

A LECTURE COMPANION

---

# **“The Unreasonable Effectiveness of the Behavioral Sciences in Developmental Biology and Biomedicine”**

---

Michael Levin

Recorded on December 11, 2025

## About this document

This document is a companion to the recorded lecture “*The Unreasonable Effectiveness of the Behavioral Sciences in Developmental Biology and Biomedicine*”, recorded on December 11, 2025. You can watch the original lecture or listen in your favorite podcast feeds — all links are on the page [here](#).

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

---

## Lecture description

This is a ~46 minute talk by Michael Levin titled “The Unreasonable Effectiveness of the Behavioral Sciences in Developmental Biology and Biomedicine” given at a developmental biology/bioengineering symposium, explaining the connections between the cognitive/behavioral sciences and the study of morphogenesis, from the lens of our framework and new technology.

---

## Follow my work

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

---

## Transcript note

This transcript is generated automatically. While we strive for accuracy, occasional errors may occur. Please refer to the video for exact wording.

---

## Want one for your lecture?

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

## “The Unreasonable Effectiveness of the Behavioral Sciences in Developmental Biology and Biomedicine”

COMMUNICATIONS OF PURE AND APPLIED MATHEMATICS, VOL. XIII, 001-14 (1960)

The Unreasonable Effectiveness of Mathematics in the Natural Sciences

Richard Courant Lecture in Mathematical Sciences delivered at New York University, May 11, 1988

EUGENE P. WIGNER

Princeton University

Michael Levin  
Allen Discovery Center at Tufts

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

<http://thoughtforms.life/>

ALLEN DISCOVERY CENTER at Tufts University

Computer-designed Organisms TUFTS UNIVERSITY | UNIVERSITY OF VERMONT

WYSS INSTITUTE

HARVARD

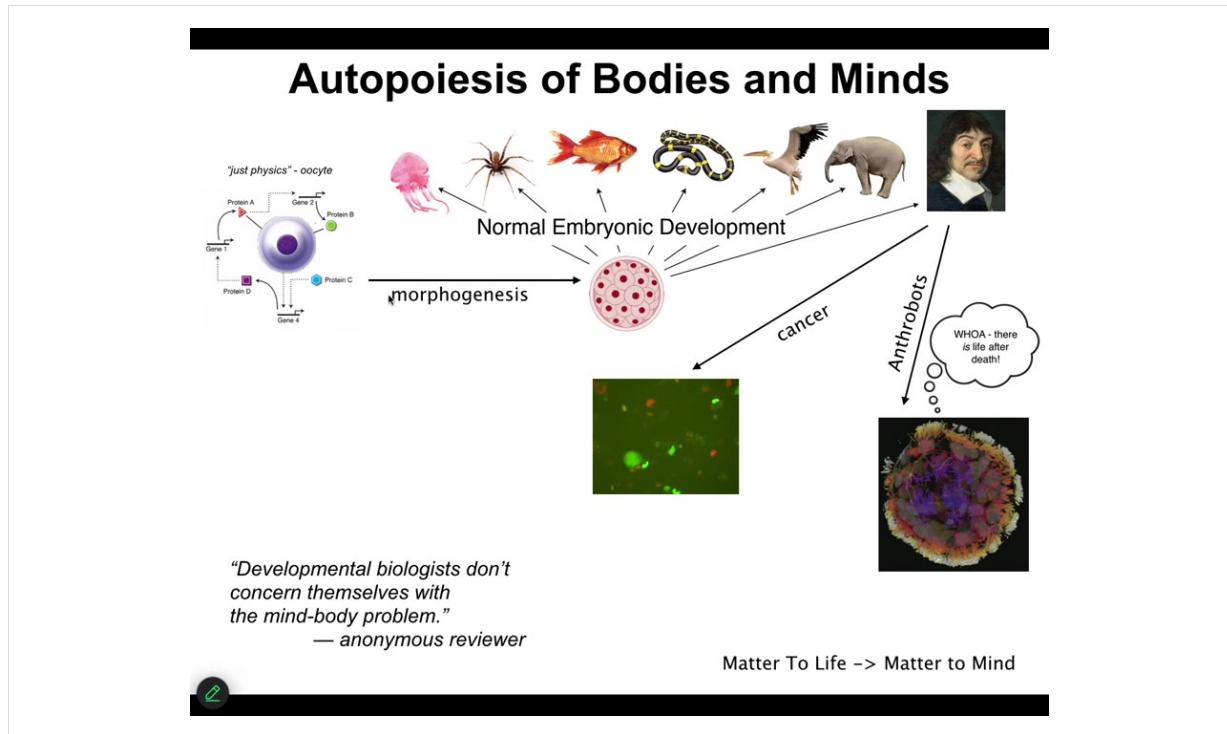
Thank you for the opportunity to speak to you all. A fascinating title for your session.

What I'm going to do today is give a talk that I call "The Unreasonable Effectiveness of the Behavioral Sciences in Developmental Biology, Bioengineering, and Biomedicine." That title is patterned on this very famous paper by Eugene Wigner on mathematics in the natural sciences.

I think it's interesting to ask why such a claim. I'm going to make the claim today that tools from behavioral science are of prime importance to developmental biology, to bioengineering, to cell biology. That's a very odd claim. People generally don't think that's true. The reason it seems unreasonable is that we tend to be educated with this ladder of sciences that everybody knows. You've got physics on the bottom, then chemistry, then biology, then behavioral science and psychology, and then whatever comes above that. That's the standard view.

I think that's exactly backwards. Behavioral science goes all the way down, and we'll talk about it. What I'm going to do is take you on a tour of data and approaches that others and we have taken over the last 30 years or so to try to explain why this is the case.

In a review on one of our papers, a reviewer said, "Developmental biologists don't concern themselves with a mind-body problem."

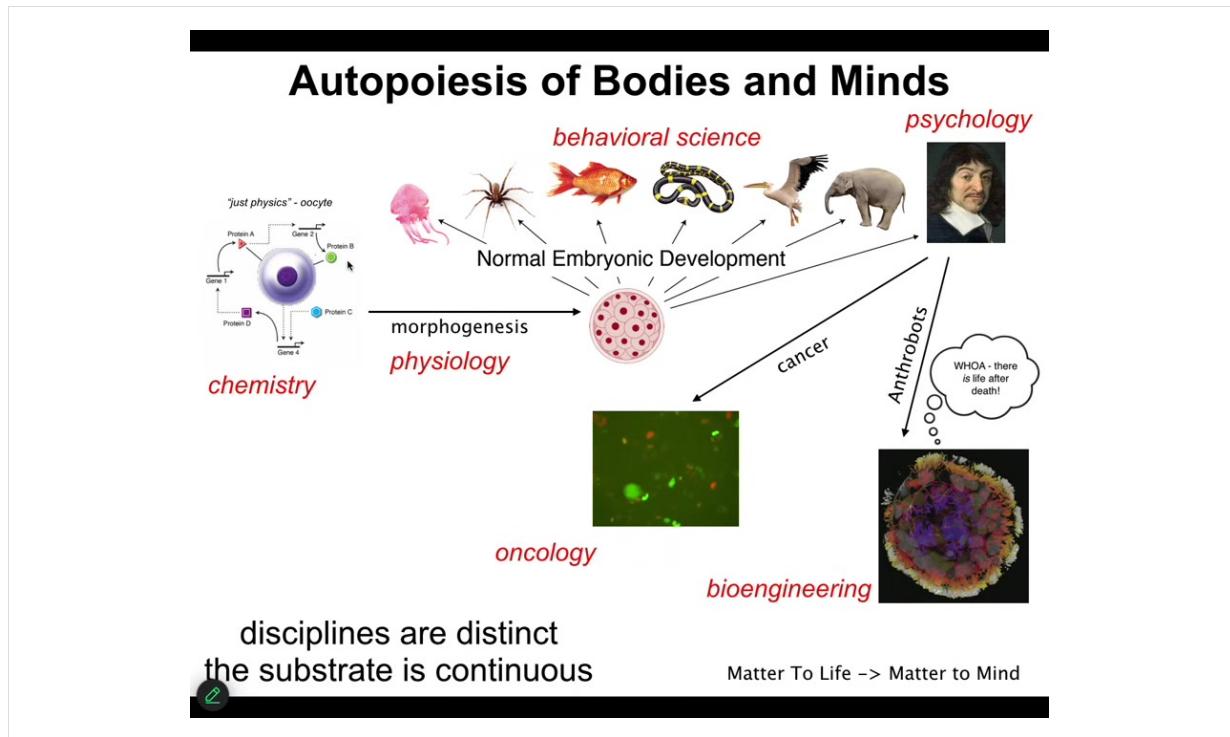


I think that's deeply wrong, even if they don't know that that's what they're doing.

If it's not developmental biologists that are going to crack the mind-body problem, I have no idea who it's going to be. Because in biology, especially in developmental biology, you have this incredible journey that we all take from a single unfertilized oocyte that looks to be described by the laws of physics and chemistry, and then there's this slow, gradual process, and eventually something like this shows up, or even something like this.

Here's René Descartes; he's going to talk about not being a machine and being something else. Even after that, some interesting things happen, which I won't get too much into today, although I will show you some anthrobots.

This is the process of matter, life, and mind. This is exactly where these things come from. Developmental biologists have to concern themselves with these issues. Not only do we take the journey, each of us, from this kind of object to these sorts of things, we also transition across disciplines.



First you are the province of chemistry and then some physiology and then eventually developmental biology and behavior science and maybe psychology and psychoanalysis.

Maybe we take a detour through oncology and parts of us take a detour through bioengineering.

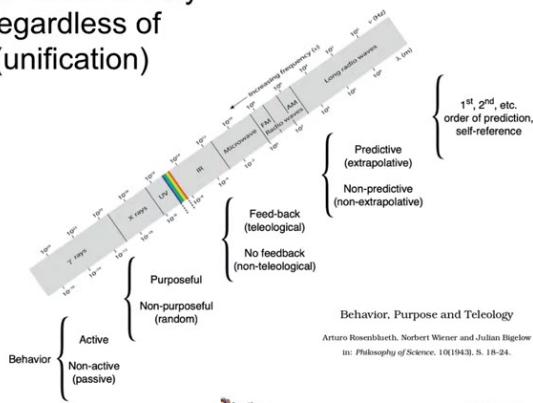
The important thing about this process is that it's smooth and continuous. There is no magic lightning flash at which point you go from being just matter to being a mind.

The disciplines are distinct, and that's just for convenience, so that funders, university departments, and journals can have some sort of boundaries that they try to enforce, but the substrate is continuous. I don't believe it obeys any of the kinds of distinctions that we try to make among these things. It's the transformation across this journey that I think we need to understand.

## My Framework Goal:

- Recognize, create, and relate to truly diverse intelligences regardless of composition or origin (unification)

- familiar creatures - us, apes, birds
- weird creatures (colonial organisms, swarms)
- synthetic biology - engineered new life forms
- AI (software or robotic)
- exo-biological agents (Earth is N=1)
- patterns within media



- moves experimental work forward - new biomedical and synmorpho capabilities, better ethics



My framework is attempting to create a structure that is tools and also conceptual kinds of things where we can learn to recognize, create, and ethically relate to, truly diverse intelligences. And that means familiar creatures like primates and birds, a whale or an octopus, but also some very weird creatures. So colonial organisms, swarms, engineered synthetic new life forms, AIs, whether embodied or purely software, even someday some exobiological agents. And then also some very exotic things that we're not going to talk about today. I just don't have time.

