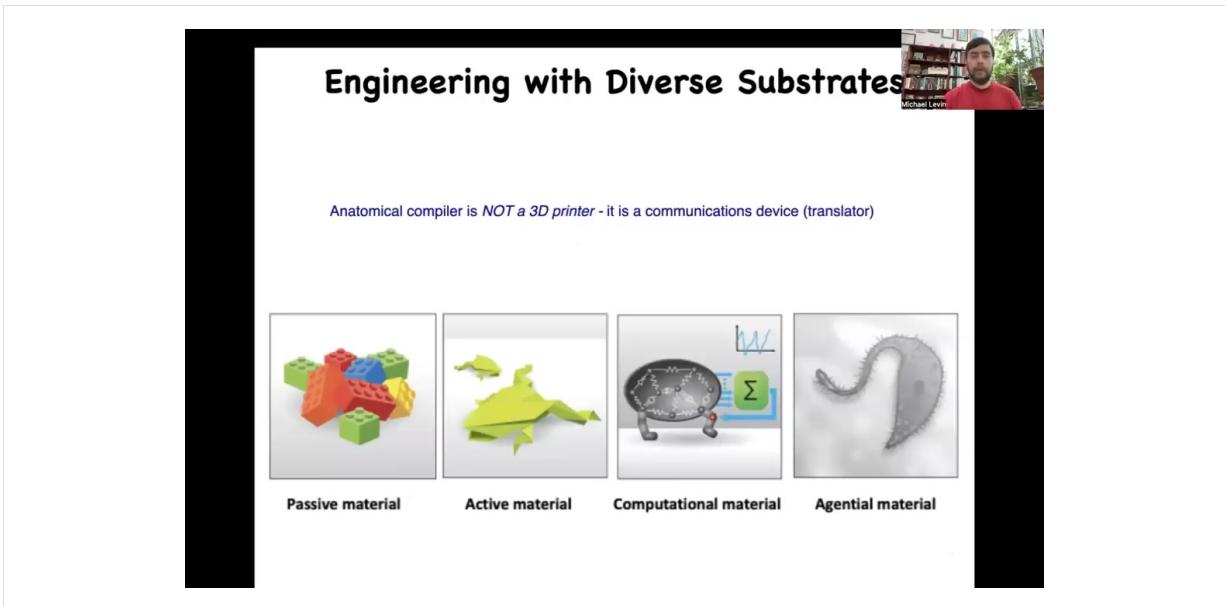
how to control emergence and scaling of cellular activity toward desirable complex, system-level outcomes?

Michael Lobo

So let's think about the end game of our field. What we would like to do is something called an anatomical compiler. This is an idea that encapsulates where we want to go with this, which is the complete control over growth and form. Someday you will be able to sit in front of a computer system, you will be able to specify the plant or animal or organ or biobot that you want, you could draw it at the level of anatomy, not at the level of molecular pathways. And if we knew how to do this, what the system would do is compile that description into a set of stimuli that would have to be given to cells to build exactly what you want to build.

The practical applications of this are obvious. If we had something like this, then birth defects, traumatic injury, cancer, aging, degenerative disease, all of these things would be solved. But there's a more fundamental problem here, which is that this is not meant to be a 3D printer. The point isn't to micromanage the position of cells or the expression of tens of thousands of genes. This is basically a communications device. It's a translator between your goals as the engineer and the goals of the collective, which is going to build something. The question is, how do we control what it is that they built?

Engineering with Diverse Substrates



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

Passive material Active material Computational material Agential material

Now, over the millennia, humans have been building with passive materials. And more recently, we've developed active matter and even computational materials. But in biology, we use something that's even deeper and more interesting, which is something we call an agential material.

We All Used to be Single Cells!

Lacrymaria olor

video by Charles Krebs

Lacrymaria = 1 cell
no brain
no nervous system

high competency
at cell-level
agendas

Normal Embryonic Development

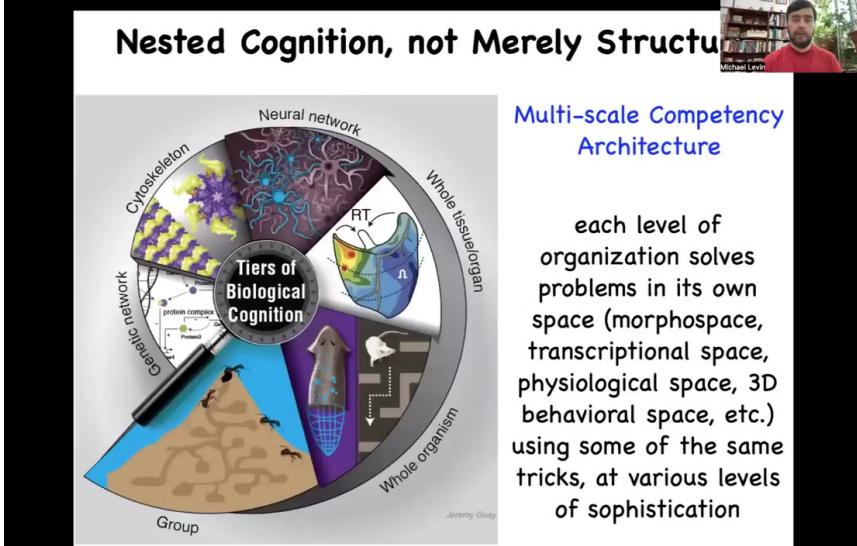
Protein A, Protein B, Protein C, Protein D, Gene E, Gene F

Michael Levin

Charles Le Brun

Now here's a single cell. This happens to be a free-living organism named Lachrymaria. You can see the amazing control that this thing has over its body and the competency with which it solves its local needs, all of its physiological, metabolic, anatomical, and so on. There's no brain, there's no nervous system. The single cell does all of the things that it needs.

The interesting thing is that we all come from a single cell. We all begin life as a molecular network inside an unfertilized oocyte. Eventually, we become something like this. We all used to be single cells at one point, and we need to understand this continuous, gradual process by which chemistry becomes mind. In the scale up from a system that is well described by the laws of chemistry and physics to a system that is really also amenable to very high level descriptions of behavioral science and psychoanalysis and these kinds of things, our job is to understand that scaling and to figure out what new competencies that gives us as engineers.



The diagram illustrates the concept of nested cognition in biology. It features a circular model with concentric layers, each representing a different level of biological organization. The innermost layer is labeled 'Tiers of Biological Cognition'. The outer layers, moving outwards, are labeled 'Genetic network', 'Cytoskeleton', 'Neural network', and 'Whole organism'. Each layer contains specific biological components and processes. For example, the 'Genetic network' layer shows a protein complex and a gene. The 'Cytoskeleton' layer shows microtubules and actin filaments. The 'Neural network' layer shows neurons. The 'Whole organism' layer shows a group of ants and a mouse. The diagram is signed 'Jeremy Guay' at the bottom. Above the diagram, the text 'Nested Cognition, not Merely Structure' is displayed, along with a small portrait of Michael Levin and the text 'Michael Levin'. To the right of the diagram, the text 'Multi-scale Competency Architecture' is written in blue, followed by a detailed explanation in black text: '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'.

Because of this, biology can be looked at as a multi-scale competency architecture, not just nested scales of size, like nested dolls of structure, but also function. All of these layers themselves are competent at solving problems in different spaces. All the way from molecular networks to subcellular components to tissues and organs and whole organisms and even swarms are able to solve problems in anatomical spaces, in gene expression spaces, in behavioral spaces.

Changing the Subject of Intelligence



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

Michael Littman

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

The stability of memories during brain remodeling: A perspective
Douglas J. Blackstone¹, Tal Moran^{2,3}, and Michael Littman¹

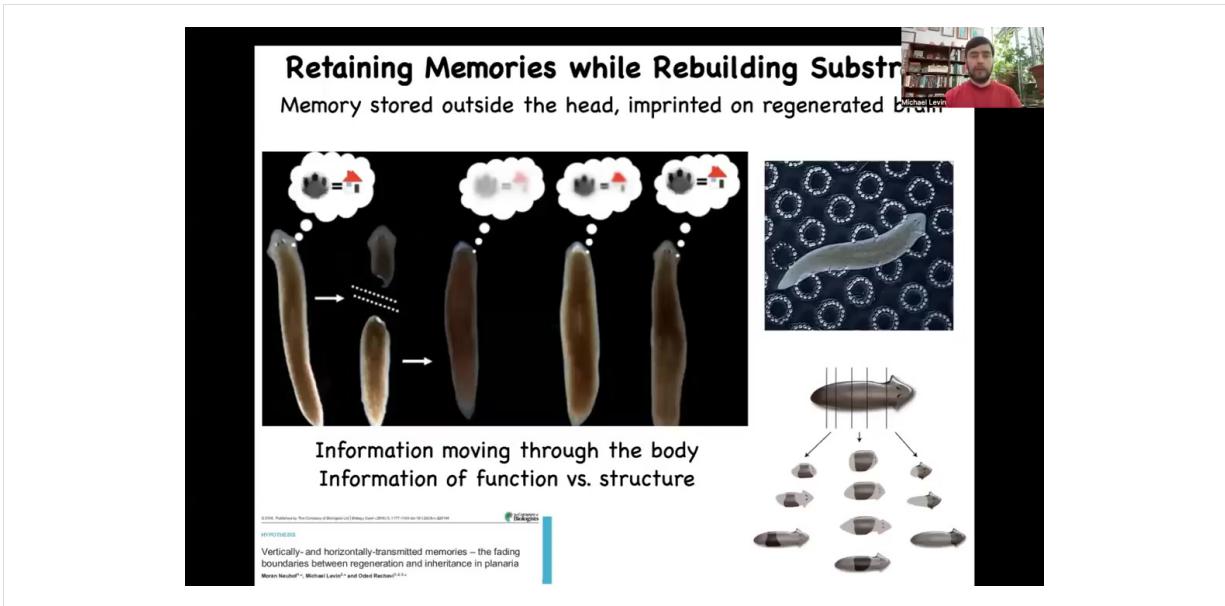
That gives rise to a number of very interesting properties. Our engineered constructs are still very far behind.

I want to describe some of the amazing things that biology does. Here's one example. This is a caterpillar, and caterpillars live in a two-dimensional world. They crawl around, they chew leaves, they have a little brain that's suitable for driving this, in effect, a soft-bodied robot without any hard elements. This animal has to turn into this animal, which is going to live in a three-dimensional world, fly around using a hard body design and drink nectar. It has a completely different brain.

Along the way, what happens during metamorphosis is that the brain is largely dissolved. Most of the cells are killed off, the connections are broken, except that we find, and this was work going back many decades, that if you train the caterpillar to eat leaves on a particular color disc, that memory persists in the butterfly. You might think that the amazing thing here is the question of where is the memory? Where do you store memories when the whole brain is being taken apart and put back together?

It's actually even much more interesting than that because the raw memories of the caterpillar, what types of muscle motions to activate to retrieve a reward on a particular stimulus, are completely useless to this animal because it doesn't eat the same thing. It doesn't move the same way. The controller is completely different. What actually happens to the information is not merely that it is stored, but it is remapped. It is remapped to a new creature that lives in a higher dimensional world and has a completely new life.

This idea of information being context sensitive and being interpreted by a new substrate is something that the material does on its own. We need to understand how this works.



In other models, such as these planarian flatworms, which I'll tell you more about in a few minutes, which regenerate, we can train the animals to find food on these bumpy substrates. This is place conditioning.

They learn to look for food there. You can cut off their head, which contains a centralized brain. The tail will sit doing nothing for about 8 days until it grows back a new brain and a new head. At that point, you can see that this animal, once again, remembers the original information. It's interesting to think about how information moves through the body. Wherever else it's stored, it has to be imprinted onto the new brain as the brain develops.

On-the Fly Adjustment to Novel Architectures

Ectopic eyes on tail provide vision!

1. Robust goal-directedness → buffers negative consequences of mutations, allowing evolution to explore positive pleiotropic effects (creativity)

2. Higher-level control circuits don't need to solve a huge-dimensional micromanagement problem – they can optimize in reward space

3. Biology is highly hackable!

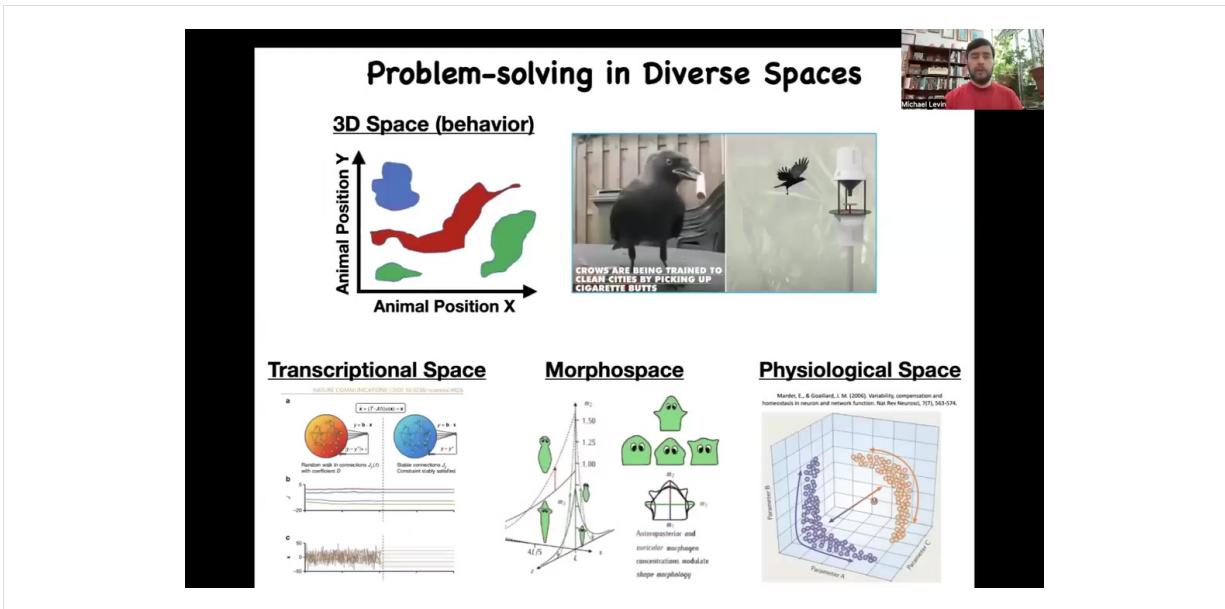
Brain dynamically adjusts behavioral programs to accommodate different body architectures

Behavioral Testing Device

And this kind of plasticity is also seen in vertebrates. This is a tadpole of the frog *Xenopus laevis*. Here's the mouth, here are the nostrils, the brain, the gut. There are two eyes missing here, but we place an ectopic eye onto its tail. These eyes make an optic nerve. The optic nerve comes out. It will synapse, for example, on the spinal cord here. It does not go to the brain. Yet these animals can see.