There are two aspects to trying to do something. First, people have tried before. This is Rosenbluth, Wiener, and Bigelow back in 1943, trying to show what the great transitions might be on a scale from passive matter to human-level metacognition. What I always have in mind is this electromagnetic spectrum, this idea that before we had a good theory of electromagnetism, we had magnets and light and lightning and static electricity. We had different names for them. Because of our own evolutionary history, only a tiny bit of that spectrum was accessible to us. But what a theory enabled us to do was have a unification, where in fact we now understand that all of these things are part of the same phenomenon, and it allowed us to make technologies to let us operate in regions to which we are normally not able to reach.

I think something very similar is going on here in science, where we have a tremendous amount of mind blindness. That is, because of our own evolutionary history, we're only good at recognizing intelligence and very narrow types of embodiments. What I'd like to do is to create tools to ameliorate that, in ways that are

not just philosophy or linguistics in terms of redefining intelligence. I want to move experimental work forward. I want this framework to drive new biomedical discoveries, new capabilities in synthetic morphology, and better ethics. If you want to read a lengthy exposition of that, this is a 1.0 version of that.

One of the things that I always think about in trying to put different types of agents on the same scale is what do they all have in common?

Slide 5 of 55 · Watch at [5:57](#)

### Axis of Persuadability:

The diagram illustrates the 'Axis of Persuadability' with four circular icons:

- Hardware modification only:** Shows a clock mechanism.
- Modify the data encoding setpoint of goal-driven process:** Shows a control system interface with a temperature setting from 68°F to 72°F.
- Training by rewards/punishments:** Shows a dog with a bell and a water drop, labeled 'CS' (Conditioned Stimulus), 'R' (Reward), and 'UCS' (Unconditioned Stimulus).
- Communicate cogent reasons:** Shows a classical painting of a philosopher (Aristotle) and a student (Plato).

Below the icons, a photograph of a man working on a complex electronic circuit board is shown, with the following text:

Intelligence is to be decided empirically, not via stale philosophical categories

Cognitive claims are interaction protocol claims

> better technology  
> better ethics

Try to think of it this way...

One way we can start to think about this is through the axis of persuadability.

In other words, as engineers, you look at a system and you ask yourself, what tools can I use to optimally interact with that system? Maybe hardware rewiring is the only thing that I can do, but maybe the tools of cybernetics and control theory will do, because we're talking about some kind of homeostatic system, or maybe training with rewards and punishments, behavioral science, maybe psychoanalysis and friendship and love.

What I think is that any kind of a claim about a system is really a protocol claim. What you're saying as an engineer is, here's this bag of tools that I plan to bring to that system. Let's all see how that works out for me. Then we find out.

That means it's an empirical question as to how much or what kind of intelligence any system has. You can't decide it from a philosophical armchair. You can't try to prop up ancient philosophical categories about what intelligence or cognition means; you have to do experiments, see how they work out, and then we can all see how that went. This is the sort of thing we're made of.

Slide 6 of 55 · Watch at [7:17](#)

## We are Made of Agential Materials



video by Charles Krebs

00:02 - 01:12

nature reviews bioengineering

<https://doi.org/10.1038/s44222-022-00001-9>

Review article

Synthetic morphology with agential materials

Jamie Davies & Michael Levin

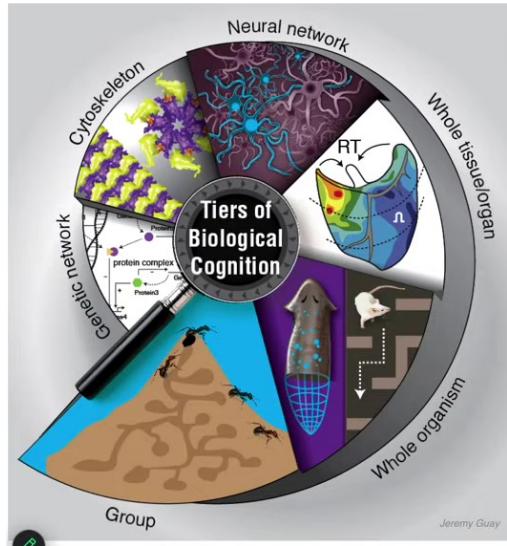
Lacrymaria = 1 cell  
no brain  
no nervous system  
high competency  
at cell-level  
agendas

We are made of an agential material, meaning that unlike wood or metal or these things that just hold their shape, we are composed of, for example, cells.

So this is a single cell called the Lacrymaria. There's no brain, there's no nervous system. You can see it has this incredible ability to get its goals met in the local micro-environment.

If you're into soft robotics, you might be drooling right now because we don't have anything remotely this capable that we can build. We talk about Jamie Davies and how you engineer with these materials, which is very different from engineering with normal matter.

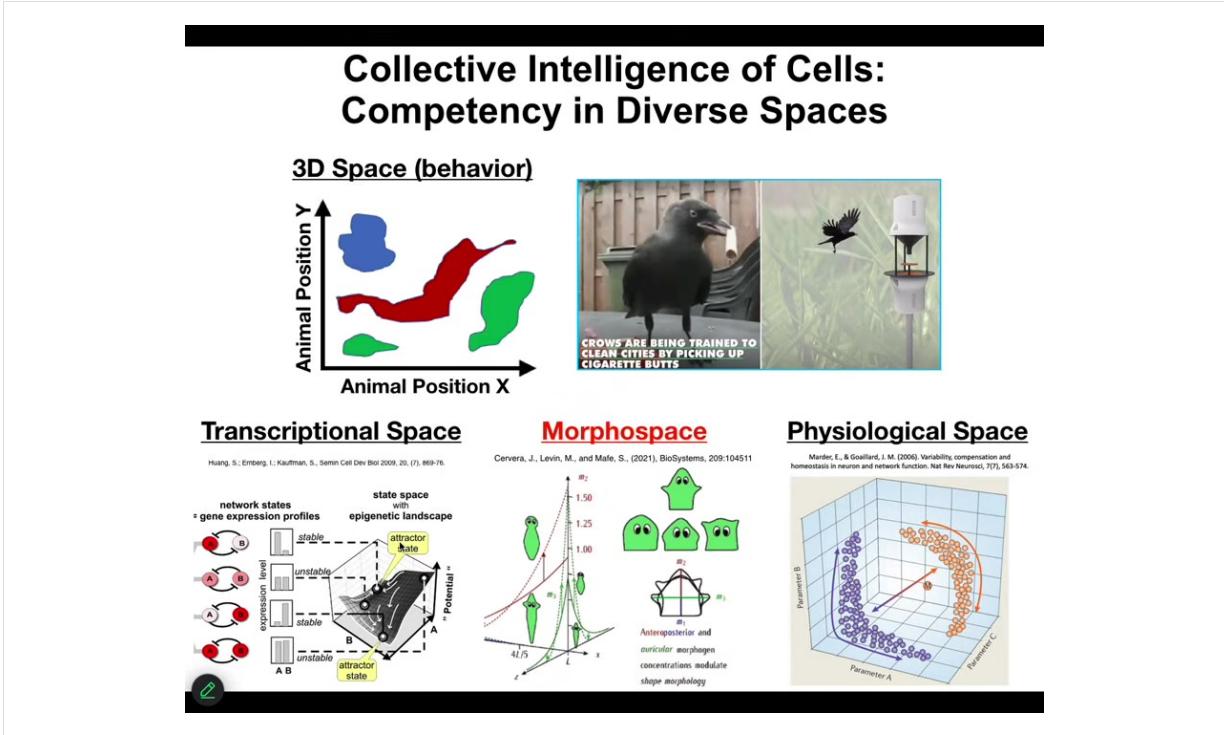
## Nested Competency, not Merely Structure



### 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

And so in biology, we actually have a multi-scale competency architecture because all the way from the molecular networks up through cells and tissues and organs and whole animals and swarms, all of these systems are not just scales of size, but they're actually different, they have different competencies projected into different spaces.



And what I mean by that is simply that people are very used to thinking about intelligence as movement through three-dimensional space. When you see crows picking up cigarettes to receive a reward or something that is a medium-sized object moving at medium speeds through 3D space, we're primed by our evolutionary origins to recognize that.

But biology has been doing that kind of sensing, decision-making, actuation loop in many spaces long before we had muscle, long before we had nerves. Lots of biological materials navigate transcriptional spaces, high-dimensional difficult landscapes, physiological state spaces, and anatomical morphospace. This is the one we'll spend the most time talking about today, which is the space of anatomical possibilities.

So what we're interested in is understanding this navigation. What kind of information processing do these systems do? What do they know about their world? How much memory do they have? All those kinds of things. I'm going to use, for the purposes of today, one definition of intelligence.

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

- William James

*Hypothesis: morphogenesis is a collective intelligence, exerting its behavioral competencies in anatomical morphospace*

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

- Different time scale (minutes, not milliseconds)
- Different problem space (morphospace, not 3D motion)
- but, same (homologous!) mechanisms - ion channels, GJs, NTs
- and, many of the same algorithms (tools of cog sci work great)



There are many. I'm not claiming this is the only definition, and I'm also not claiming that it captures everything. It doesn't capture play and exploration. I'm going to focus on one particular aspect of it, which was emphasized by William James, which is some degree of the ability to reach the same goal by different means. This is nice because it focuses on a measurable thing, that is the ability to reach certain goal states under different circumstances. It doesn't talk about brains. It doesn't talk about a specific problem space. It's a nice general cybernetic definition, very helpful.

And so the middle part of this talk is going to be exploring the data that we've generated with this hypothesis, that morphogenesis, so pattern formation out of cells in vivo or in vitro, is a collective intelligence. That collective intelligence, it's like all of us, it's made of parts. All intelligence is collective intelligence. It exerts its behavioral competencies in an unusual space, an anatomical amorphous space. It's a nice stepping stone towards understanding diverse or unconventional intelligences because it's related to us. It's made of the same stuff, so it's not completely alien. And yet it operates in a different scale with different kinds of goal states that it can undertake in anatomical amorphous space. It significantly stretches our imagination and our ability to generalize from our ancient pre-scientific experience. So I think this is a very valuable exercise.

So what we're going to look at next is some examples of what kind of collective intelligence do cellular swarms deploy. What we're going to see is that it works on a different time scale. Instead of milliseconds in brains, we have to think about minutes

and hours. It's a different morph, different problem space. Not three-dimensional motion, but changes in morphospace.

Interestingly, the same exact mechanisms as your brain uses to do this. Ion channels, gap junctions, neurotransmitters, we'll talk about that momentarily. Many of the same algorithms, the tools of cognitive science work great in this field. Thomas O'Brien developed this amazing tool that you can go to this website called "Field Shift," which automates something that I used to have my students do manually, which is to take a neuroscience paper and just replace words. Every time it says neuron, I make it say cell. Every time it says milliseconds, say minutes. If you do that replacement, generally you have yourself a pretty reasonable developmental biology paper because there is incredibly deep symmetry between the question of where bodies come from and the question of where minds come from.

So let's look.