We know they can see because we built a device to automate the training and the testing of these animals in a visual assay. We train them to avoid or follow certain lights, and they can do quite well. This is telling us that this material does not need new cycles of evolutionary search. It does not need more generations of generate and test of blind mutation and then selection in order to adapt to a new sensory-motor architecture. No eyes in the brain. Now you've got this eye on the tail connected to the spinal cord. The whole thing works. The brain is able to use these signals.

This kind of plasticity has many implications for evolution, but during today's talk we're going to talk about what it means for the engineer.



So one of the key things that we have to do in order to start utilizing these amazing properties of the living material is to widen our ability to think about other problem spaces. Humans are pretty good at recognizing intelligence of medium-sized objects moving at medium speeds in three-dimensional space. So we see apes, we see some birds, maybe an octopus or a whale doing interesting things in three-dimensional space, and we can understand this as intelligence. But there are all these other spaces. So there's the space of gene expression, there's the space of physiological states, and what we're going to talk about most of all, the space of anatomical states. And in all of these spaces, you can think about different types of systems at different time scales, in different embodiments, navigating those spaces adaptively, that is making decisions, taking measurements, having memories, and trying to reach specific regions of those spaces.

Where do Anatomies Come from?

stem cell
embryonic
blastomeres

self-assembly

earthymission.com

Michael Levin

- 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?

So let's talk about anatomical space. We'll have time today to mostly focus on that. So all of us begin life as a collection of blastomeres, but if you take a cross-section through a human torso, you see something like this, incredible order. Everything is the right shape, the right size, the right orientation next to the right thing. Where does this incredible pattern come from? Arranging this pattern in the space of all possible configurations is a very specific journey. So how do the cells know where they're going?

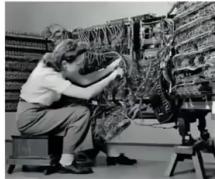
Now, you might be tempted to say it's in the genome, it's in the DNA, and that's what most people would say. But the problem with that is that we can read genomes now. We know what's in the DNA, and it's nothing like this. What's directly in the DNA are protein sequences. So they are specifications of the nanoscale hardware that every cell gets to have. After that, the cells have to engage in a process we call developmental physiology, which allows them to build something like this.

Now, that means we need to understand how do cell groups know what to make and when to stop? How could we convince them to repair or rebuild if something is missing? And as engineers, we would also like to know how much plasticity is there? Could we get them to build something completely different with the exact same genome? How reprogrammable is this hardware? What else is it willing to do? The information on the final outcome is no more in the genes than the exact structure of the termite mound or the spider web is encoded in the genome of these animals. All of these things are outputs of behavior. This anatomy is no less an output of behavior than are these constructions.

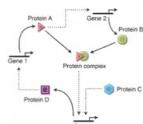
The State of the Art: biological information

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)



Where we are today in the biosciences is this. We're getting very good at manipulating cells and molecules, what the genes interact with, what the genes are, and so on. We're really a long way away from large-scale control of form and function. If somebody needs their limb back after an injury, we really don't know how to make them grow back. Or, we really don't even know simple things like why you have two bones in your forearm instead of one, despite all the molecular biology.

I think that what's happening here is that biomedicine has been stuck where computer science was in the 40s and 50s. This is what it used to look like to program a computer: you'd have to physically rewire it. The reason that today you don't have to get out your soldering iron every time you switch from PowerPoint to Microsoft Word is that we've understood the power of reprogrammable hardware, the power of using the information-processing competencies of your material to control it with stimuli at a higher level, not just rewiring.

Biomedicine, molecular medicine anyway, is largely focused on the hardware. Today, all the exciting advances are around genomic editing, protein engineering, pathway rewiring. It's all down at the level of the molecules. I think the reason that we're still not close to being able to control what we actually want to control, which is form and function, is because we've barely started to take advantage of the higher-level information processing, or the intelligence and the decision making, of the material.

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

- William James

Behavior

- Active
- Non-active (passive)

- Purposeful
 - Feed-back (teleological)
 - No feedback (non-teleological)
- Non-purposeful (random)

- Predictive (extrapolative)
- Non-predictive (non-extrapolative)

- 1st, 2nd, etc. order of prediction, self-reference

Michael Levin

Behavior, Purpose and Teleology
Arturo Rosenblueth, Norbert Wiener and Julian Bigelow
in: *Philosophy of Science*, 10(1943), S. 18-24.

When I say intelligence, what I mean is what William James defined it as: "the ability to reach the same goal by different means." This is a very cybernetic definition. It doesn't focus on how big your brain is or whether you are naturally evolved or engineered. It talks about what level of competency you have to call upon a varied bag of tricks to get to the same goal when things change. People have used this kind of ladder. Here's Rosenblueth, Wiener, and Bigelow with their ladder of three different kinds of systems going all the way from passive matter up to different levels of complexity of control and goal directedness to human level metacognition and so on. Intelligence doesn't mean that you know that you're intelligent and that you have this kind of self-awareness, it just means you have some degree of problem solving to get your goals met. What level of degree is the operant question.

Let's talk about what kind of intelligence the living material possesses, specifically in this case, an anatomical space. I could tell you many stories about intelligence and problem solving in other spaces. One thing we know is that developmental self-assembly is reliable. You go from a single cell to a normal human target morphology. Most of the time it works correctly, but that isn't what we mean by intelligence. We don't mean an increase in complexity. It's not enough to say that something becomes complex because we all know there are fractals, there are cellular automata, there are many systems that follow simple rules in a feed-forward, open-loop fashion, and something complex results. That's just the beginning of the story. What's more interesting is the fact that if you cut an early embryo into pieces, you don't get half bodies, you get perfectly normal monozygotic twins, triplets, and so on. You can reach the ensemble of goal states corresponding to a normal human from all kinds of starting positions and avoiding some local minima. That gives you the

ability to get to where you're going from different starting positions. That's one type of competency here.

Slide 14 of 47 · Watch at [15:01](#)

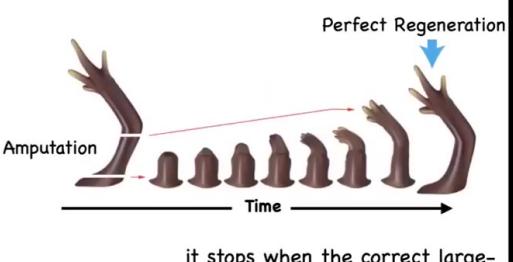
Same anatomy, from different starting st



Michael Levin

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





Perfect Regeneration

Amputation

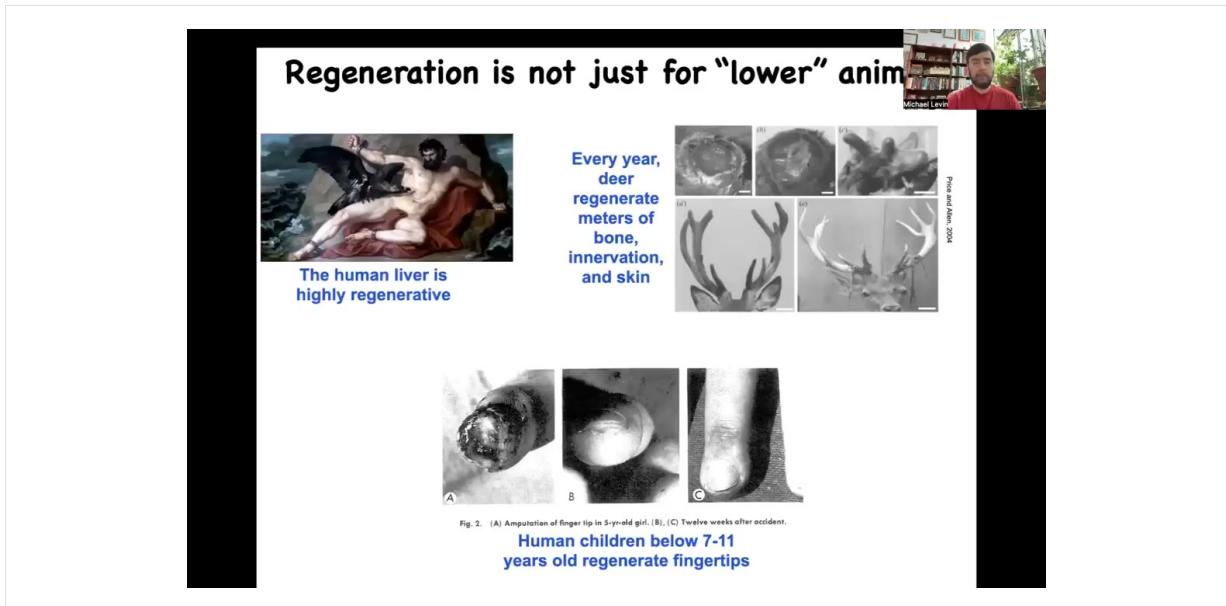
Time

Anatomical homeostasis:

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

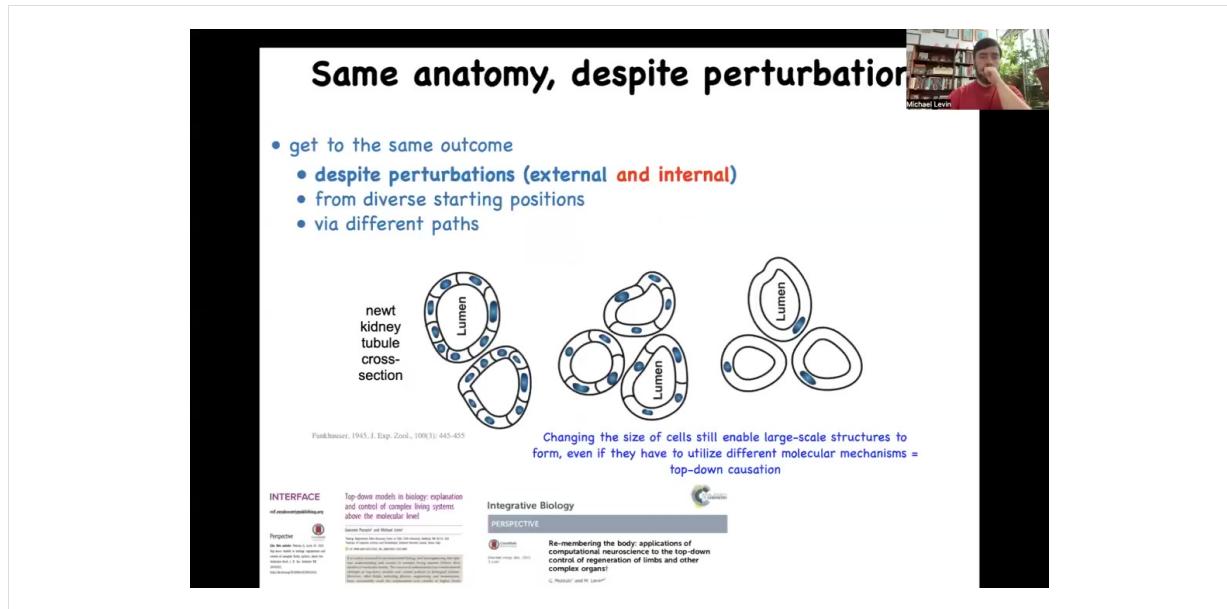
The diagram illustrates the process of limb regeneration in an axolotl. It shows a limb being amputated at the bottom. Over time, the limb regenerates through a series of stages, represented by a sequence of limb shapes. A blue arrow points upwards from the final regenerated limb towards the text 'Perfect Regeneration'. A red arrow points from the text 'Amputation' to the initial stage of regeneration. The text 'Anatomical homeostasis:' is written in red, and below it, in black, is the explanatory text: 'it stops when the correct large-scale setpoint (target morphology) has been reached'.

Another type is the idea that later on in adulthood, some animals maintain this amazing capacity for regeneration. This is an axolotl. These animals regenerate their limbs, their eyes, their jaws, their spinal cord, portions of the heart and brain. If you amputate anywhere along this limb, they will regenerate. The cells will grow, they will undergo morphogenesis, they will regenerate the limb, and then they stop. The most amazing thing about regeneration is that it knows when to stop. How does it stop? It stops when a complete salamander limb has been achieved. You can think about this as a means-ends analysis or basically an error-reduction scheme. It's also known as homeostasis. This is anatomical homeostasis. This is the goal state. When you deviate the system from that goal state, it will work really hard until the error is within acceptable limits and then everything stops. That brings up an obvious question: how does it know what the correct pattern should be?



I should point out this isn't some special thing about worms and frogs, although I am going to spend most of my time talking about those models. This is very widespread. Adult humans can regenerate their liver. Human children can regenerate their fingertips below a certain age. Deer, large adult mammals, can regenerate huge amounts of bone, vasculature, innervation, skin every year as they regrow their antlers. They grow about a centimeter and a half of new bone per day to make the same pattern every year.