---

Slide 10 of 55 · Watch at [12:00](#)

## Components of Intelligence:

- Homeostasis: the basis of goal-directed activity
- Homeostasis<sup>2</sup>: Same ends by different means
- Hierarchical, non-local control
- Hackability (software, not just hardware)
- Learning
- Creative problem-solving toward default goals
- Novel beings, novel goals: never give up



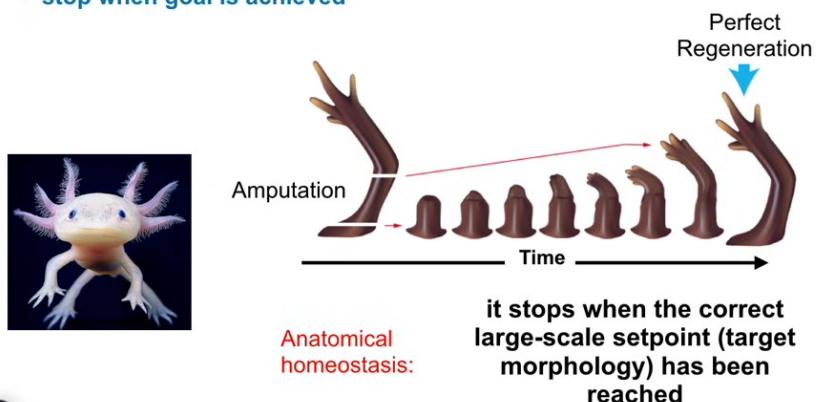
I'm going to take you very briefly through these examples. If you want to see the papers and the underlying data, everything's on our website. We're going to run one

by one through a bunch of well-established components of intelligence. First, goal-directed activity. Here's a simple example.

Slide 11 of 55 · Watch at [12:17](#)

## Same anatomy, from different starting states

- get to the same outcome (maintain set point)
  - despite perturbations
  - from diverse starting positions
  - via different paths
  - **stop when goal is achieved**



Here's an axolotl. We know that it is not simply a matter of feed-forward open-loop emergence, how you get from an egg to an axolotl, because if you amputate the limb, then the cells will very rapidly grow, they will pattern, and then they will stop. The most amazing part is that they know when to stop. When do they stop? They stop when a correct axolotl limb has appeared.

This is a context-sensitive homeostatic system. You can think about it as trying to keep to the same region of anatomical space, this particular region. If you're deviated from it, you'll do your best to try to get back there. This is not simply a progressive unrolling of some kind of cellular automaton. This is very much context-sensitive, and it has a set point. We know it has a set point because we can see it, and if we deviate from that set point, that is the state towards which it will come back.

## Components of Intelligence:

- Homeostasis: the basis of goal-directed activity
- Homeostasis<sup>2</sup>: Same ends by different means
- Hierarchical, non-local control
- Hackability (software, not just hardware)
- Learning
- Creative problem-solving toward default goals
- Novel beings, novel goals: never give up



Another interesting thing about intelligence is not simply that it can reach the same goal every time, but it can do so by different means. One simple example that I can show you is the transition from tadpoles to frogs.

## Same anatomy, via different paths

**Normal**

**Picasso-like**

**normal development**

**"as needed" remodeling**

**Laura Vandenberg**

**Craniofacial mispatterning**

**Metamorphosis**

**Morphometric analysis and modeling**  
reveals: **faces fix themselves!!**

**dorsal** **profile**

**unaligned eyes**

**aligned eyes**

**A** **B**

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.

Cannot just follow a rote set of steps. How does it know when it's "right?"

Tadpoles have to rearrange their face. The eyes, the nostrils, the mouth, all of these structures have to rearrange. You get from a tadpole to a frog. That was thought to be a hardwired process. All tadpoles look the same. All frogs look the same. If you do the same motions in the same direction you should get from here to there.

What we created were so-called Picasso tadpoles. We scrambled the face. The eyes were on top of the head, the mouth was off to the side, everything was completely scrambled. What we found is that basically these creatures make pretty normal frogs because all of these structures don't just go in the same pattern they always go. They move in novel ways, sometimes going too far, having to come back. They have all kinds of rearrangements to get to the correct final state. We're going to talk in a few minutes about how they know what the correct state is.

What the genetics gives you is not a hardwired set of movements. It actually specifies a program that executes a flexible ability to recognize unexpected states and take corrective action toward a specific goal. We're all told in chemistry and cell biology classes that none of these systems know anything. We somehow know things, but cells and tissues definitely don't know anything. They just follow hardwired rules and eventually interesting things happen. That's a claim. That shouldn't be an axiom. That should be an empirically testable claim.

The claim is that open-loop emergent models, like cellular automata, are going to be sufficient to explain and harness morphogenesis. It turns out, as a matter of empirical testing, that claim is incorrect. Those are useful models in some cases, but they

definitely don't give you everything you need. There are many other kinds of tools in the study of goal-seeking systems, cybernetics, psychology, and so on, that are very useful in our field. You should not think that there's some sort of rule that you have to stick to open-loop kinds of models.

The next thing that's really interesting and relevant to what I was just saying has to do with hierarchical non-local control.

---

Slide 14 of 55 · Watch at [15:30](#)

## Components of Intelligence:

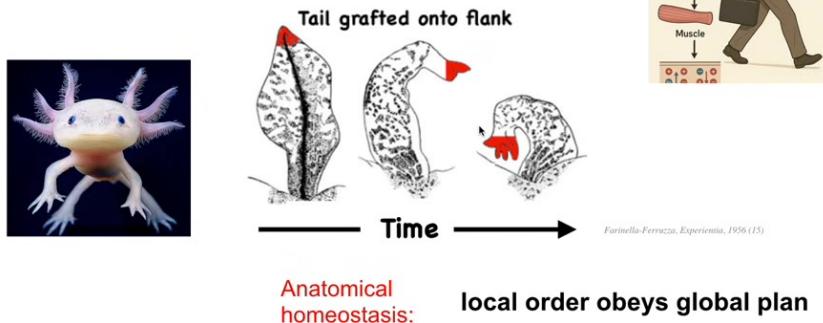
- Homeostasis: the basis of goal-directed activity
- Homeostasis<sup>2</sup>: Same ends by different means
- Hierarchical, non-local control
- Hackability (software, not just hardware)
- Learning
- Creative problem-solving toward default goals
- Novel beings, novel goals: never give up



I want to point out this experiment done first in the 1950s, when you take a tail and you graft it to the side of an axolotl or a newt or some other kind of amphibian.

## It's not Just about Damage: Holistic Order

- get to the same outcome (maintain set point)
  - despite perturbations
  - from diverse starting positions
  - via different paths - modify as needed!
  - stop when goal is achieved
  - parts aligned toward top-down goal



What happens then is that over some period of time, this thing metamorphoses into a limb. In other words, the tail turns into a limb structure. It remodels. Now, pay attention to the cells at the tip of the tail. They eventually turn into fingers. Why? There's nothing wrong with the tail. There is no damage. There's nothing wrong here. There are tail tip cells sitting at the end of the tail. Why would these things suddenly start turning into something else? It's because there is a very important top-down control where the parts are aligned towards a goal that they don't know anything about. Individual cells have no idea what a tail is or what fingers are or anything like that, but the collective absolutely does. And it does because if you deviate the large-scale form by putting a tail where a limb should go, over time the system will correct to a large-scale target morphology that is more correct. In other words, having a limb here instead of a tail. And in doing so, it filters down all of the control signals to change the molecular biology of the cells at the tip of the tail, even though locally there's nothing wrong, to turn them into fingers. So while the individual cells have no idea why they're being reprogrammed and remodeled and killed off and so on, to become these fingers, the collective is doing what collectives always do, which is to deform the option space for their parts so that the parts will then do things that are aligned with goals of the large scale in some other space that the parts don't need to know anything about.

Notice the interesting similarity between voluntary motion. Every morning when you wake up, you might have very abstract research goals, social goals, financial goals, whatever. These are incredibly abstract kinds of things. But in order for you to act on them, ions have to cross your muscle membranes. This is the essence of mind-matter

interaction right here. Your goals and your hopes and dreams and your plans have to eventually be able to make the biochemistry do things in your muscles, right? And your whole body, using the electrical system we're about to talk about, is a set of transducers that allows that to happen. For the exact same reason that the ions in your muscle membranes dance due to high-level abstract goals in weird spaces that you might have, for that same reason, these molecular events inside these cells are obeying a large-scale anatomical plan.

Slide 16 of 55 · Watch at [18:12](#)

## Components of Intelligence:

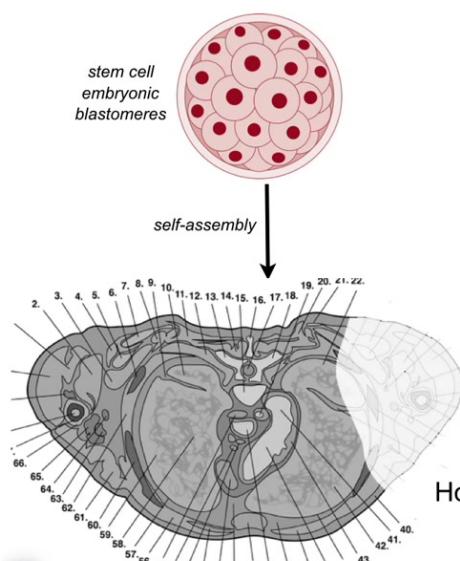
- Homeostasis: the basis of goal-directed activity
- Homeostasis<sup>2</sup>: Same ends by different means
- Hierarchical, non-local control
- Hackability (software, not just hardware)
- Learning
- Creative problem-solving toward default goals
- Novel beings, novel goals: never give up



One other interesting component of intelligence is hackability. That is, intelligent systems generally have a kind of software, not just hardware that they're made of, but they actually have a kind of software that can be exploited to get them to do things they don't normally do.

In order to look at that, this is where we're going to talk about bioelectricity and the interface that my group uses to access the reprogrammability of life.

## Where do Anatomies Come from?



The diagram shows a cross-section of a human torso with various numbered structures. At the top left, a cluster of red circles is labeled "stem cell embryonic blastomeres". An arrow labeled "self-assembly" points from this cluster to the cross-section of the torso below. The numbered structures in the torso cross-section range from 2 to 66, representing different anatomical features.