Same anatomy, despite perturbation



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

newt kidney tubule cross-section

Funkhouser, 1945. J. Exp. Zool., 109(3): 445-455

INTERFACE and environmentpublications.org

Top-down model in biology: regulation and control of complex living systems above the molecular level

Genetic Puzzles and Molecular Cues: How the Newt Kidney Tubule Grows and Gains Its Distinctive Structure

Integrative Biology PERSPECTIVE

Re-imaging the body: applications of computational neuroscience to the top-down control of regeneration of limbs and other complex organs

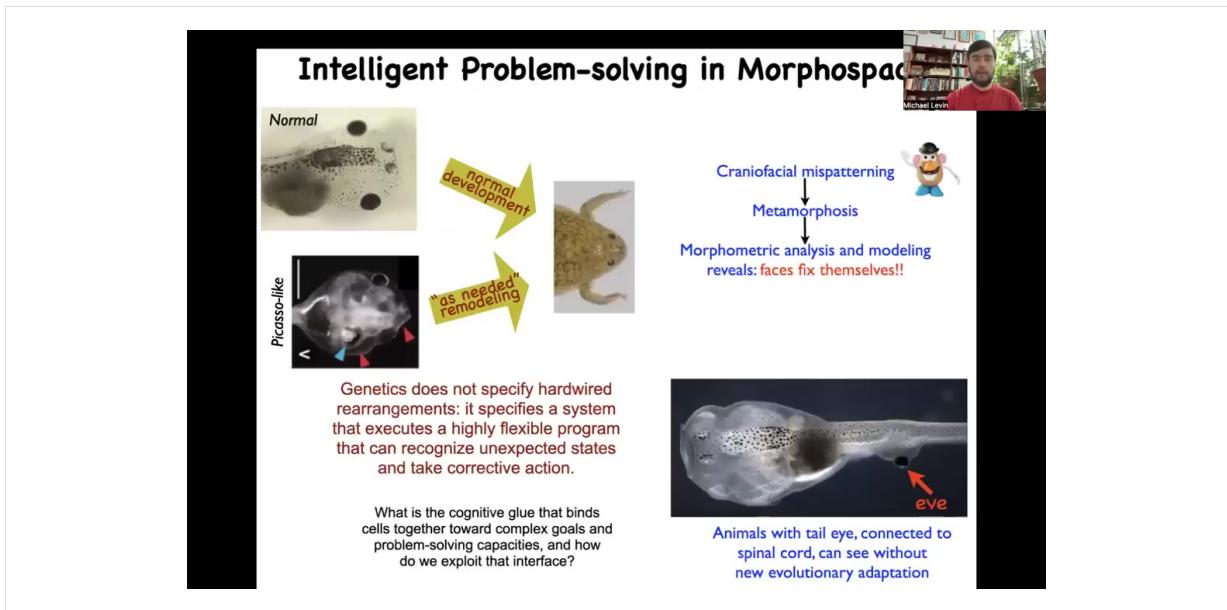
Michael Levin

One of the most important things to realize about this kind of anatomical homeostasis is that it isn't just about external injuries. The stories that I've told you are all about damage, external kinds of damage. But there's something much more profound here, and I want to illustrate this with one of my favorite examples, which is the newt kidney tubule. If you take a cross-section of the newt kidney tubule, you see that about 8 or 10 cells work together to create this tubule, and then there's a lumen inside. One thing you can do with these animals is increase the copy number of their genome. You can make polyploid newts that are not $2n$, but $3n$, $4n$, $5n$, and so on. If you do this, the first amazing thing is that you still get a perfectly normal newt. The second thing you find is that when you increase the amount of DNA, the cells actually get bigger to maintain a proper ratio of cell size to nucleus size. And that means that in order to make the same size newt, which is what they do, you have to have fewer number of larger cells. The cell number adjusts to the cell size. The most amazing thing of all is that if you make the cells truly gigantic, and I believe this is like $5N$ or $6N$ newts, the cells get so big, what they have to do is bend around themselves to make the same structure. This is a different molecular mechanism. This is not cell-to-cell communication as before. This is now cytoskeletal bending. This is an interesting kind of top-down causation. In the service of a particular anatomical goal, the cells are able to call up different molecular mechanisms to get their job done.

A couple of things going on here. One is that there's that definition of intelligence, meaning they reach the same goal by different means when the circumstances change. This is a novel scenario. They've never seen this before. Just think about what this means as an agent coming into the world. As a newt coming into this world, you cannot overtrain on your evolutionary priors. You can't assume how many copies of your DNA you're going to have. You can't assume the number or the size of your cells.

You have to get your job done despite not only variations in the environment, not only potential injuries, but in fact variations of your own parts. When we build robotics and computerized systems, we spend a lot of effort making sure that the parts are exactly what we think they are, the tolerance on the parts is extremely high, they never change, and then we can engineer. Biology does the exact opposite. It commits right from the beginning to the idea that both evolutionarily and ontogenically, your parts are unreliable, they're going to change, there's going to be noise, there's going to be death, there's going to be error. Despite all of that, it has to work. It's a completely different way to engineer. It uses this multi-scale competency architecture, this idea of goal directedness and decision making at every level to make this whole thing work.

Slide 17 of 47 · Watch at [19:38](#)



The final example that I'm going to show you is something we discovered a few years ago. Here's a tadpole with some eyes and some nostrils and the mouth here. In order to turn into a frog, what they have to do is rearrange their face. All these organs have to move, the eyes have to move forward and so on.

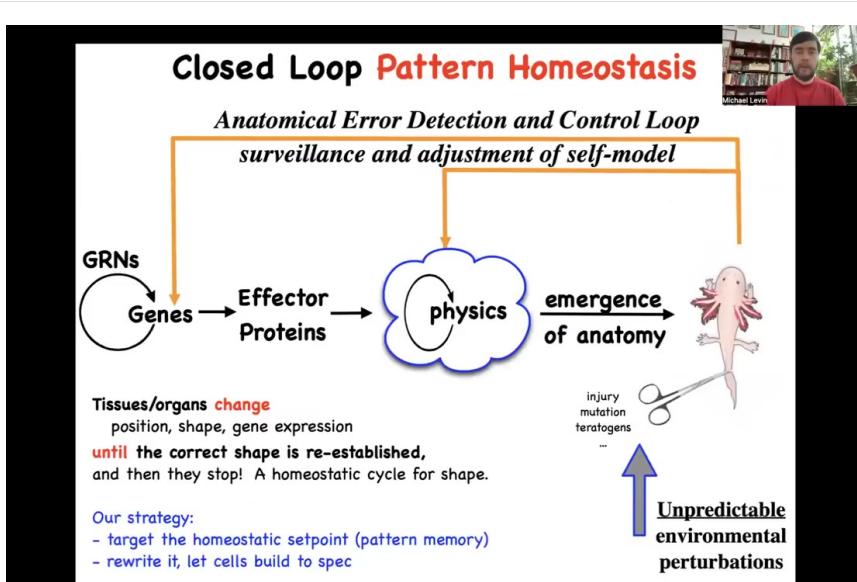
It used to be thought that this was a hardwired process. After all, every tadpole looks the same, every frog looks the same. To go from a tadpole to a frog, you just know what direction and how much every part is going to move. We wanted to test this hypothesis. The intelligence of any system cannot be determined by philosophical pre-commitments. You can't just decide how intelligent things are, nor can you

determine it from pure observation. You have to do experiments. You have to do perturbative experiments to really know.

What we did was we said, okay, what is the competency of this thing? Let's scramble all the organs. Here you can see an example. We call these Picasso tadpoles. You've got an eye on the top of the head. The mouth is off to the side. Everything's scrambled. It's a big mess. Then what you find is that these things give pretty normal frogs because all of these organs will move in novel paths to get to where they need to go. In fact, sometimes they go too far and have to come back a little bit.

What you see here is that the genetics does not specify a hardwired set of movements, it gives you a system that is an error reduction scheme. It has a set point, it has an error reduction capacity. This goes to my major claim that what evolution produces is not solutions to specific environments, it produces problem-solving agents. That's what it produces.

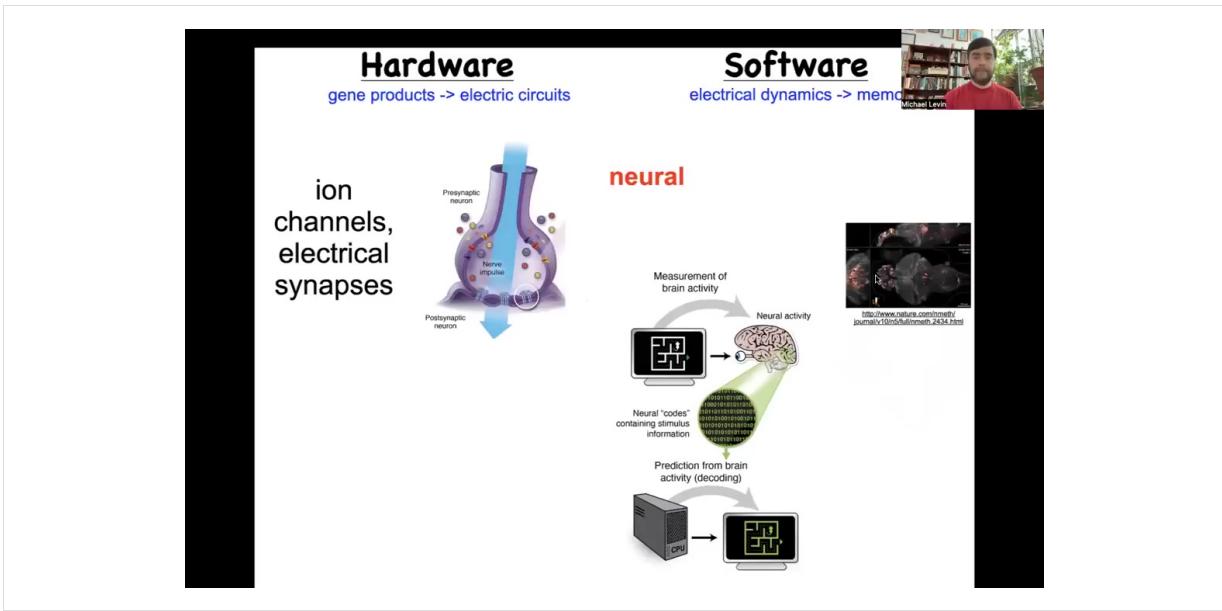
Slide 18 of 47 · Watch at [21:23](#)



You have here this main loop that you see in a developmental biology textbook that goes from genes and gene interactions to various effector proteins, and then this process of emergence. And all of this is true. It does happen. This complexity does emerge from these local interactions. But there's something really important missing here, which is this pattern homeostasis scheme. This idea that this is not purely a

feedforward process. There's actually a set of feedback loops that, if you get deviated from this goal — that's injury, mutations, teratogens, and so on, depending on the time scale that you're talking about — then these other mechanisms kick in and try to reduce the error.

So this makes a very strong prediction. What this suggests is that there should be an explicitly recorded set point, that is, the cells should have access to some kind of biophysical structure that actually contains that large-scale set point, and that if we were to change that set point, we could control what the system builds without having to change the hardware. This is important because under traditional approaches, if you want to make that change up here, you've got to figure out which genes to edit. And there are a few low-hanging fruits, single gene diseases and single gene traits and so on, where it's pretty clear what to change. But for the vast majority of things we're interested in, this process is not reversible. We're not able to figure out what genes to edit. In that case, it would be convenient if we could find the set point and directly edit that information structure and let the cells do what they do. We've been doing this approach for some years now and following the strong predictions of this model, which is that it should be possible to find it, should be possible to decode it, and it should be possible to rewrite it if it's true that these pattern memories exist.



Where would we find this? Is it crazy to think that a collection of cells could store a memory of what to do? Neuroscience for centuries has been studying exactly that, a collection of cells that stores memories of what to do.

The way they do it is by this interesting architecture where you have hardware that consists of cells communicating with each other through electrical synapses by electrical signals that they establish through the movement of ions in these ion channel proteins. That's the hardware. The software, this group made an amazing video of a brain of a living fish as the fish is thinking about whatever it is that fish think about. The idea is that the process of neural decoding, so this is the commitment of neuroscience, is that the cognitive goals, preferences, beliefs, or behavioral repertoires, everything that exists in the mind of this animal can be decoded from the electrophysiology that goes on in the nervous system.

It turns out this is an extremely ancient discovery by evolution. In fact, evolution found that electrical networks are good for scaling and processing information back around the time of bacterial biofilms. This is very old. Every cell in your body has ion channels. Most cells have gap junctional electrical synapses with each other, and they build networks.

Whereas neuroscientists try to decode the electrophysiology of neural cells to understand how they move your body through three-dimensional space, could we do the same thing? Could we understand how information propagates through electrical networks to determine how they help you navigate anatomical space? We use a lot of the same tools as neuroscientists do, practical tools as well as concepts, to try to understand this because the tools actually do not distinguish between these two cases. Neuroscience is not just about neurons. All of the kinds of concepts that we learn in

multi-scale neuroscience, including things that have to do with very complex system level behaviors and perceptual illusions, carry over. In fact, we've made AI tools that help scientists make that mapping and design new experiments.