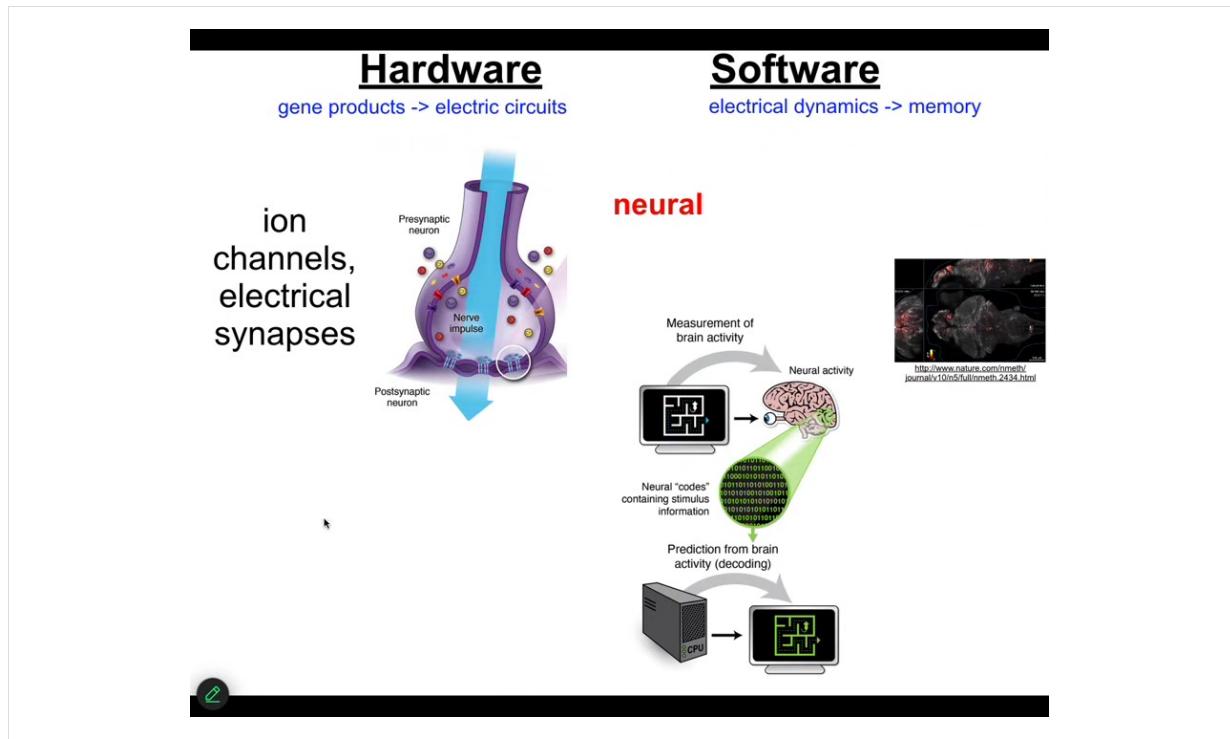


earthlymission.com

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

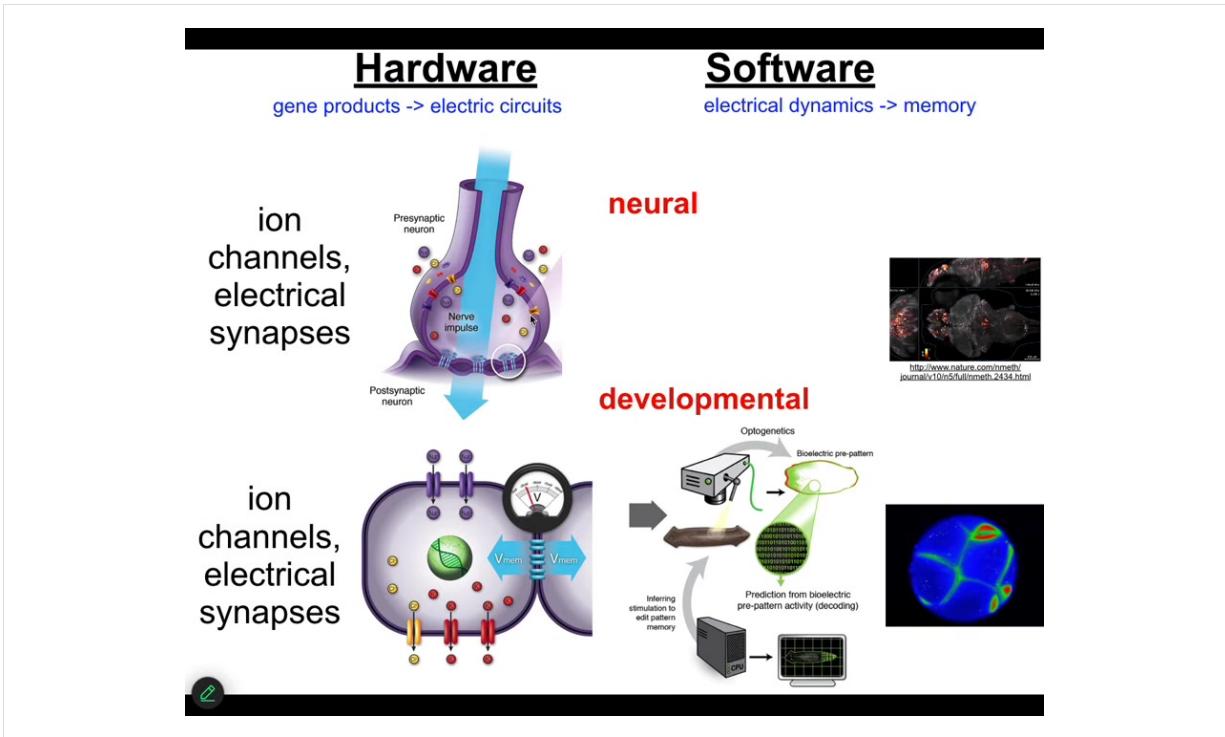
How to repair  
(edit) it?

Let's first ask, where do anatomies come from in the first place? Here's a cross-section through a human torso. Lots of amazing complex structures. Where was this pattern specified? Most people immediately will say the genome—it's in the genes. We can read genomes now. We know that none of this is directly in the genome. What the genome specifies are proteins, the tiny molecular hardware that every cell gets to have, and some timing information around where certain things will be expressed. But this is all the function of the software that's encoded by that molecular hardware. This kind of information isn't in there any more than the structure of termite colonies or the specific structure of spiderwebs is in the genome of those species. We need to understand how this works. What I'm going to tell you is that it works very similarly to another system that we all know, which is the brain.



In the brain, the hardware is made of ion channels. These are proteins that let charged chemicals in and out: calcium, potassium, sodium, chloride. Because of that, it acquires a voltage gradient, and it can pass on that voltage information through electrical synapses onto the neighbors.

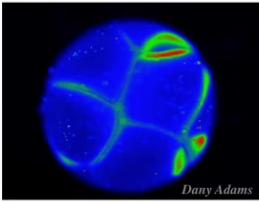
This group did an amazing video of a living zebrafish brain doing whatever it is that fish think about doing. The cognitive content, the memories, the preferences, all of that stuff is encoded in the electrophysiology of this network. That's the commitment of neuroscience. Neural decoding is the idea that if we were able to read and decode these electrical signals, we would know what the creature is thinking about, what its memories are, and so on. This is an amazing system, but it turns out that it's incredibly ancient.



This is not really about nerves and muscles. In fact, I don't think neuroscience is about neurons at all. It's about this cognitive glue, the processes that make us all a collective intelligence. It's why you know things your individual neurons don't know, because it makes a computational network that can operate in other problem spaces. Biology discovered this a very long time ago, around the time of bacterial biofilms. Now every cell in your body has ion channels. Most cells have these electrical synapses to their neighbors. And we have been running that same decoding program, just not in neurons, looking at embryos to ask, what did these electrical networks think about before there was motion such that you could think about positioning yourself in three-dimensional space? To do that, we developed some tools.

## How we Detect and Model Bioelectric Signals:

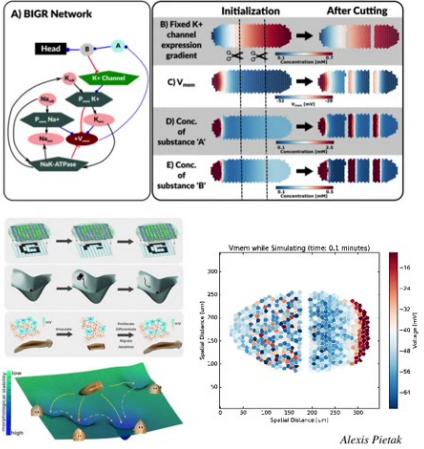
**Characterization** of endogenous voltage gradients - direct measurement and correlation with morphogenetic events



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



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



A) BIGR Network  
B) Fixed K<sup>+</sup> channel expression gradient  
C) V<sub>mem</sub>  
D) Conc. of substance A'  
E) Conc. of substance B'

Initialization      After Cutting

V<sub>mem</sub> while Simulating (time: 0.1 minutes)

Alexis Pietak

Back in 2000, we developed the first molecular tools to read and write electrical information into tissues, into non-neural tissues. We have voltage-sensitive fluorescent dyes, and we can do this *in vitro*. The colors here reveal voltage. These are not models. These are real data. The colors reveal voltage in this early frog embryo. You can see all the conversations that the cells are having with each other.

There are lots of computational tools to try to understand all the way from the molecular networks that express channels up to tissue properties of what happens to these electrical patterns, and then how they do things like pattern completion, AKA regeneration, how they store memories as attractors in these electrical networks. Pattern memories. I'll show you what they actually look like.

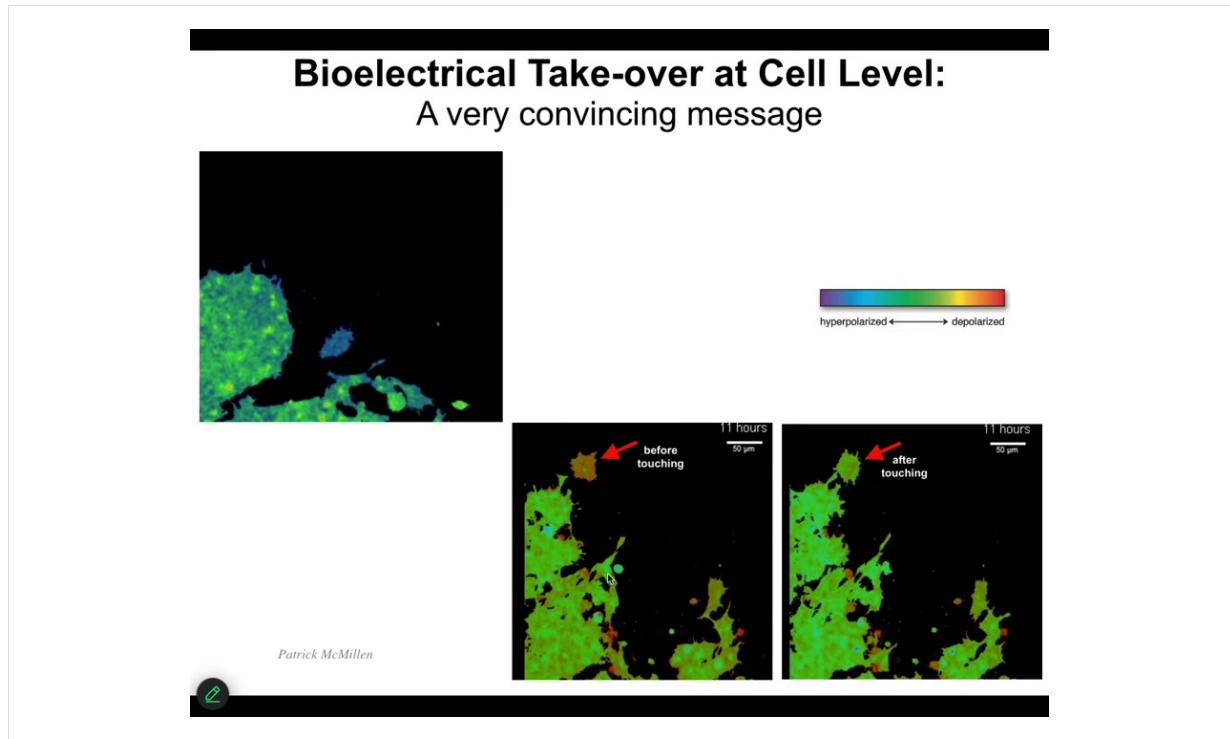
## Endogenous Bioelectric Prepatterns: reading the multi-scale mind of the body



So this is a video of what we call the electric face, defined by Danny Adams in my group years ago. Here's one frame from that video. And you can see there's a bioelectric pre-pattern. Here's where the eye is going to go. Here's the mouth. It's going to be here. The placodes are out there. This is well before the genes come on to pattern the face. If you change this electrical pattern, the pattern of the face changes.