Slide 20 of 47 · Watch at [25:48](#)

How we detect and model bioelectric signals

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

Voltage reporting fluorescent dye in time-lapse during *Xenopus* development

Dany Adams

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

More Positive Cell (Higher pH; Depolarized) More Negative Cell (Lower pH; Polarized)

V (Ventral) K⁺ KCNQ1 KCNE1 K⁺ K_{ATPase} K⁺ K_{ATPase}

L (Left) R (Right) D (Dorsal)

Gap Junction

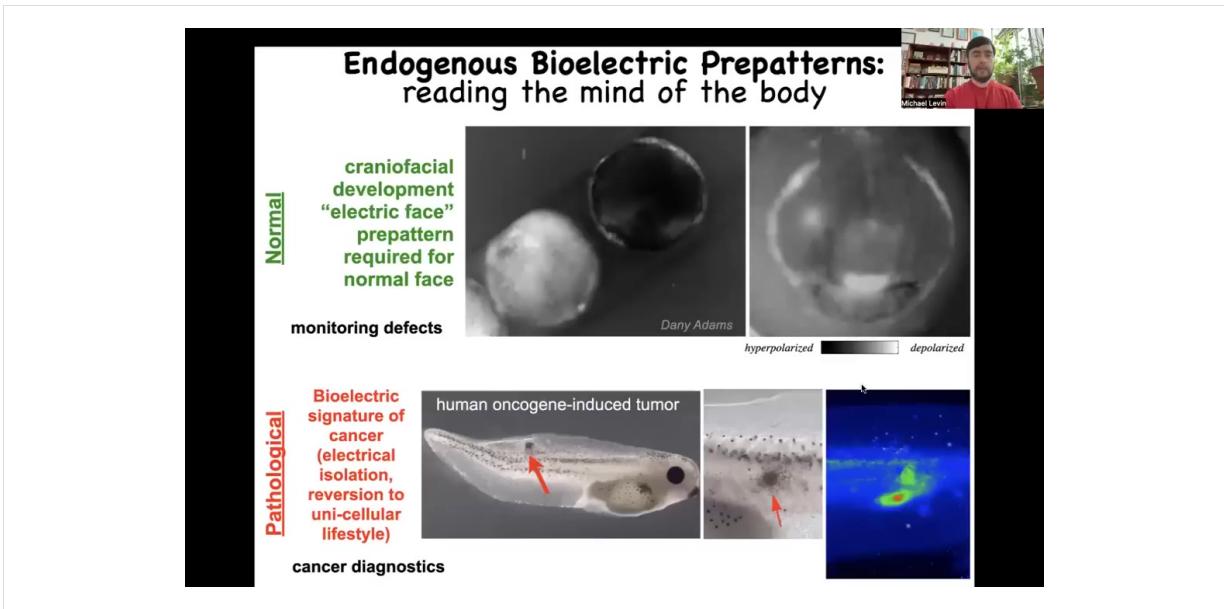
H⁺/K⁺-ATPase V-ATPase

Si-HIF

Michael Levin

In order to show that this actually works, we had to develop some new tools. We had to develop tools to read and write the electrical information of these networks so that we can ask, what do they think about? We know what brain networks think about mostly. What do these somatic networks think about?

First we developed voltage-sensitive fluorescent dyes that allow us — this is an early frog embryo in time-lapse. This is not a simulation. This is real data. We can try to understand how all of the cells are talking to each other in terms of the electrical patterns that are there. We do a lot of computer simulation to try to understand where these currents come from given the ion channels and pumps that are expressed.



I'm going to show you an example of a couple of these patterns. This is something we call the electric face. This is one frame from a voltage dye movie of an early embryo putting its face together. What you see in this frame is that long before the anatomical structures start to be developed and the genes come on to regionalize the face, already these cells are showing a pre-pattern of what's to come. Here's where the animal's right eye is going to be, here's the mouth, here are the placodes. This is a pre-pattern memory of what's going to happen. It's a crucial one because if you disturb this memory, then you change the anatomy of the face. I'll show you that momentarily.

There are also pathological patterns which you can induce by introducing, for example, a human oncogene. Here's some cells injected with a human oncogene. They make a tumor eventually that metastasizes. Before that, you see very early on that these cells have electrically disconnected from their environment and are basically rolling back to their ancient unicellular roots. They're going to be single-cell amoebas and treat the rest of the body as external environment because they are no longer able to connect to that electrical network that remembers what to do at a large scale.

Manipulating Bioelectric Networks' Cont

Non-neuronal cell group

hyperpolarized ← → depolarized

Gap Junctions (electrical synapse)

V_{mem}

Ion channels (setting V_{mem})

Na^+ K^+ Cl^-

Intracellular

Extracellular

V_{mem}

Neurotransmitter (moving via V_{mem})

- Transporter or receptor mutant overexpression
- Drug agonists or antagonists of receptors or transporters
- Photo-unmasking of neurotransmitter

Tools we developed
(no applied fields!)

- Dominant negative Connexin protein
- GJC drug blocker
- Cx mutant with altered gating or permeability

Synaptic plasticity

- Dominant ion channel over-expression (depolarizing or hyperpolarizing, light-gated, drug-gated)
- Drug blocker of native channel
- Drug opener of native channel

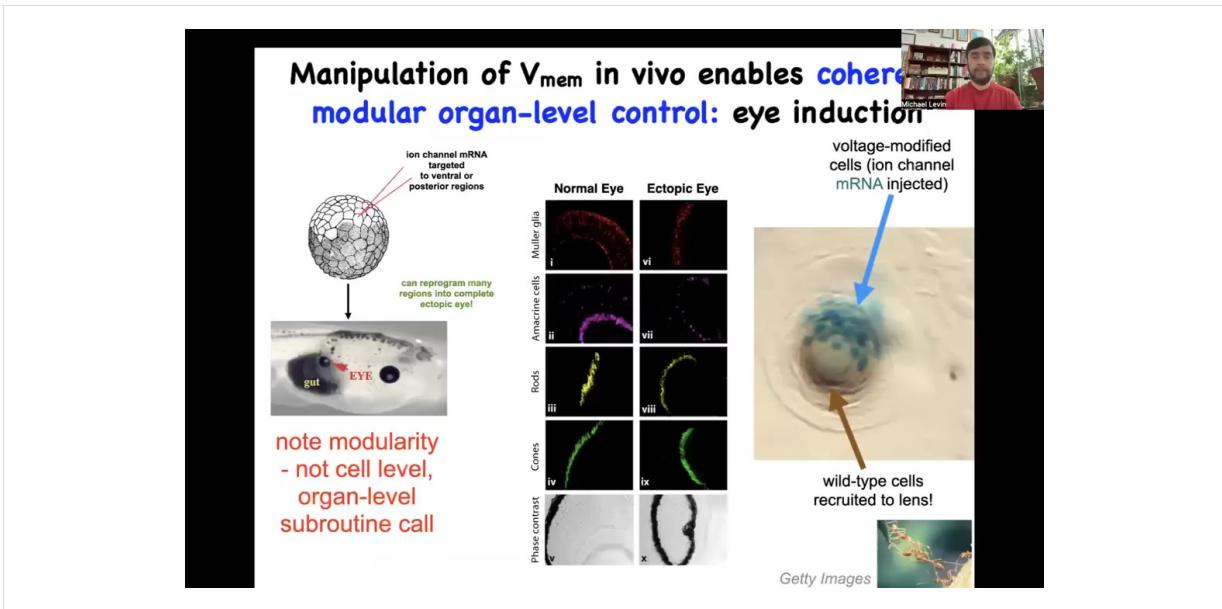
Intrinsic plasticity

Michael Levin

So tracking patterns is nice, but what's really critical are functional tools. So you need a way to manipulate the bioelectrical information in the individual cells and/or to control which cells talk to which other cells through these gap junctions.

For this, we do not use any kind of applied fields. There are no applied electrodes. There are no magnets, no waves, no electromagnetic radiation. What we do is the same thing that neuroscientists do, which is to manipulate the interface that the cells use to talk to each other. That means we can open and close these gap junctions using drugs, by using mutant channels, using optogenetics. This is all molecular physiology.

So the idea is that can we develop a neuroscience beyond neurons to understand what the electrical patterns are and control them to read and write information into the collective intelligence of these cells. Our goal is to treat morphogenesis as the behavior in anatomical space of a collective intelligence. Just like you and I are a collective intelligence of neurons and some other cells, the morphogenetically active body is also a collective intelligence trying to make its way through anatomical space. And it's held together by the same cognitive glue that holds us together, by electrical communication between cells.



Let me show you some examples. Why do we think that these electrical states actually store behavioral patterns in morphogenesis? First I showed you a little electrical eye spot, the thing that causes eyes to form in the early phase. We asked if we could reproduce that pattern somewhere else. One thing we did was inject potassium channel RNA into a different region, in this case a region that was going to give rise to the gut. If you do that, it makes eyes and these eyes can be perfectly normal, exhibiting the lens, retina, optic nerve, all the same stuff.

Notice a few interesting things. First of all, this is a highly modular trigger. That is, we do not provide all the information needed to build an eye. We provide a simple pattern of voltage. We don't tell the cells what stem cells to have or what the pattern of a particular eye needs to be. In fact, we have no idea how to do that, but that's okay because the tissue does. What we're doing is using a high-level subroutine call to tell the cells that an eye belongs here. And then we don't micromanage the molecular states. We let the cells figure it out.

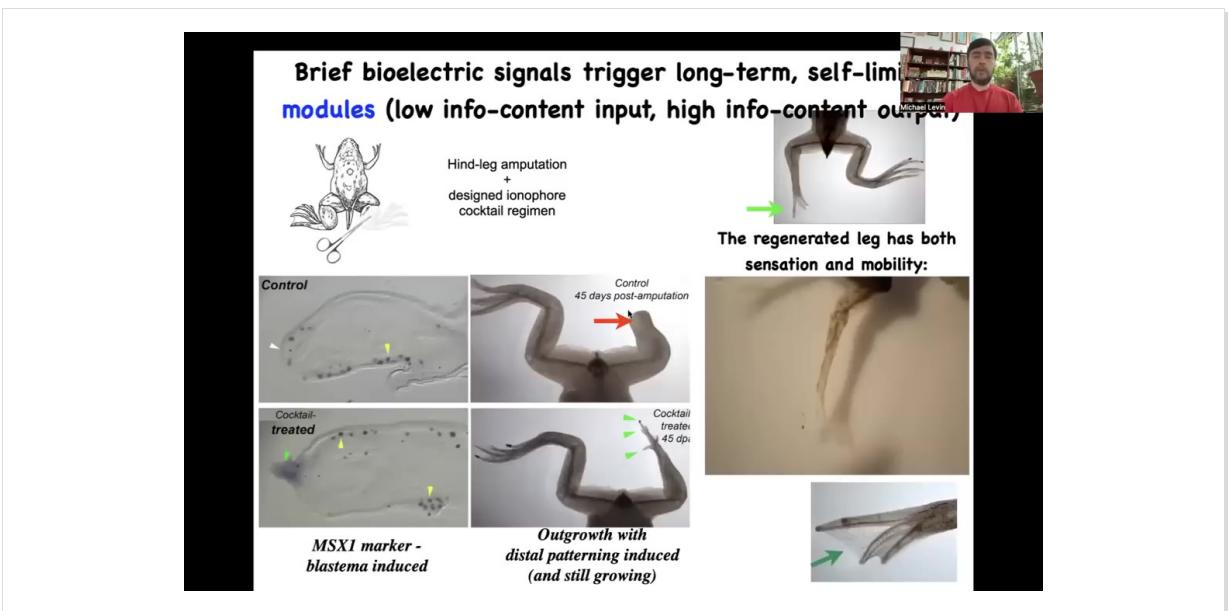
So not only is the bioelectric signal functionally causal, it controls what organ you get at a particular location. It's not just an epiphenomenon. It's a very modular high-level subroutine call that can pattern at the level of organs, not at the level of gene expression or at the level of a single cell fate. This is not about telling stem cells what to do. This is about dialing in very large scale patterns. Big, big movements in morphogenetic space, not tiny states.

The other thing that's fascinating about this is, look at this example here. This is a lens of an eye sitting out in the tail of a tadpole somewhere. And the blue cells, this is the lineage label. These are the ones that we actually injected with our potassium channel. There's not enough of them to make a proper eye. So what they do is they recruit all

their friends, they recruit these local cells to help them complete the journey. And that is something that we didn't have to teach them. They already do that. This is secondary instruction. We tell these cells make an eye. They do the computations that lead them to believe that there's not enough of them and that they should recruit their neighbors. So that is something that other collective intelligences do, like ants or termites. When a few of them find something that's too big to move, they recruit their neighbors to help. So this is something that is apparently general to various collective intelligences.

This idea of using high-level triggers is a regenerative medicine roadmap. In frogs, which normally do not regenerate their legs the way salamanders do, 45 days after amputation, there's nothing.

Slide 24 of 47 · Watch at [32:03](#)



We designed a bioelectrical cocktail where we can induce an electrical state that immediately triggers pro-regenerative genes such as MSX1. Then by 45 days, you already have some toes, you've got a toenail, eventually a pretty respectable leg that's touch sensitive and motile. You can see the pattern. In our latest experiments, 24 hours of stimulation with a wearable bioreactor and this drug leads to 18 months of leg growth. We do not micromanage the states. We don't tell it how to build a leg. This is part of the competency of the material.