Not only is bioelectricity the memory medium that helps complex structures know what they should look like, it also works to glue individual cells into large-scale structures, but actually across embryos, so multi-scale.

Here's an example. We poke this embryo, and within minutes, all of these other guys find out about it. If we poke this one, these two, even though they're not touching, that information spreads. In fact, we now know that groups of embryos form a kind of hyper embryo that can solve problems better than individuals. In other words, large groups resist teratogens better than small groups. They have unique gene expression.



So now we're going to get to this reprogrammability angle. Having seen this idea that bioelectricity is how cells normally communicate with each other to make all kinds of patterning decisions, we wanted to try to exploit that. I want to show you one simple example. This is a video which I'll play. Here are the before and after frames. Again, the color represents voltage.

You can see this cell, before it touches, has a different voltage than all of these guys. But as soon as it touches, it has the same. Here it is. It's crawling along, minding its own business. This guy is going to reach out, touch it, bang. That's it. Now it's green like the rest of them. The ability to inject an electrical signal through even an extremely tiny contact like this has the ability to change the cell voltage and change the cell behavior. We wanted to be able to do this as well, and so we developed tools now.

## Rewriting Somatic Pattern Memories

The diagram illustrates various tools for manipulating cell membranes:

- Non-neuronal cell group:** Shows a cluster of cells with a color scale from purple (hyperpolarized) to red (depolarized).
- Gap Junctions (electrical synapse):** Shows a junction between two cells with a voltage-gated channel ( $V_{gj}$ ).
- Neurotransmitter (moving via  $V_{mem}$ ):** Shows the chemical structure of a neurotransmitter (dopamine) and a green circle indicating movement through the membrane.
- Ion channels (setting  $V_{mem}$ ):** Shows a membrane with  $V_{mem}$  and channels for  $Na^+$ ,  $K^+$ , and  $Cl^-$ .

**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**

We don't use any electrodes. There are no magnets, no waves, no frequencies. We don't do anything like that. We use all the tools of neuroscience. In other words, we manipulate ion channels and gap junctions. That means neuropharmacology. It means optogenetics. It means being able to mutate these channels and close the gap junctions. So we can control the electrical state of these cells and we can control who talks to whom. I'll show you two examples of what this means.

## Communicating New Goal Top-Down via Bioelectric Interface

The diagram shows a cross-section of a Xenopus laevis tadpole with a red arrow pointing to the gut and a red box labeled 'EYE'. Above the tadpole, a cell is shown with text: 'ion channel mRNA targeted to ventral or posterior regions' and 'can reprogram many regions, even outside "competency zone", into complete ectopic eye!'. Below the tadpole, a list of four points is provided:

1. BIOE is instructive
2. modularity - not cell level, organ-level subroutine call
3. higher-level prompt reveals higher tissue competency than Pax6 prompt
4. self-scaling of system to task

On the right, a grid of 10 panels (labeled i-x) shows immunofluorescence images of Muller glia, Amacrine cells, Rods, and Cones in both Wild Type (WT) and Ectopic conditions. A large image of an ectopic lens cross-section shows 'voltage-modified cells' (blue) and 'wild-type cells recruited to ectopic lens!' (orange). A small image of ants is at the bottom right.

Here's the side view of a *Xenopus laevis* tadpole. Here's the gut. Here's the animal's right eye. Here's the mouth over here. The brain is up here.

If we inject RNA encoding a particular ion channel into cells that are going to become these gut cells here, what we can do is set up a little voltage spot that looks just like that eye spot in that electric face that I showed you. What was done here by Viphof Pai and Sherry Au in my group was to inject the ion channels to cause the voltage spot to be there, and the cells interpret it as a "make an eye" signal. Sure enough, they make an eye, and here it is.

Those eyes have lens, retina, optic nerve, all the same stuff. Notice two interesting things. First of all, it's instructive, but it's very modular. In other words, we didn't have to specify how to build an eye. That's the benefit of dealing with an agential material: you don't micromanage it. You give it cues or prompts. We gave a very simple voltage signal that said, "make an eye here." All the downstream stuff, all the regulation of stem cells and morphogenesis, everything was taken care of after that.

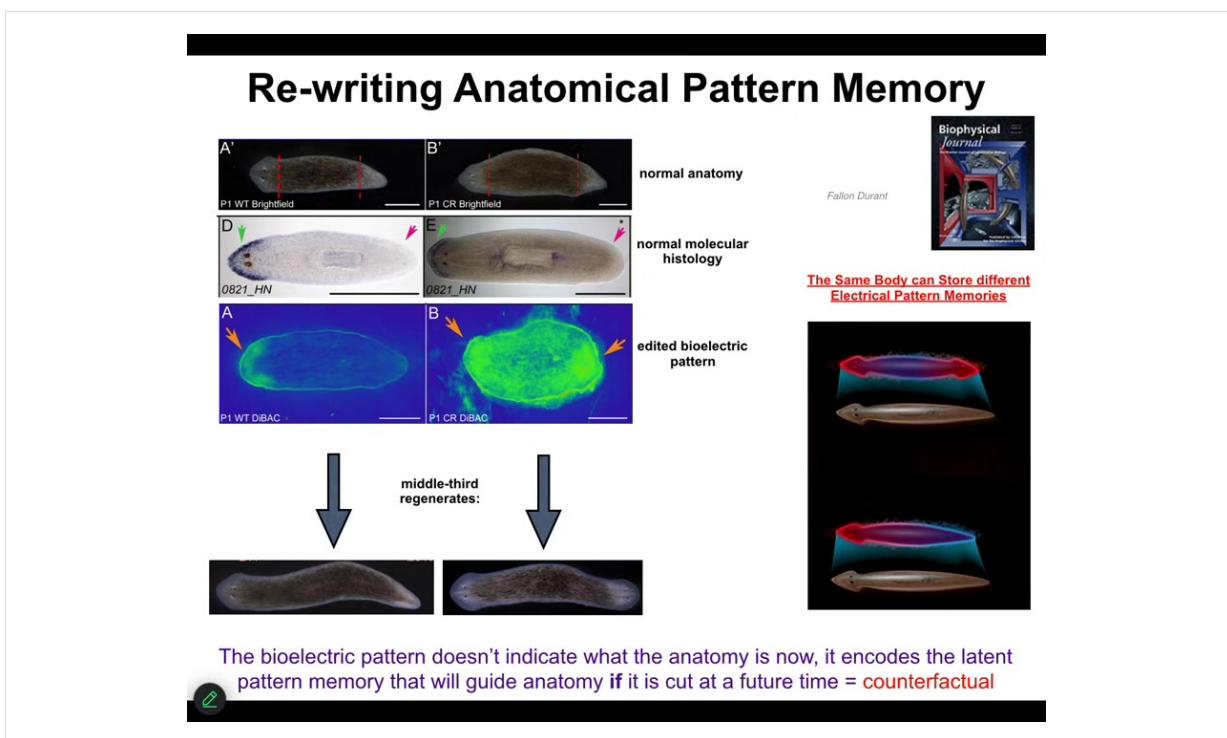
It's exactly what happens in cognitive systems. When I'm talking to you now, I don't need to worry about putting your synaptic proteins in particular places. You'll do that yourself. I'm just giving you text prompts. That's the joy of communicating with a multi-scale agential material.

I'll also note that the blue cells, this is a cross-section through an ectopic lens, the blue cells are the ones we injected, but there's tons of others that are contributing to this

morphogenetic process. We didn't have to tell them to do that. We didn't have to warn them that there's not enough injected cells to actually make an eye. They will do that themselves. They recognize that there's not enough. They recruit their neighbors. There's a bunch of back and forth because the neighbors resist, because this is also how tumors can form. The neighbors don't immediately snap into line. But when things go well, and the signal is sufficiently convincing, which is what we want as workers in regenerative medicine, we want to be able to motivate these cells. They will do all kinds of other things like recruit their neighbors and undergo morphogenesis and so on.

Prompt, not micromanagement.

Slide 25 of 55 · Watch at [27:12](#)



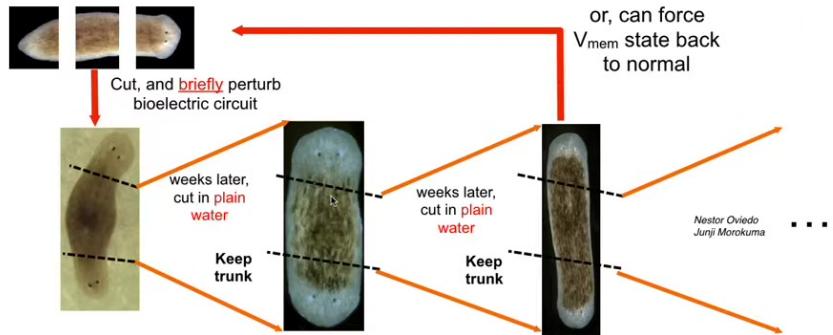
One other example I want to show you is planaria. So this is regeneration. Here's what these worms look like, one head, one tail. And when you cut them into pieces, they reliably, every piece, so let's look at this middle piece, it reliably makes one head, one tail. So you might wonder, How do they know how many heads they're supposed to have? And so we looked and it turns out that there's this voltage gradient that says one head, one tail.

And what we can do is take this animal and reprogram the bioelectronics. And this is again done with drugs, with ionophores and ion channel drugs. And we can establish a different state. So 2 heads. You can see it's messy. The technology is quite messy still. But what will happen then is that the anatomy is still quite normal. The molecular biology is still quite normal. In other words, the anterior markers are expressed up here, not in the tail. The molecular biology doesn't know anything's happened. The anatomy doesn't know anything that's happened. It's a latent pattern memory.

But if you then cut this animal, this is what it does. It will go and make two-headed worms. Then this memory becomes active and it will do what the memory says. So one normal planarian body is able to store two different representations of what a correct planarian should look like. And that representation is a counterfactual memory. It's not true right now, because I only have one head right now, but it is what I'm going to do in the future if I get injured.

This is the pattern memory. If you want to know how all these systems know what they should build, they have a bioelectrical encoding that they maintain of what the correct pattern is, and then they will simply try to reduce error towards that encoding, but it's rewritable. And the reason I keep calling it a memory is because it is stable, but it's rewritable.

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



Nestor Oviedo  
Junji Morokuma

### Basic properties of memory

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



If I take these two-headed worms and I cut off this ectopic secondary head, it doesn't go back to normal, even though the genetics are completely normal. We haven't touched the genome. It will continue to make two-headed worms because that's where the memory is. It is not genetically specified.