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*
Ana Dominguez*, Justin A. Garry*, Carla Heras-Rescorla**, Michael Levin*, David Kaplan*

Wearable Bioreactors for the Study of Limb Regeneration

Frog Device Schematic
Outer silicone sleeve
Silk hydrogel insert
10 mm

Mouse Device Schematic
Metal protective ring
Outer attachment sleeve
Polymerizable scaffold insert
10 mm

Experimental Design

Feedback

Wearable Bioreactor

Biological Outcomes

David Kaplan's lab, Tufts BME

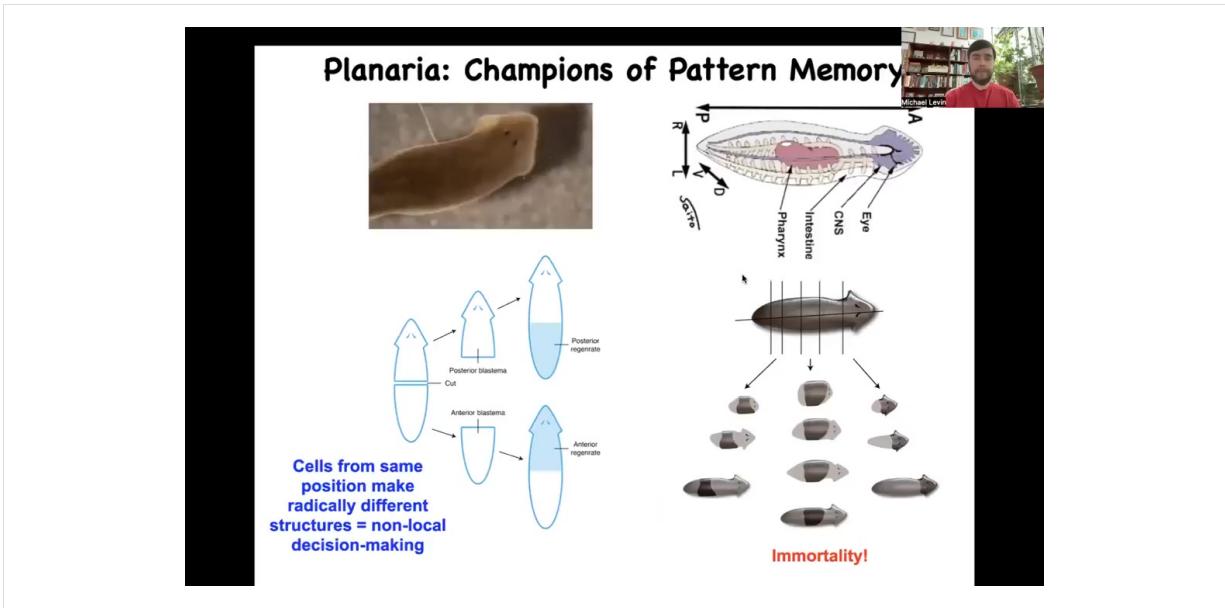
Ion channel modulators

Jay Dubb

Disclosure:
Morphoseuticals Inc.

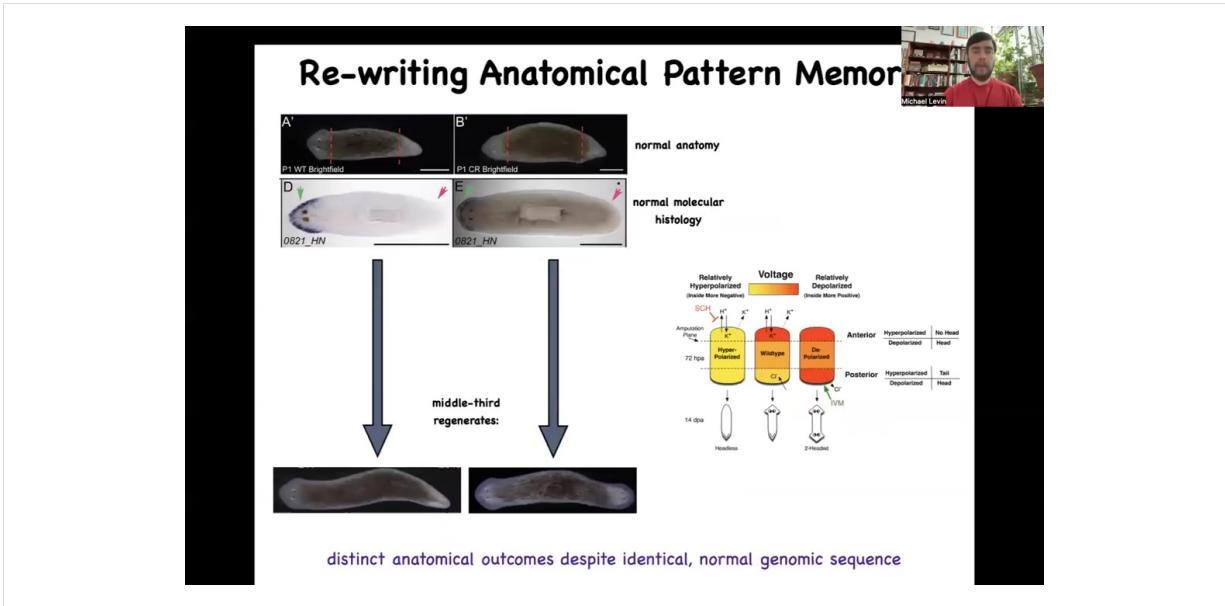
Michael Levin

Our job is to trigger it. I have to disclose that David Kaplan and I are co-founders of this company, Morphoseuticals, Inc. What we're trying to do is develop this technology using his bioreactor and our bioelectric payload to try to trigger this in mammals.

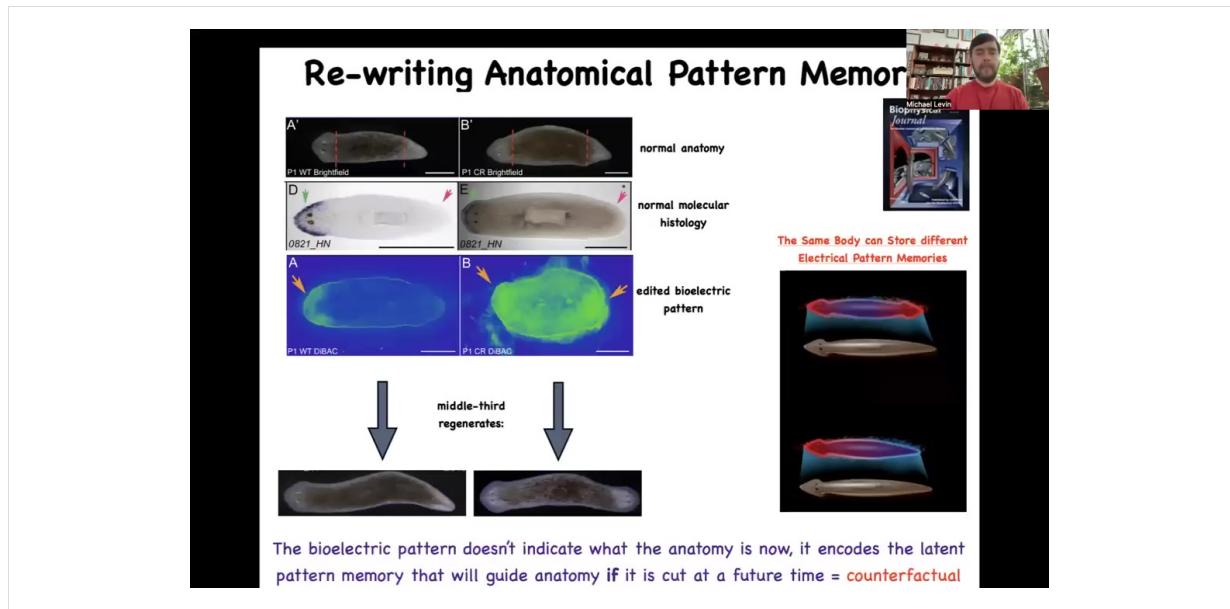


I want to now switch to a different organism and tell you the next part of the story, which has to do with actually observing and rewriting these patterned memories. This is a planarian. These flatworms are incredible because among other things, you can cut them into many pieces. Each piece has this holographic property. It remembers exactly what a whole planarian is supposed to look like, and then you get normal worms.

They're also immortal. If anybody's interested in that, you can ask me after the talk. We can talk about why these asexual worms have no aging. There's no such thing as an old asexual planarian, which is remarkable. It's because of their incredible morphogenetic control.



The thing about these planaria is that extremely reliably, if you cut off the head and the tail and you just have this middle fragment, 100% of the time it gives you a one-headed worm. How does it know how many heads it's supposed to have? How does this fragment know to make a head here, but not to make a head there? It's not just position, because this fragment at the exact same position is going to make a head here. How does this middle fragment know what to do? How do you know how many heads you're supposed to have? We discovered this interesting electrical circuit that determines this.



And the way it works is this: there's a voltage pattern that is a standing pattern of resting potentials present in this tissue that says, and this is one of the first ones we've decoded, so that's why I'm showing it to you, one head and one tail. And that pattern can be rewritten. We now have a way using ion channel drugs to rewrite that pattern so that we can say two heads. This is pretty messy. The technology is still being worked out, but you can see two heads here. If you do that, you get a two-headed animal. This is not Photoshopped. These are real.

This electrical map is not a map of this two-headed animal. This electrical map is a map of this perfectly normal looking one-headed animal. The anatomy is normal. The molecular biology is normal. It has head markers in the head, no head markers in the tail, but it has the incorrect memory of what to do if it gets injured in the future. The set point for anatomical homeostasis has been rewritten. That doesn't matter until it gets injured, which is why there's a mismatch between this pattern and the actual anatomy. If it gets injured, this is what it's going to do.

A couple of interesting things here. First of all, this is a counterfactual memory. It is like the kind of mental time travel you see in neuroscience where brains are good at remembering things that are not happening now because they either happened before or they might happen in the future. That's the beginnings of that kind of capacity. This creature can store at least two different ideas of what a correct planarian is going to look like if it gets injured. And those multiple different representations can live in the same kind of body. This is something else that happens in brains. You don't need to change your brain genetics in order to learn something new. It's amenable to multiple different goal states. That's what you're seeing here.

At the beginning, when I told you that there was a recorded memory of what the shape is supposed to be so that all the cells can agree and cooperate to build it, this is it. You're looking at the morphogenetic memory of how many heads there should be in planaria.

Slide 29 of 47 · Watch at [36:16](#)

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

Michael Levin

Basic properties of memory

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

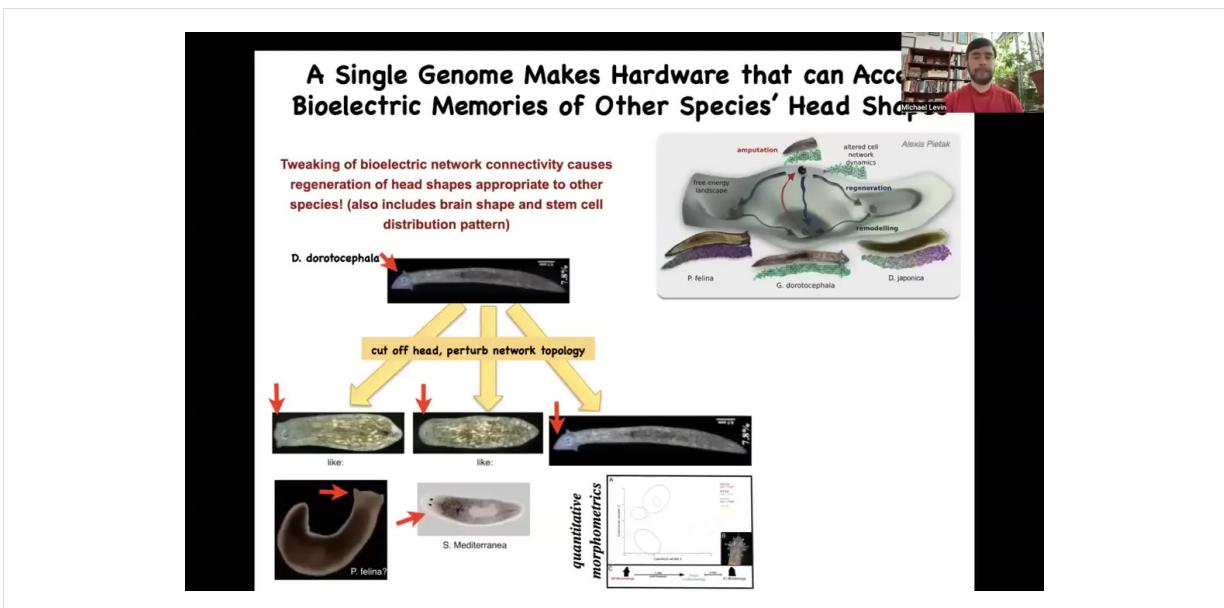
... (ellipsis)

I call it a memory for many reasons. One of them is that it has all the properties of memory. So if we take this two-headed worm and we cut off the primary head, we cut off this ectopic secondary head in plain water, no more manipulations of any kind, we let it regenerate. You might think that it should go back to normal because the genome is normal. We haven't touched the genetics. The hardware is completely wild type. It makes another two-headed worm. If you recut again and again and again, it continues to make two-headed ones. The memory of what to do, like any good memory, is long-term stable. We think it's permanent until we change it back, and we do know now how to change it back. Once you change it, it keeps. You have long-term stability, but it's rewritable if you know how to rewrite it. It's got latency, which I showed you, and it has these two possible outcomes. You can see these two-headed planaria hanging out. This is what happens.

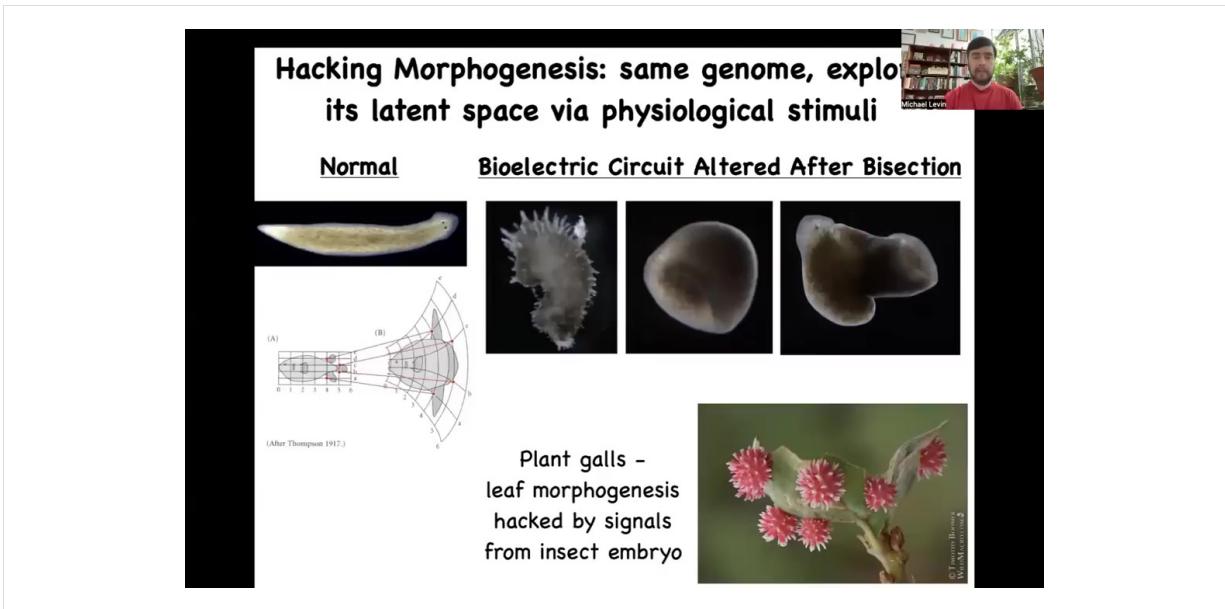
So what we are doing now is making computational models of the electric circuit and all of the different states that lead to the different anatomies and the properties of that circuit, such as pattern completion, which you see in a lot of connectionist models of

neural networks, where if part of the information is erased, how does the network recreate the entire information? We're trying to merge that with models of active inference to understand how the collective actually makes decisions on an energy landscape, which is an anatomical energy landscape that is powered by bioelectric effectors that the system can trigger to get to where it's going.

Slide 30 of 47 · Watch at [37:58](#)



It's not just about head number. For example, what about the head shape? This animal with a nice triangular head can be caused to make the heads of other species. It can make flat heads like a *P. falina*; it can make round heads like an *S. mediterranean*. The distance between these species and this guy is about 100 to 150 million years. This hardware has no trouble visiting the other attractors in anatomical state space that are normally occupied by these other species; you can go there if the bioelectric network guides you there. Not only would you have a normal head shape for these other species, but you would also have their brain shape, which is quite different, and their distribution of stem cells, specific to these other species, without any genetic change. You're starting to see the incredible reprogrammability of this hardware. It's willing to do many interesting things.



In fact, we can make worms that don't even look like flatworms at all. There's these crazy spiky forms, cylindrical things, these hybrid shapes. You can explore regions of the latent space of that circuit that evolution normally avoids. These are not particularly ecologically competitive, but the cells don't mind making them. The latent space of possibilities is quite huge.

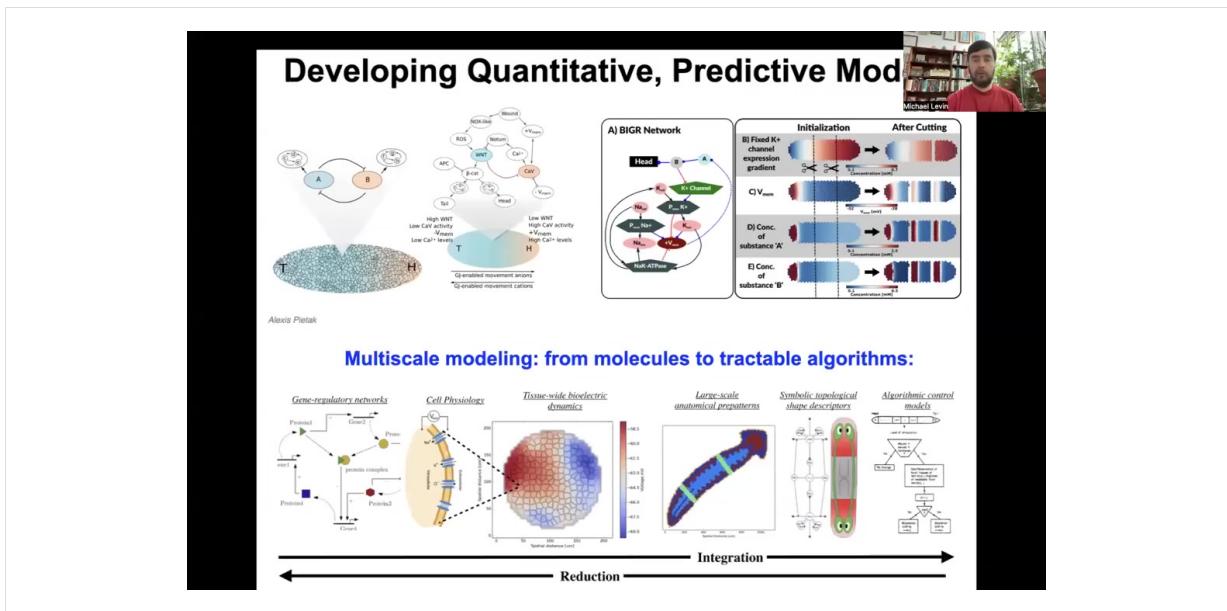
When you think about bioengineering and people say there are developmental constraints and there are things you will never be able to build, I'm not sure that's true at all. Of course, there are developmental constraints, but I think that those are more constraints of our imagination and our knowledge than of the cells themselves. I think we're imposing that on the cells. I think they can build almost anything.

In fact, this thing right here, this is an amazing example. There is a plant genome that normally very reliably makes these flat leaves. You would look at those leaves and you would think that's what the genome knows how to do. It knows how to make this nice flat structure. But there's another bioengineer around, not us, these wasps. What the wasps do is put down an embryo and some signals, and those signals hack the morphogenetic cues of the plant. Just as we did with the frog and the worm, they hack these cells to get them to build something completely different, this crazy round, spiky red thing. If not for them prompting these leaves with new stimuli, we would have no idea that these cells are even capable of doing this. Who would have known? If you hadn't seen this, who would have known that leaf cells can be prompted to do something like this? Again, not by changing the genome, but by providing signals and stimuli the way that we want to do with our anatomical compiler.

Biology is eminently hackable. I think living things hack each other all the time. This is what we as engineers need to understand. We need to understand the interface, and

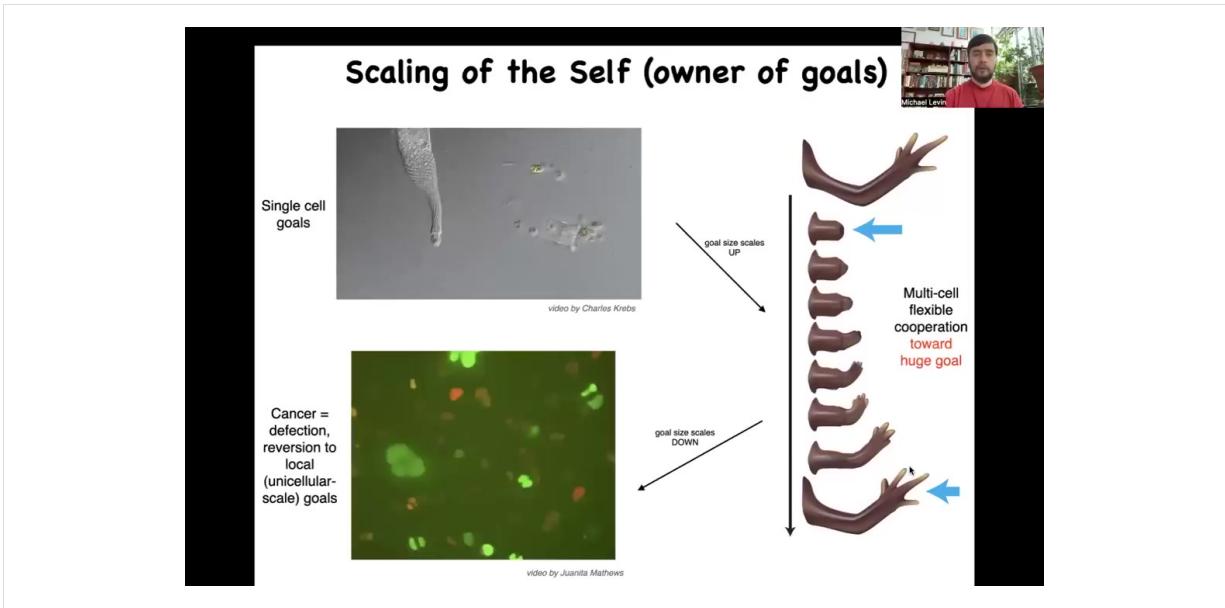
we need to understand the competency of our material so that we can begin to collaborate with it on building things and not try to micromanage it.

Slide 32 of 47 · Watch at [40:55](#)



So we build tools, from these kind of molecular biology things, where we understand which channels are expressed where, through bioelectric simulators to higher level kinds of models of cellular error detection and minimization decision-making, so that eventually you get this kind of algorithmic model of how these anatomical decisions get made so that you can interact with it and hopefully be able to control anatomy in a much better way.

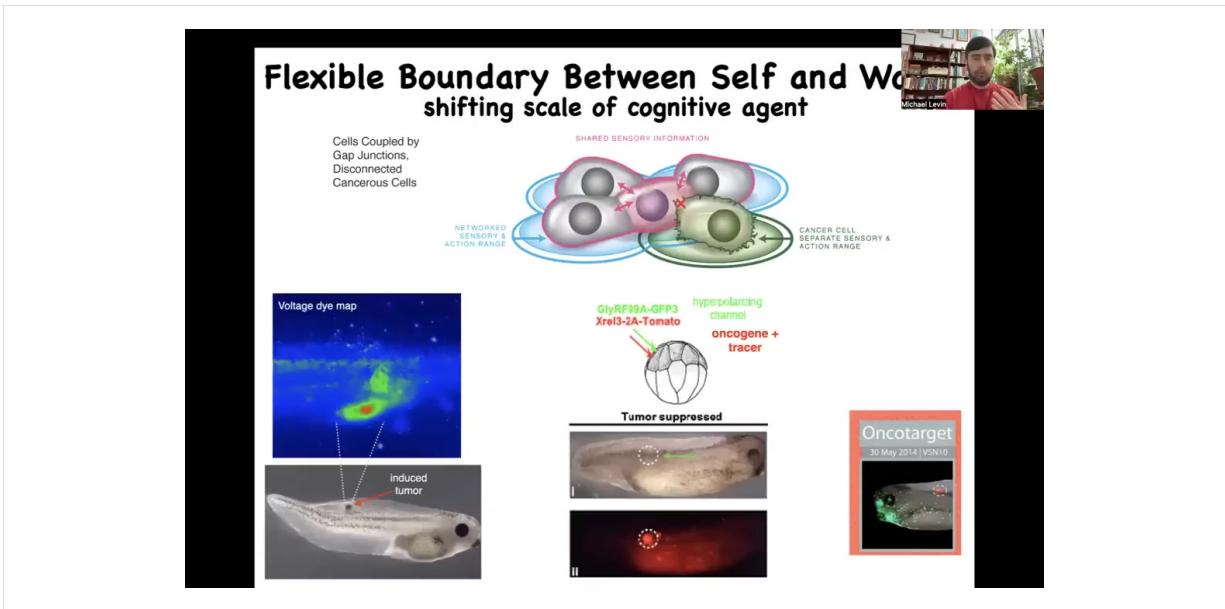
This is something that we're working on, this kind of full stack of modeling, the way neuroscience goes from the proteins that are inside of synapses to psychiatry and psychoanalysis, where you can make very large-scale system-level changes.



I want to very briefly point out that one of the things that comes out of this way of thinking is that during all these processes, not only do the cells pursue goals, but the size of those goals changes.

Originally, both in early development and in evolution, unicellular organisms have little tiny cognitive light cones. All of their goals are very small. This cell really is just only interested in the states pretty much around itself. That's a small goal. But during evolution, these bioelectric networks and other forms of communication allow those goals to grow. The size of the goal becomes huge because these cells are committed to this enormous, enormous thing, this structure in anatomical space. In fact, they're perfectly willing to die for it. Many of these cells will die. They will apoptose during this journey. The collective itself is working very hard to maintain this goal. If you deviate it by cutting it, it'll get right back there and then it'll stop. The size of the goals towards which they work grows radically. But that process has a failure mode and that failure mode is cancer because occasionally those cells get disconnected. What you're looking at here is human glioblastoma. These cells are not any more selfish than these cells. It's just their selves are smaller. What's happening during cancer is a constriction of the cognitive light cone where the size of the self, the border between self and world, goes from this very large thing or even larger in the body to back to here, back to their evolutionary ancient past as unicellular amoebas.

That means that if we really think about the scale of this cognitive light cone, what's the size of the kinds of goal states that cells and tissues can entertain, we now have a new entry point to controlling these things in clinical settings.



This is our example where what we've done, we inject a nasty human oncogene, a KRAS mutation, and that'll make a tumor. But instead of trying to kill those cells, what we do is we co-inject an ion channel that will hyperpolarize them and keep them functionally connected, electrically connected to their network. If you do that, even though the oncogene is blazingly strong up here, you can see it everywhere, this is the same animal, and there isn't any tumor, because it's not the genetics that drives.

It's not that these cells are irrevocably broken by this mutation. They were simply disconnected. If you force them to be reconnected, they're now, instead of pursuing little amoeba scale goals of proliferation and migration to wherever life is good, metastasis, they're continuing to work in a group towards making nice skin, nice spinal cord muscle and all of that. We're trying to move that now to some human medicine.

Making Novel Bodies and Minds: 

How much real-time (software) plasticity – what goals are cells capable of forming, pursuing?