What the genetics does specify is a hardware that by default acquires a voltage memory that says one-head, but it's totally changeable. We can change it back. This question of where the fact that they should have two heads is stored is subtle. The genetics matters because you need to have hardware that's capable of acquiring those states. But after that, it is actually stored in the bioelectrical state and you don't need to change the genome any more than we need to change our genomes when we learn new things in the brain. Not only can you reset this animal to have two heads, you can also ask it to have heads of other species.

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

Species-specific shapes = attractors in morphospace = memories as attractors in neural network

*D. dorotocephala*

cut off head, perturb network topology

like: *P. felina?*

like: *S. Mediterranea*

quantitative morphometrics

amputation, free energy landscape, altered cell network dynamics, regeneration, remodelling, Alexis Pietak

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

We can take this triangular form and ask it to make flat heads like a *P. falina* or round heads like an *S. mediterranea* by perturbing the bioelectrical pattern memory that encodes which of these attractors in anatomical space they normally go to. These other species live in these attractors, but the hardware is perfectly capable of visiting them if you push it in that direction, and the distribution of stem cells and the shape of the brain become just like these other species.

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

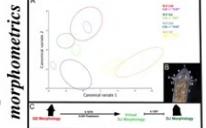
Species-specific shapes = attractors in morphospace = memories as attractors in neural network

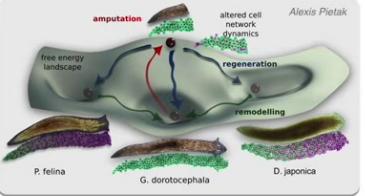
*D. dorotocephala* 

cut off head, perturb network topology

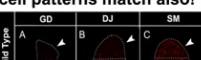
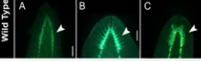
 like:  like: 

**quantitative morphometrics**





brain shape and stem cell patterns match also!

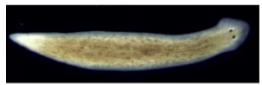
	GD	DJ	SM
Wild Type			
Pseudo			

There is something between 100 and 150 million years of evolutionary distance between this guy and these actual things, but no problem traversing the amorphous space.

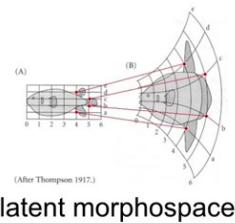
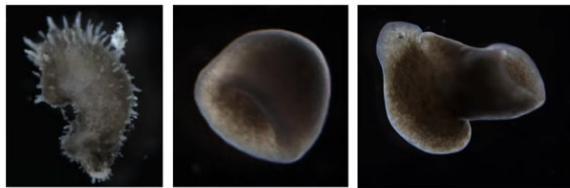
In fact, you can ask them to do more exotic things, things where you can go further in that latent space that aren't even flat or planar at all.

## The Hardware is Extremely Reprogrammable

Normal worm



Bioelectric Pattern Altered



You can make these spiky things, you can make round cylindrical things, you can make hybrids. And plants do this too.

## Good Old Reliable Development like “instinct” - fixed patterns?



White Oak Leaf - Photo by Chris Evans, River to River CWMA, Bugwood.org



This is the only thing I'll say about plants in this talk, but they absolutely have this kind of plasticity.

And you can see that the reliability of development is obscuring a lot of important things, because we see these acorns and we know that they very reliably make these oak trees, and we think that is what the genome knows how to do. The oak genome knows how to make these oak trees, these oak leaves with a very specific shape. And the reliability of it hides the fact that if prompted by a good bioengineer, and these wasps are amazing at putting down prompts that cause the plant cells.

## Messages from Hacker Reveal Flexibility:



Photo Credit: Andrew Deans

### Hedgehog Gall

*Acraspis erinacei*

August - November



© TIMOTHY BOOMER  
WILDMACRO.COM



Parasite hacks host to induce new anatomy (bio-prompting)



They don't build these themselves. The wasp isn't building this. These are plant cells. They make these galls with these incredible structures. We would have absolutely no idea that these cells are even capable of doing this. The plasticity and the other aspects of the morphospace are revealed when you start probing it. And so bioelectricity is how we do that.

## Components of Intelligence:

- Homeostasis: the basis of goal-directed activity
- Homeostasis<sup>2</sup>: Same ends by different means
- Hierarchical, non-local control
- Hackability (software, not just hardware)
- Learning
- Creative problem-solving toward default goals
- Novel beings, novel goals: never give up



The last couple of things I want to show you: learning. You might be surprised to know that the molecular networks inside of cells are capable of six different kinds of learning.

## Collective Intelligence Below the Cell Level

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

Biomedicine:  
- drug conditioning

communications biology Article [Associative conditioning in gene regulatory network models increases integrative causal emergence](#) [View Article Online](#)

Association, associative conditioning, habituation, sensitization, just from the network itself. You don't need special proteins for this. The properties of all kinds of networks can do this. We're taking advantage of this for drug conditioning and other applications. Large-scale memory.

## Trophic Memory in Deer Antlers:

The diagram illustrates the process of antler regeneration over time. It shows three stages: Year 0, Years 1-5, and Years 6+. In Year 0, an 'injury' is indicated by a red arrow on a single-tined antler. In Years 1-5, the antler has regenerated with an 'ectopic' tine (indicated by a green arrow) at the site of the previous injury. In Years 6+, the antler has further regenerated, and another 'ectopic' tine is shown. A question 'What molecular pathway model would you draw?' is posed next to the Years 1-5 stage. A photograph of George Bubenik working with deer antlers is shown on the right. A small circular icon with a green arrow is at the bottom left.

Year 0

Years 1-5

Years 6+

injury

ectopic

What molecular pathway model would you draw?

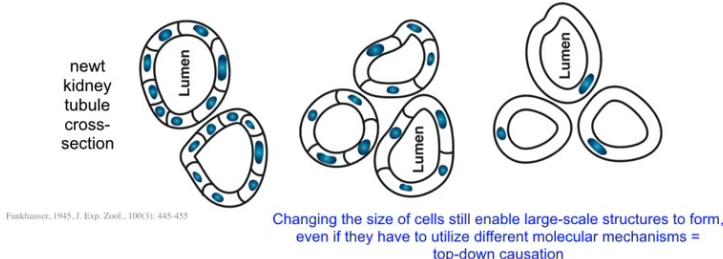
George Bubenik

This is an amazing phenomenon called trophic memory in deer antlers. Discovered by Bubenik: in one year, if you make an injury, you get a little callous, and then this whole thing falls off. Next year they grow a new rack, but it has an ectopic tine at the location of last year's injury. The system has to remember where the damage was, store the information in the stem cells or somewhere, and then when it rebuilds the bony structure, it takes a detour in that same location. Try to think about what molecular pathway model you could possibly draw for something like this. These kinds of arrow diagrams that we have in all our cell biology papers. We don't have good tools yet for dealing with things like this.

Sometimes you can't get to the goal, the target state, by your normal means. This is the kind of example where if a kidney tubule is normally built by a large number of small cells working together, if you make the cells gigantic, it will reduce the number of cells, but eventually one single cell will bend around itself and still give you that same large-scale structure, different molecular mechanisms.

## Creative Problem-Solving by Cell Groups

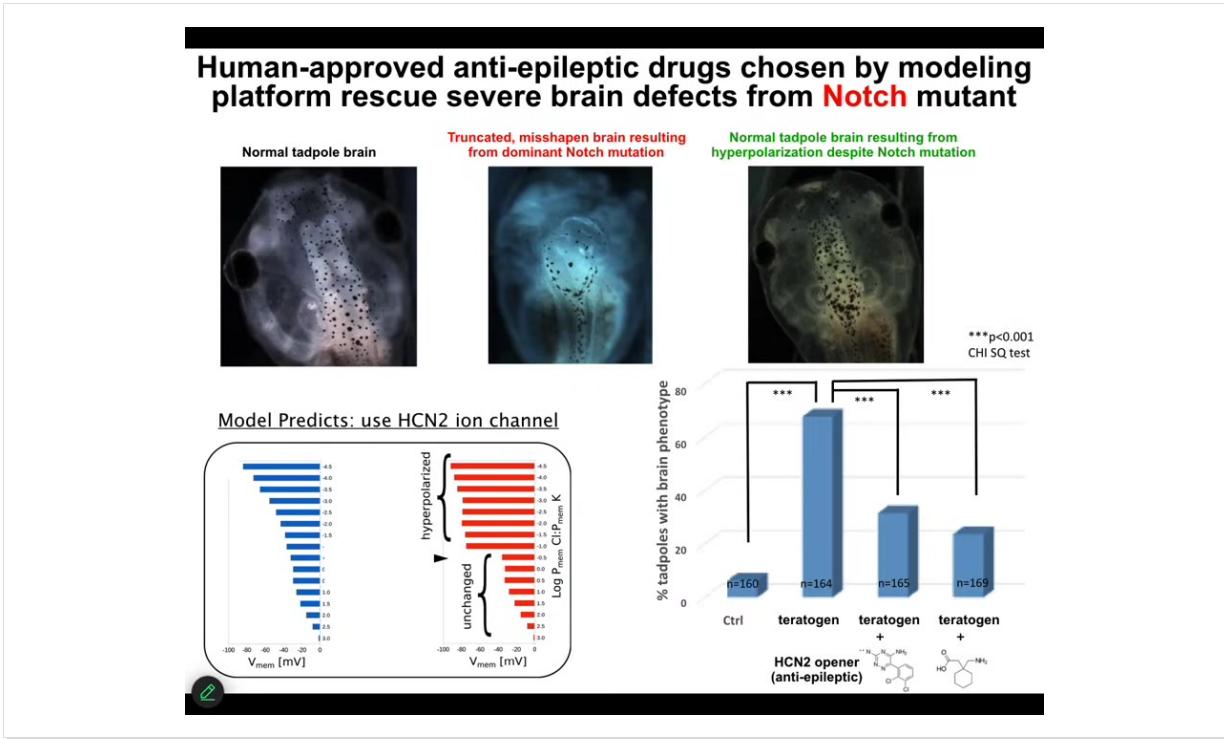
- get to the same outcome
  - despite perturbations (external **and** internal)
  - from diverse starting positions
  - **via different molecular mechanisms!**



Changing the size of cells still enable large-scale structures to form, even if they have to utilize different molecular mechanisms = top-down causation



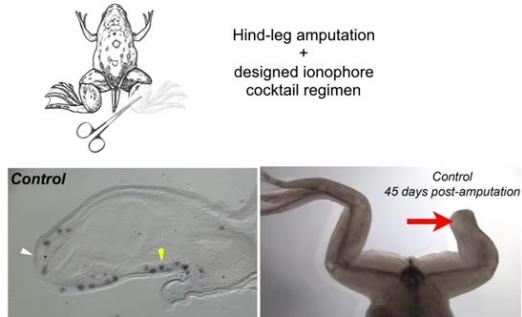
This is cell-to-cell communication. This is cytoskeletal bending. Use the tools you have to solve the problem in a novel circumstance. My cells are too big. I have tools that I can use to get here. That's the component of every IQ test. And so we use these kinds of insights to address birth defects.