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

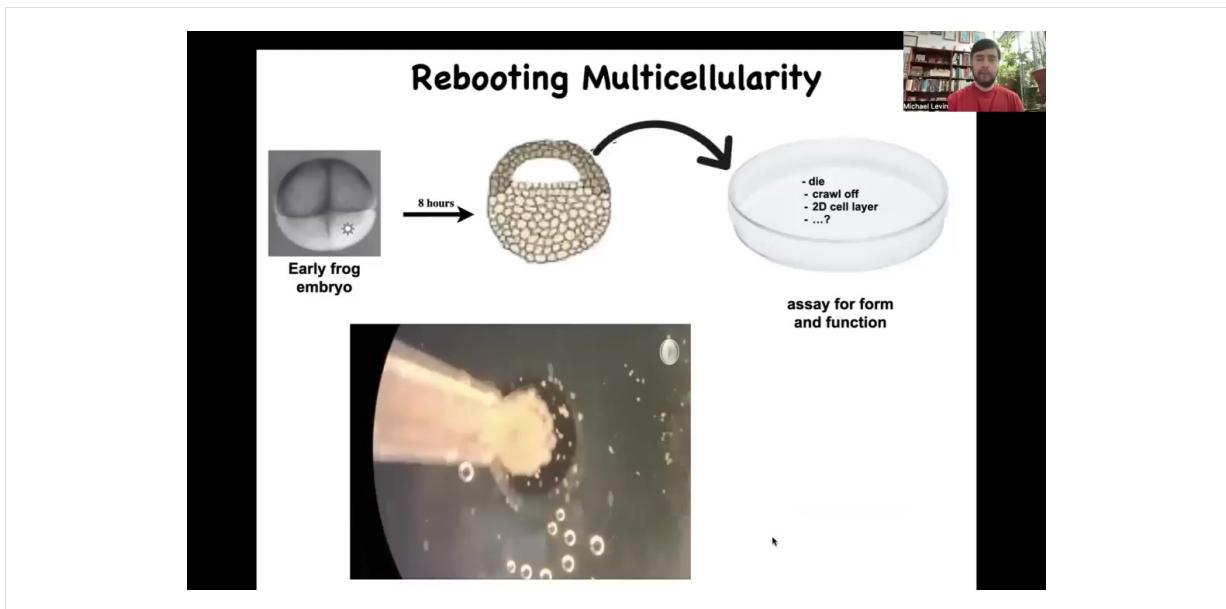
can cells, liberated from their normal boundary condition, reboot multicellularity? Do they like to cooperate, and what will this swarm agent decide to build?

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

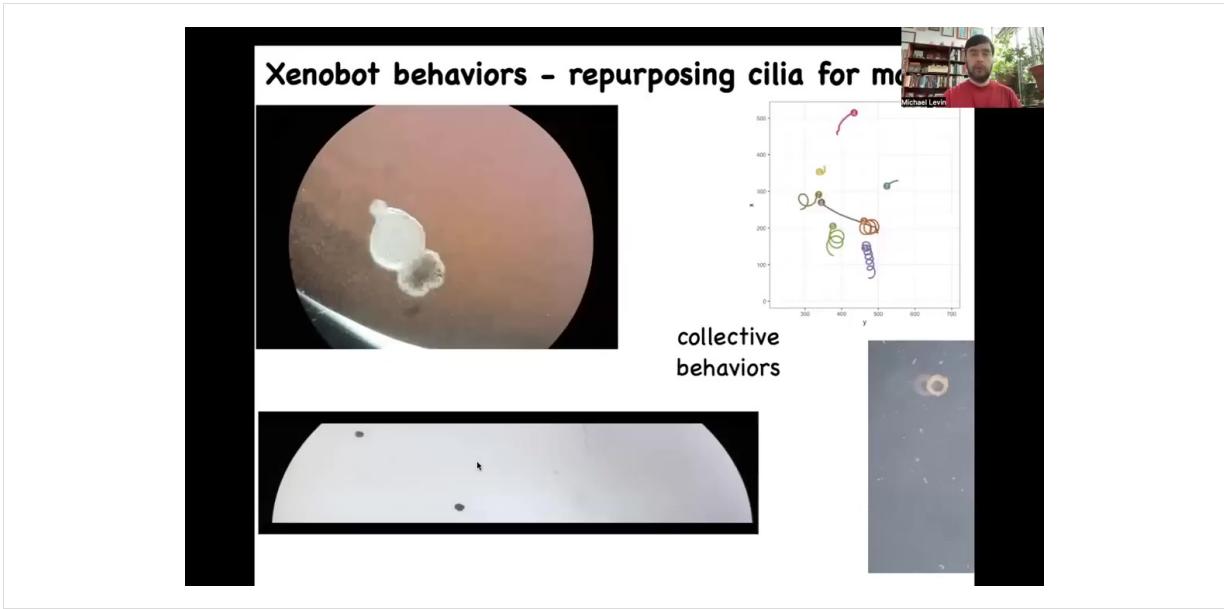
 Douglas Blackiston

The last story I'm going to tell today has to do with a kind of synthetic bioengineering. So far, what I've told you are our various attempts to control normal structures that these cells and tissues build for regenerative medicine approaches. Now let's focus on what plasticity is there. What else are they capable of forming and pursuing?

This is mostly the work that we do in our Institute for Computationally Designed Organisms. Doug Blackiston did all the biology that I'm going to show you, but this is a close collaboration with Josh Bongard at the University of Vermont and Sam Kriegman, who was his student when we did some of this early work.



Let's just ask this question. Take an early frog embryo; we're going to collect some cells up here that are destined to become skin. These are ectodermal cells that are going to become skin. We liberate them from the rest of the body. We dissociate them. We put them in a Petri dish. What's going to happen? Many things could happen. They could die. They could crawl away from each other. They could form a nice two-dimensional monolayer like cell culture. Instead, overnight they come together and coalesce into this interesting little thing we call a Xenobot. Why is it a Xenobot? *Xenopus laevis* is the name of the frog, and we call it a biobot because it is potentially a platform in which we can learn to program their form and function.



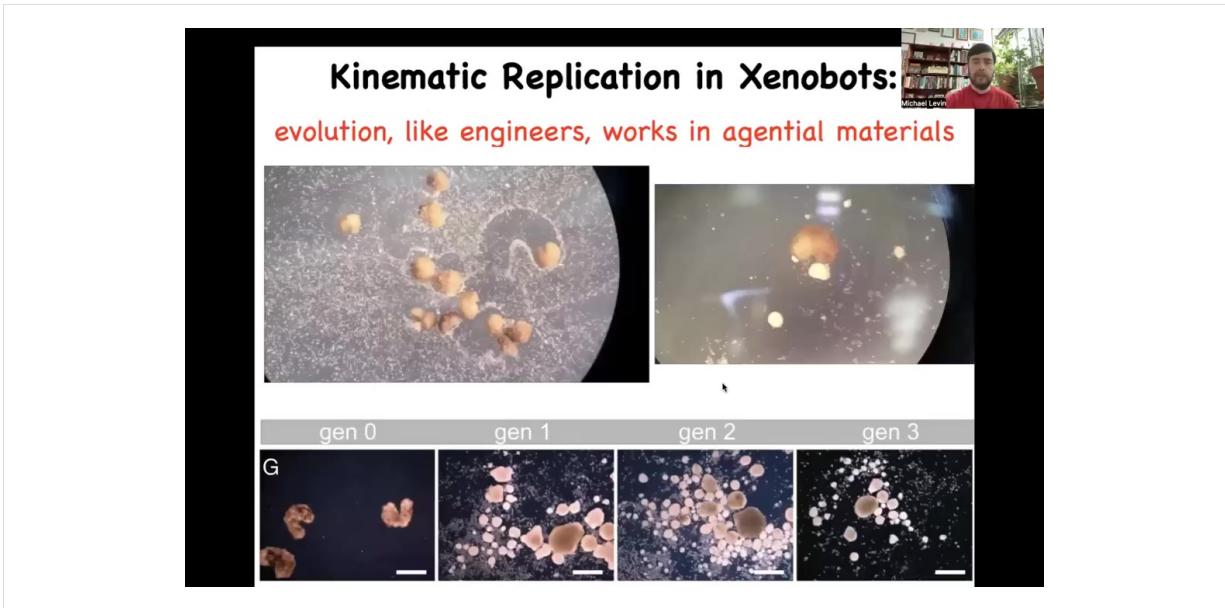
We're going to try to understand how to control this. Here's what it looks like. It's swimming along. It has little cilia, little hairs that propel it against the water. These cilia are normally used by tadpoles and by frogs to move the mucus down the side of their body. But here they're using it to row against the water. They can go in circles and they can patrol back and forth and they can be made into other shapes and they have these collective behaviors. They can interact with each other. They can go on these longer journeys.



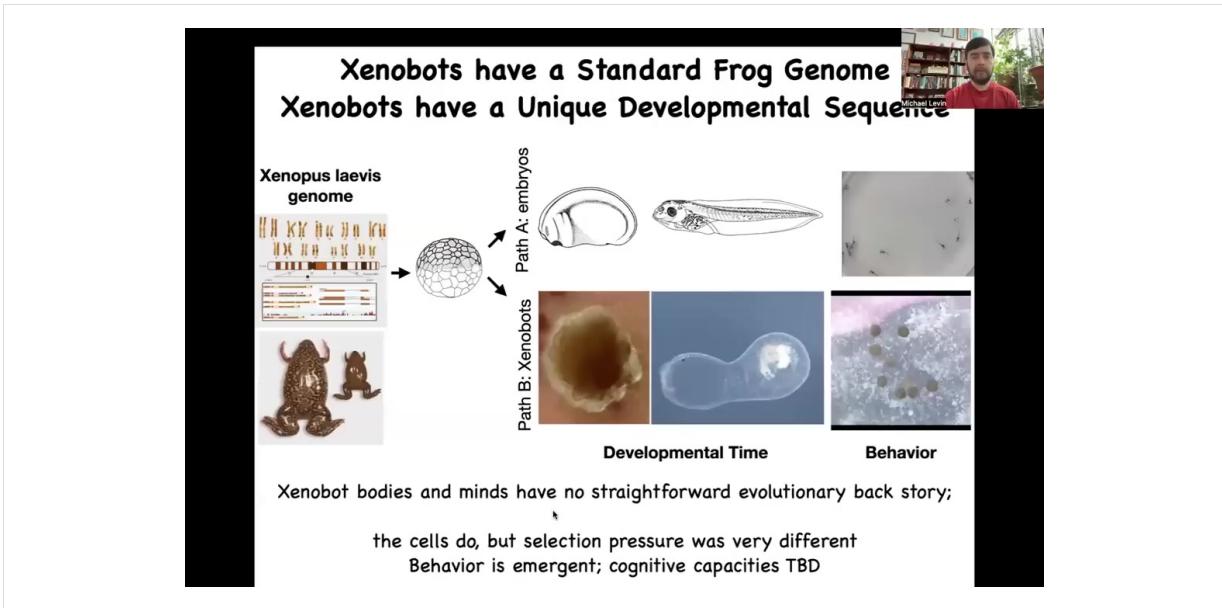
Here's one running a maze. What you see here is that it goes down this path. It takes the corner without bumping into the opposite wall. It takes this corner here, decides to turn around and go back where it came from. We're not actuating them, we're not pulsing them with electricity or anything like that. They swim on their own. They have all kinds of fascinating behaviors.



They can self-heal. If you cut them almost in half, they will form back into their Xenobot shape.



One of the most amazing things they do is they fulfill von Neumann's dream, which is to have a construct that builds copies of itself from material it finds in its environment. If you provide these Xenobots with loose skin cells, what you see them doing is corralling these cells into little balls, polishing these little balls. Because they're working with an agential material, these are not passive pellets; these are cells. What happens is these little balls mature into the next generation of Xenobots. They run around and collect, create little balls themselves, which makes the next generation. They do this kind of kinematic self-replication and continuously make copies of themselves.

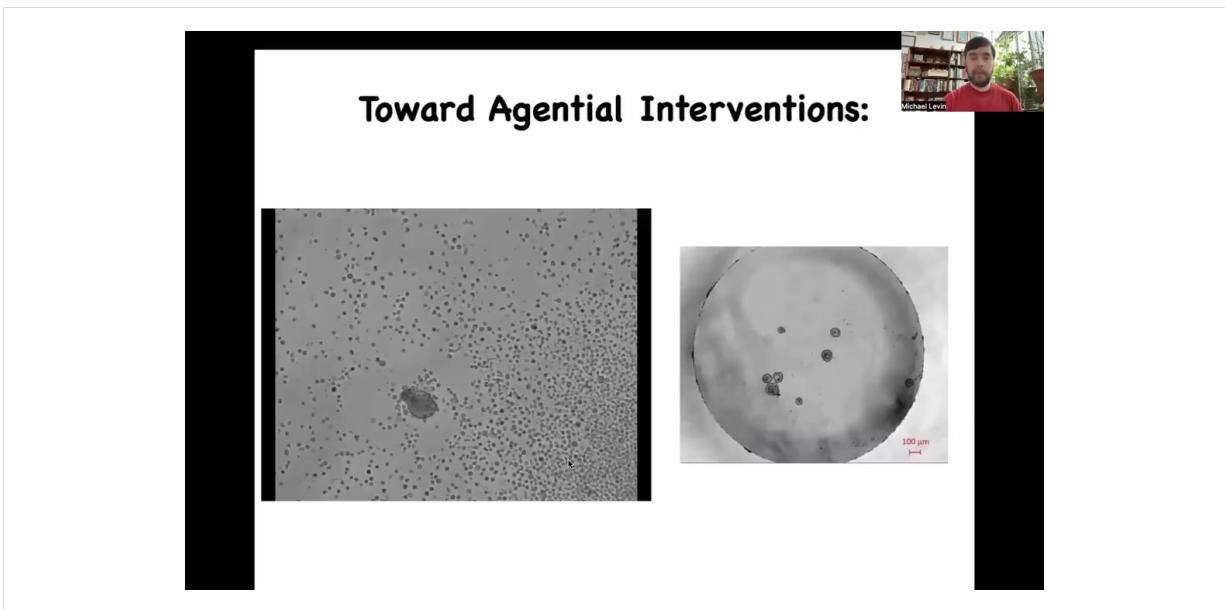


You might think that what the frog genome learned how to do is this developmental sequence and these tadpoles with specific behaviors. But now we find out that by liberating these cells, we didn't add anything to them. We didn't give them new genes. We didn't do any genomic editing. There are no weird nanomaterials. There are no scaffolds. All we did was engineer by subtraction. We liberated them from the normal signals that are coming from these other cells.