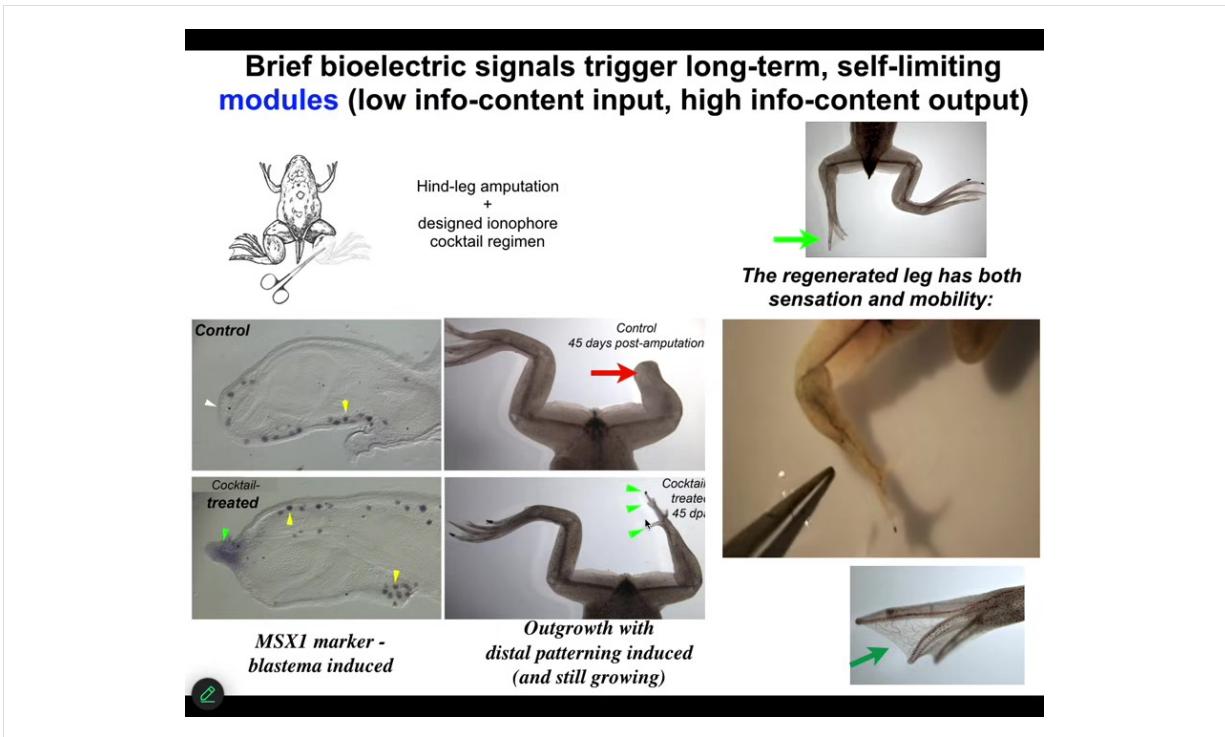
Here's a really severe brain defect in a tadpole caused by a mutation of the notch gene. Here they are bearing the mutation of the notch gene, but we have a computational platform that told us how to give them the appropriate bioelectrical pattern memories to get them to rebuild a correct brain, even though the notch is still mutated.

Some hardware defects are fixable in software. If you understand the pattern that is here and what the cells are doing, you can then ask them to build new structures, very complex structures, again, without micromanagement.

**Brief bioelectric signals trigger long-term, self-limiting modules (low info-content input, high info-content output)**



Frogs, unlike axolotls, do not regenerate their legs. What we were able to do is to provide a bioreactor with a bioelectric payload that triggers pro-regenerative genes and growth cascades.



By 45 days, you've got some toes, you've got a toenail. Eventually a pretty respectable leg.

These legs grow for a year and a half after the intervention. Our intervention is one day: 24 hours. After that the legs can grow for a year and a half, and we don't touch it during that time. That's because we're not trying to micromanage cell or molecular events. We're trying to signal to the system to commit it to the leg-building path instead of the scarring path at the beginning.

The final thing I'll show you is our work on cancer.

## Cancer Therapeutics via Resetting the Computational Boundary of the Self

The diagram illustrates the concept of cancer therapeutics via resetting the computational boundary of the self. It shows two states of cells: 'Cells Coupled by Gap Junctions, Disconnected Cancerous Cells' and 'Cancer Cell Separating Sensory & Action Range'. The coupled state shows shared sensory information and a networked sensory and action range. The separated state shows a cancer cell with a separate sensory and action range. Below the diagram, a 'Voltage dye map' shows a localized red signal (induced tumor) in a zebrafish embryo. A schematic shows 'GlyRF19A-GFP3' and 'Xrel3-2A-Tomato' with a 'hyperpolarizing channel'. Below this are two panels: 'Tumor suppressed' (labeled I) showing a zebrafish embryo with a small tumor, and 'oncogene' (labeled II) showing a zebrafish embryo with a large tumor. To the right is a 'Oncotarget' journal cover from 30 May 2014, featuring a zebrafish embryo with a tumor.

And so the idea is that when cells disconnect electrically from these pattern memories, they roll back to their evolutionary precursors and they become like unicellular creatures as far as they're concerned. The rest of the body is just external environment to them. They go where they want, they reproduce as much as they want, and we can detect this non-invasively with the bioelectrical imaging.

You can see before that's going to happen, and then if you artificially force them to remain in connection with the other cells, you don't fix the oncogene, you don't kill the cells, no chemotherapy, but you control the bioelectrics, then even though the oncoprotein is blazingly expressed here, it's all over the place. This is the same animal. There's no tumor. That's because the outcome is not directly a matter of the genetics. The outcome is a matter of the physiological decision-making that happens.

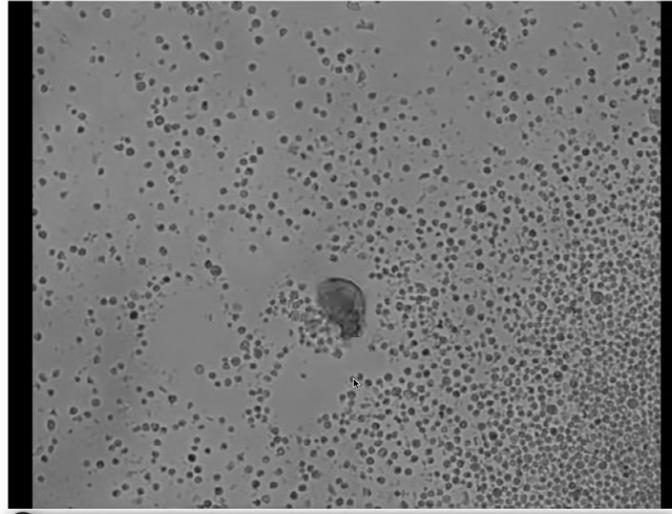
## Components of Intelligence:

- Homeostasis: the basis of goal-directed activity
- Homeostasis<sup>2</sup>: Same ends by different means
- Hierarchical, non-local control
- Hackability (software, not just hardware)
- Learning
- Creative problem-solving toward default goals
- Novel beings, novel goals: never give up



The final thing I want to show you is this notion of novel beings. Everything that I've shown you so far is creating evolutionarily provided structures. This little creature, which looks like something I might have found in a pond or a lake somewhere, if you were to sequence the genome, you would see 100% *Homo sapiens*.

## Anthrobots; can you guess the genome?



Where do  
the properties  
of novel  
systems come  
from if not  
eons of selection  
or explicit  
engineering?

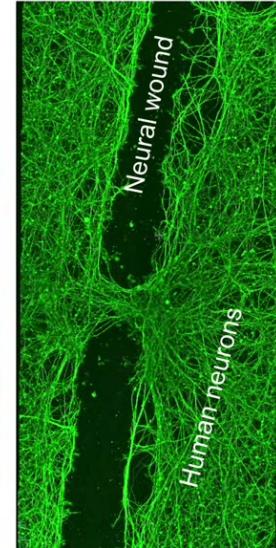
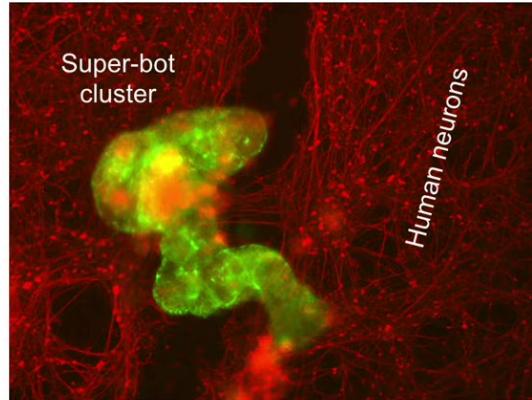
Could you guess  
the genome from  
these data?

Could you guess  
behavior and form  
from the genome?



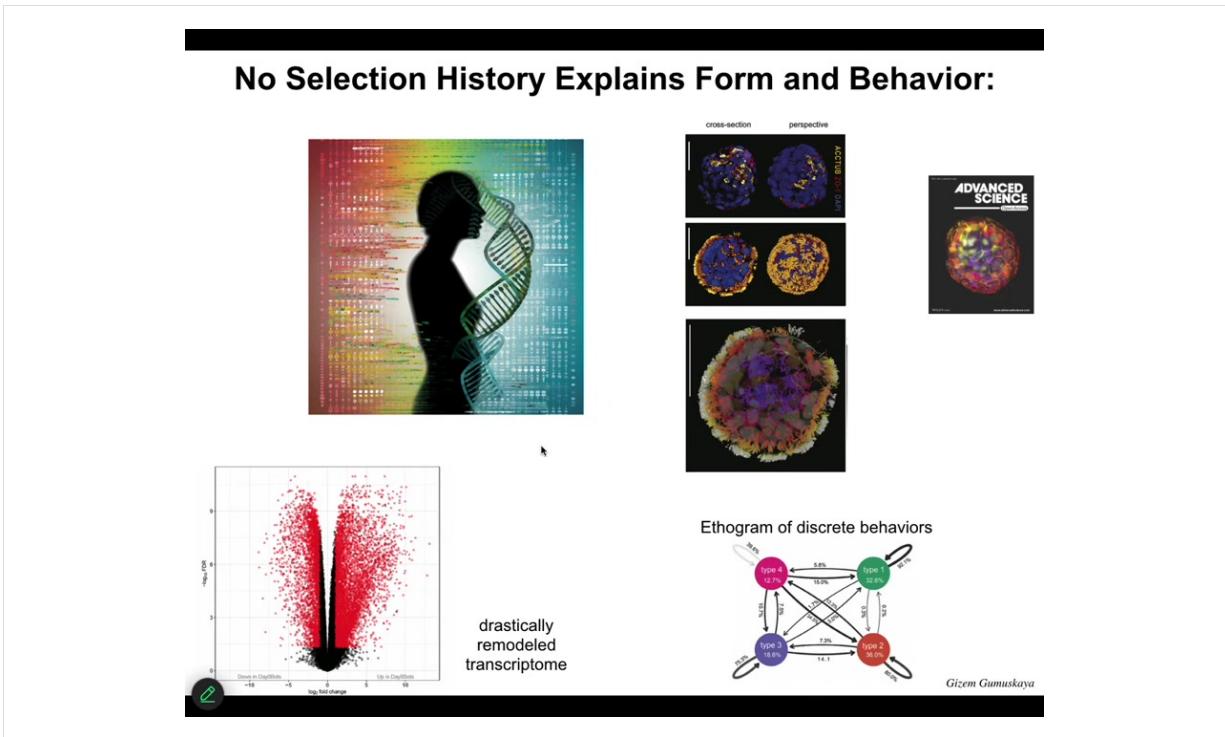
So we call these anthrobots. We allow them to self-assemble from adult, not embryonic, human tracheal epithelia. They put themselves together into this little motile creature. It has all kinds of interesting capabilities.