What these signals are doing is hacking these cells to force them into a very boring two-dimensional life as the outer skin layer of an embryo sitting there quietly repelling the bacteria. But here you get to find out what these cells are actually capable of. In the absence of these other cues, they do something completely different. They make a Xenobot. These Xenobots have their own developmental sequence over time. This is almost three months old. I have no idea what it's trying to become, but there it is. It's got this weird structure, and they have different behaviors that you're seeing here.

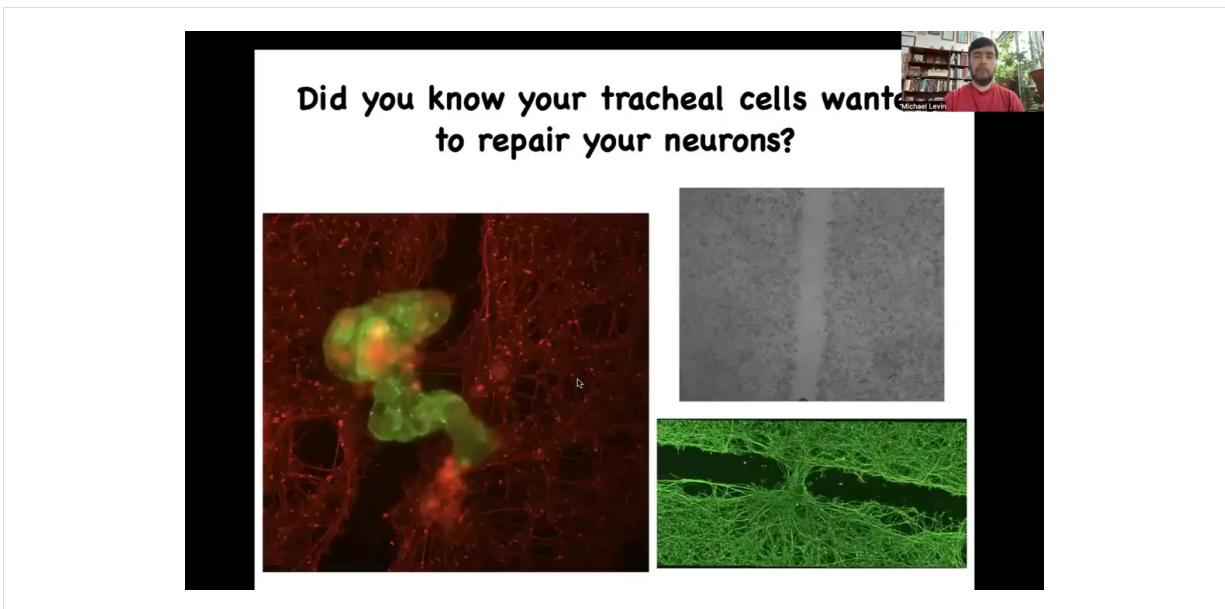
So this branch has no straightforward evolutionary backstory. There's never been any Xenobots. There's never been selection to be a good Xenobot. No other creature that we know of reproduces by kinematic self-replication. So where did this all come from? It wasn't selection. But this is an emergent feature of these cells acting in a new kind of a configuration where they're liberated from the normal forces that are hacking them into their conventional form that we see.

Their behaviors are super interesting. They can learn and have other behaviors that are not published. We're studying their ability to perceive signals and remember them.



The very last thing I'm going to show you is this. Having seen those xenobots, one thing that you might decide is that amphibians are pretty plastic. Embryos are definitely very plastic. We know that animal caps make the ciliated epithelium. Maybe this is just a frog-specific thing.

I want you to take a look at this little creature and ask yourself, what do you think this is? You might guess that this is something I got out of a pond somewhere. It's a primitive organism that came from a waterway somewhere. I can tell you that if we were to sequence the genome, what we would get is *Homo sapiens*. This is 100% human genome. These are biobots, we call them anthrobots, made of adult human tracheal epithelial cells. Patients, many of them in their 70s and 80s, donate tracheal epithelial cells during biopsies, and these cells can be coaxed into these kinds of bots.



These anthrobots — here's one moving down a scratch that we made in a culture of iPS-derived neurons. There's a wound. What you see is that if a bunch of them — we call this a superbot because a bunch of them coalesce together to make this structure — over four days they start to heal this wound. They start to knit together the two sides of this damage.

Now, who would have thought that these tracheal cells, which sit there quietly in your airway for decades, if given the opportunity can reboot their multicellularity, become a self-motile little creature that actually has this capacity? This is the first thing we found. They probably have hundreds or thousands of other behaviors that are interesting.

So this idea here is that all of these cells and tissues have numerous competencies, both individually and in groups.

Next Generation Bio-robotics: communication with collective intelligence

The slide features a diagram illustrating engineering tools across different levels of complexity:

- Hardware modification only:** Shows a clock with gears.
- Modify the data encoding setpoint of goal-driven process:** Shows a smartphone displaying '57.2°F' with a water drop icon.
- Training by rewards/ punishments:** Shows a dog in a water bowl with a bell and a UCS/CS icon.
- Communicate cogent reasons:** Shows a group of people in a meeting.

Synbio is hard if we insist on micromanaging at lowest level

Novel Problem-solving: A 3D bar chart showing the spectrum of persuadability:

Decision-making	Acting	Motility
Learning	Adaptation	Reporting
Discrimination	Generalization	Energy generation
Amplification		Self-repair
Sensing		

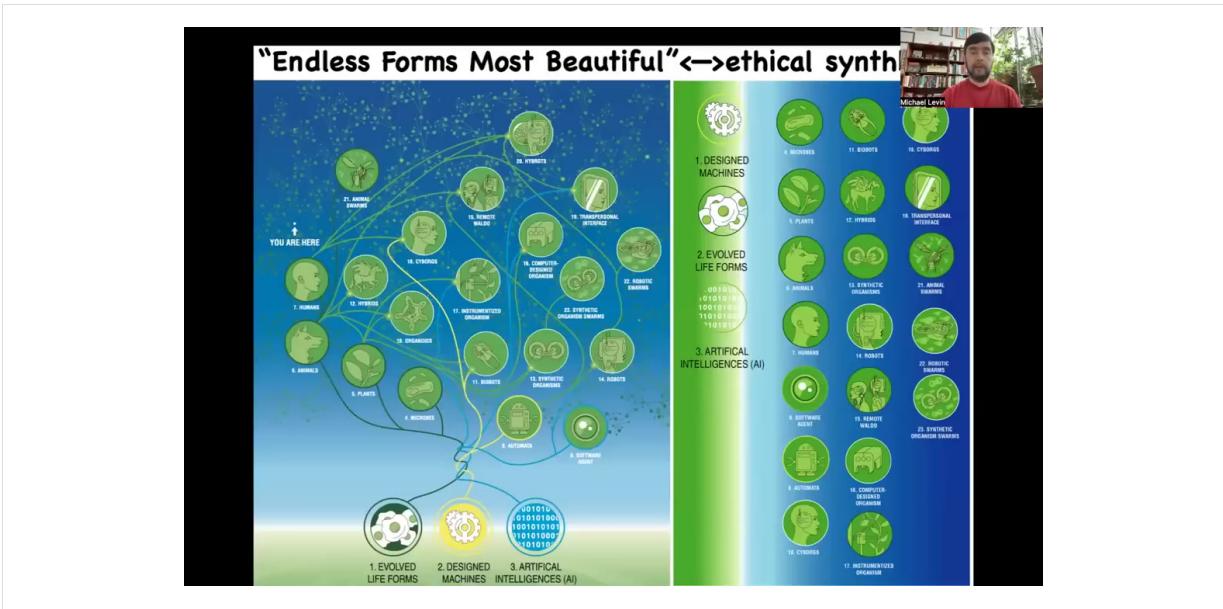
AI-discovered bio-prompts

traditional synbio and bioengineering

Michael Levin

Our job is to start to understand them and start to program them because with any complex system, the goal isn't to control them necessarily bottom up. You can see here, these are the different engineering tools that we have for different levels of complexity. I call this the spectrum of persuadability because you use different tools to manipulate these different kinds of systems. Synbio and morphogenic engineering are really hard if you insist on micromanaging at the lowest level. We're dealing with a material that has all of this stuff built in. It's already here for you. We can add some of these things with synthetic biology, but enormous power arises from controlling these kinds of capacities up here.

For example, those anthrobots are patient-derived agential interventions that could be injected back into the body. You don't need immune system suppression because they're the same cells as the patient. In fact, they have the same priors as we do about what health and disease are, what inflammation is, what cancer is. We don't have to teach them this. We don't have to construct all the sensors that they need and all this other machinery.



It's already there. Because of this amazing ability of biology to assume from the start that the parts are unreliable, that you cannot count on being a proper embryo. Who knows what configuration you're going to come into the world as. Biology is incredibly interoperable. Any combination of evolved material, designed or engineered material, and software is some kind of agent. So cyborgs, hybrids, chimeras of different kinds — some of them are already being made, but many of them are coming. In the future, what we're going to see is that all of Darwin's "endless forms most beautiful" — the whole variety of life on Earth — is a tiny corner of the possible state space that we're going to explore. For the young people in the audience, in your lifetime you are going to be living in a world where you are not going to be able to judge other beings based on what they look like or how they got here, meaning evolved versus engineered, because we're going to be seeing every combination of biology and technology moved in every direction. All of these are viable bodies and minds. We're going to have to develop strategies for an ethical synth biosis with these beings and understand how we can relate to other beings when they are not on the tree of life with you. Questions about their structure or provenance are not the key parameters that you need in order to understand how to relate to them.

Summary:



• Anatomical compiler - a communications device with cellular collective intelligence

• Bioelectricity - the reprogramming interface of anatomy (we can see, and rewrite pattern memories = top-down, rational control)

• Immense plasticity - exploit native competencies, extract knowledge of cells → bidirectional communication with tissues via AI interfaces

• Unlock CRISPR, synbio, and biorobotics by programming agential materials top-down

I'm going to summarize by talking about this idea that we've learned many lessons from the cellular collective intelligence navigating in anatomical space.

Bioelectricity is a cool interface by which we can communicate with the primitive cognition or the problem solving of those kinds of agents. We can exert, top-down, some degree of rational control over where they go in space. There's immense plasticity. We are working on some AI interfaces to help communicate with those tissues.

We're going to unlock the potential of CRISPR, synthetic biology, and biorobotics by understanding the intelligent, gentle nature of the material, not just the technology to control specific molecules, but to understand the world from the perspective of the material. What does it remember? What does it measure? What are its goals? What competencies does it have? What is its stress level about various things?

If anybody's interested in this stuff, here are some papers that go into all of this in great detail.

Thank you to:

Post-docs and staff scientists in the Levin lab:
Douglas Blackiston - brain-body interface plasticity, synthetic living bodies
Devon Davidian, Evgenia Shmeleva, Nirosha Mungar, Celia Herrera-Rincon, and AiSun Tseng - limb regeneration
Vaibhav Pai - voltage gradients in eye/brain induction and repair
Johanna Bischof - electrotaxis
Tal Shomrat - persistence of memory in regenerating brains
Nestor Oviedo, Wendy Beane, Johanna Bischof, Junji Morokuma - bioelectronics of planarian regeneration
Patrick McMillen - bioelectric imaging and transitions to multicellularity

Ph.D. Students:
Fallon Durant - V_{mem} and pattern memory in planarian regeneration
Gizem Gümüşkaya - Antibiotics

Undergraduate Students:
Maya Ernstrom-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
Erin Switzer - vertebrate animal husbandry
Emma Lederer - Xenobots behavior
Joan Lerner - molecular biology

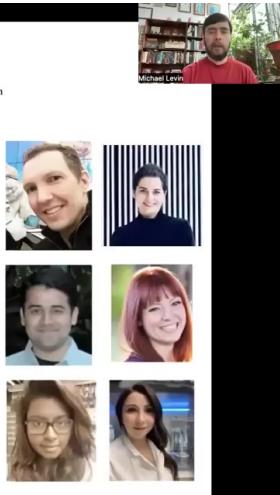
Collaborators: Allen Institute
Dallen Adams - bioelectric face prepatterning
Alexis Pietak - computational modeling of bioelectronics
Joshua Bongard - Xenobots simulations and AI
David Kaplan - V_{mem} and human MSC differentiation, regenerative sleeves
Simon Marlow - physics of the mechanics of Antibot form and function
Chris Fields - physics of sentience and sentience of physics
Richard Watson - computational models of cognitive scaling and evolutionary learning
Giovanni Pezzullo - cognitive science applied to morphogenesis
Don Ingber - cell biophysics and biocomputing
Jack Tuszyński - cell biophysics and drug discovery

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

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

Illustrations: Jeremy Guay @ Peregrine Creative

Disclosures: Morphochemicals, Fauna Systems, Astonishing Labs



I want to thank the people. These are the postdocs and grad students who did most of the work that I showed you today. We have lots of amazing collaborators in our technical support, lots of funders to thank for supporting our work, including these companies. We have commercially sourced funding from these labs. In particular, the animals really are the most important component here. I thank you for listening.

Thank you for reading.

More lectures

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

Follow my work

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

Want one for your lecture?

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