## Anthrobots Exert Neural Repair



Gizem Gumuskaya

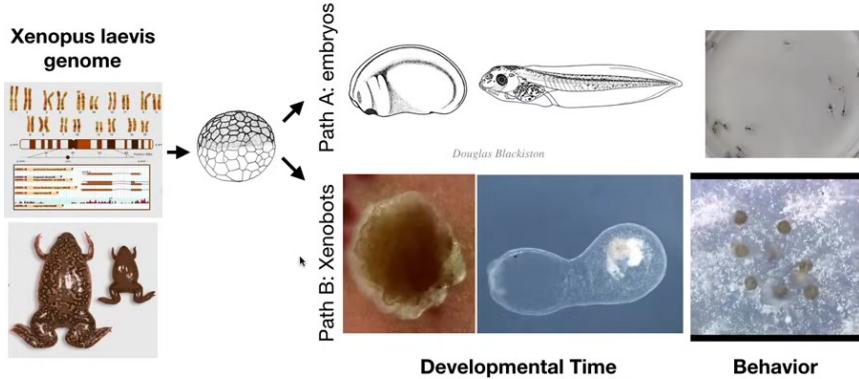
For example, if you put them down on a bed of human neurons that you've put a big scratch through with a scalpel, you make a wound, they will assemble into this superbot cluster and they will start to knit together the gap. This is what it looks like when you lift them up. They start to repair. Who would have thought that your tracheal cells, which sit there quietly in your airway, have the ability to be a self-multi-little creature?



This is what they look like with the little cilia on top. This is how they swim around. They have over 9,000 differentially expressed genes. They have their own transcriptome: about half the genome is different. We didn't touch the genome. They have exactly the same unedited human genome. There are no synthetic circuits, no transgenes, no weird nanomaterials. But what they do have is a different lifestyle. This is releasing native plasticity and they have four discrete behaviors that they can transition through.

The one thing I didn't show you, which is what all this is, is Xenobots, a similar idea made from frog epithelia.

## Biobots have Standard Genomes Novel Morphology, Transcriptomes, Behaviors



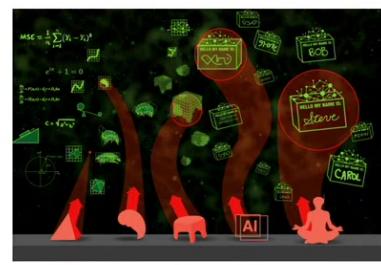
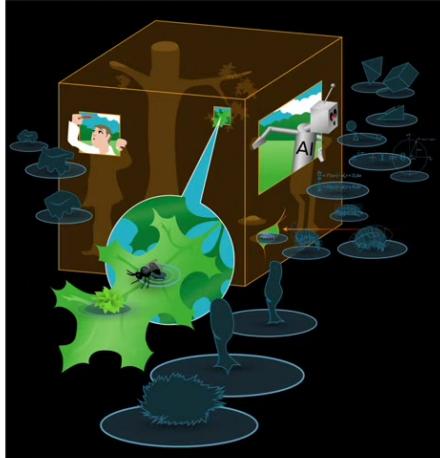
- Xenobot bodies and minds have no straightforward evolutionary back story; whence their specific competencies?
- We know when computation was done to make a frog; when was it done for Xenobots/Anthrobots?



But the question I want to ask is this: xenobots and anthrobots have no straightforward evolutionary backstory. There have never been any xenobots or anthrobots. No anthrobot has ever been selected for the ability to swim or to repair neural defects or any of the other things it's doing. Where do their properties come from? If it's not from a story about past selection forces, where does it come from?

And we know the computational cost to design a human or a frog was paid in the eons of that genome bashing against the environment and being selected. When did we pay the computational cost to make good xenobots and anthrobots? I don't have time to go into it except to point out that there's some exciting work from a number of people, and you can see all this at our symposium.

## Latent Space of Dynamic Patterns, for Physiological, Anatomical, and Behavioral Spaces



<https://thoughtforms.life/symposium-on-the-platonic-space/>



mind:brain :: math:physics



This is an asynchronous symposium on the Platonic space with the idea that the latent space of mathematics provides some very important information that impinges not only on physics, but also on developmental biology and on cognitive science. There are some interesting things to be said about the mind-brain relationship, and I encourage all of you to come and check that out.

## AI for Diverse Intelligence (and ALife, Bioengineering, biomedicine, etc. etc.) research:

tools → colleagues

nature reviews bioengineering <https://doi.org/10.1038/s44222-022-00001-9>

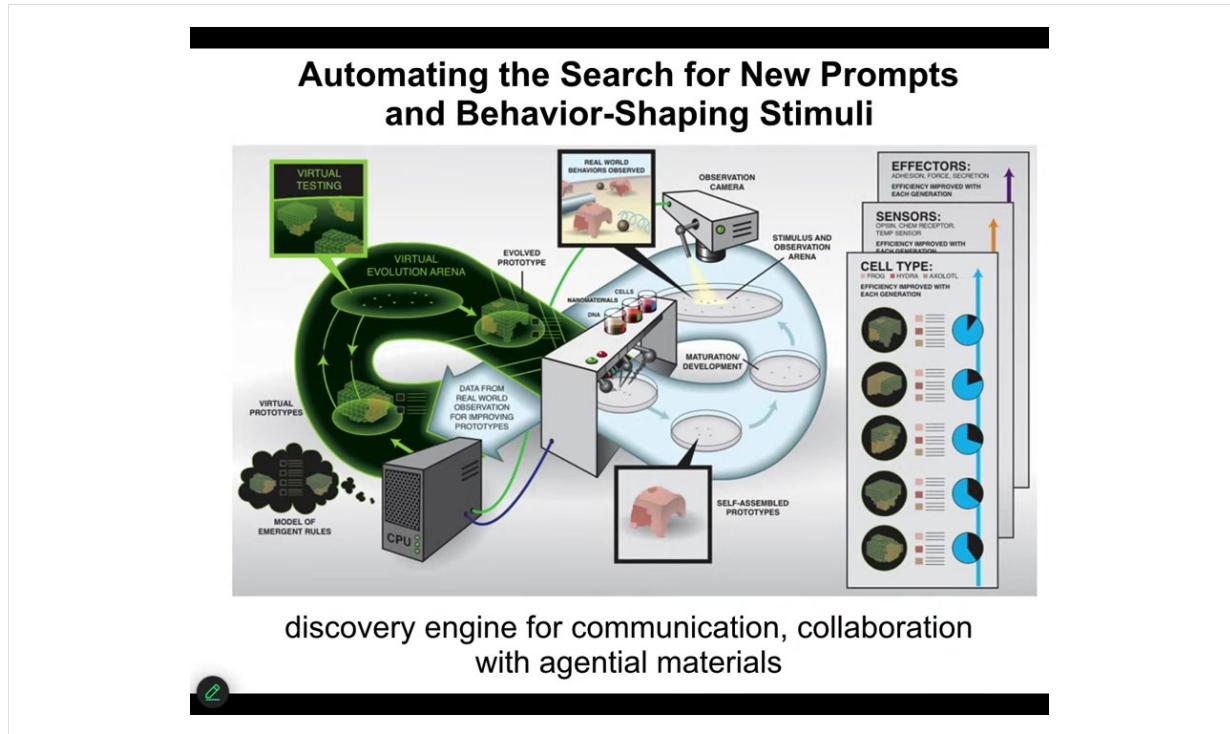
Review article

Synthetic morphology  
with agential materials

Jamie Davies\* & Michael Levin  

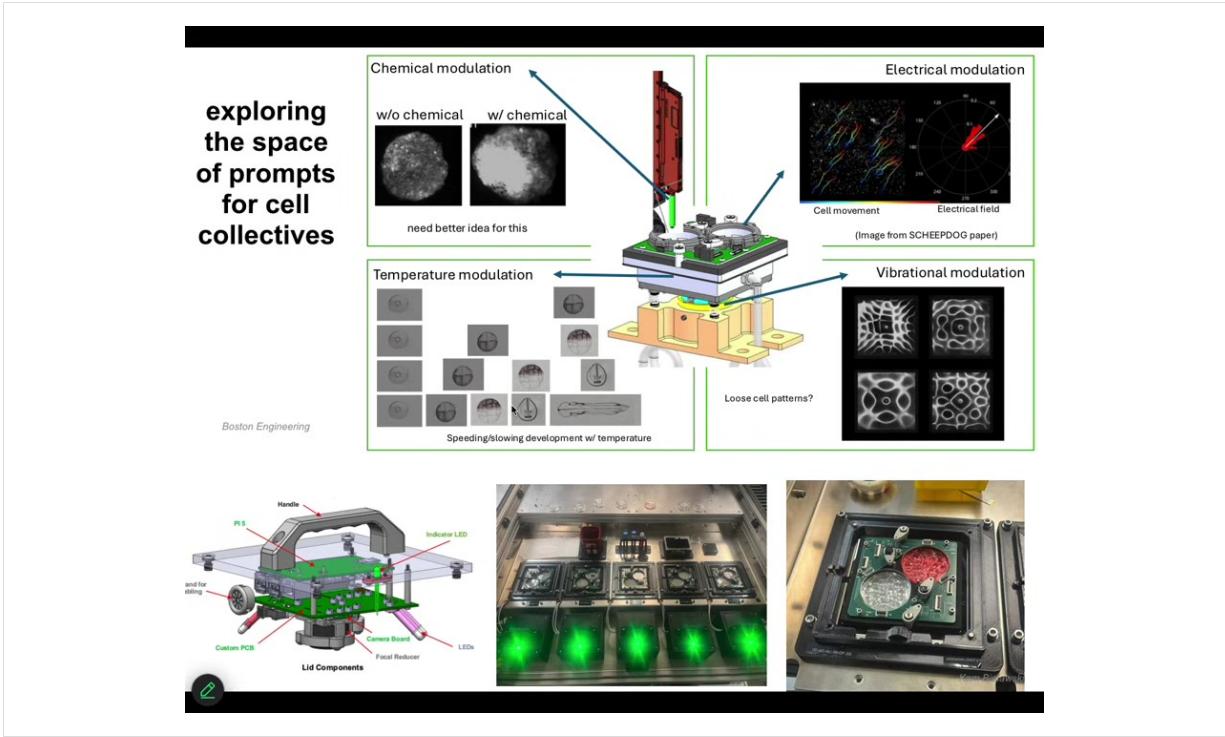
The very final thing I'll point out is our efforts to breach the two areas I was talking about today. I was talking about the plasticity of life itself, the development of cognition in minimal media, and our tools to study that.

What tools and apparatus do we have to explore that space? I want to merge those two categories and show you our latest piece of engineering, which is a robot scientist platform that makes hypotheses about the kinds of signals that have to be given to cells to get them to do whatever.



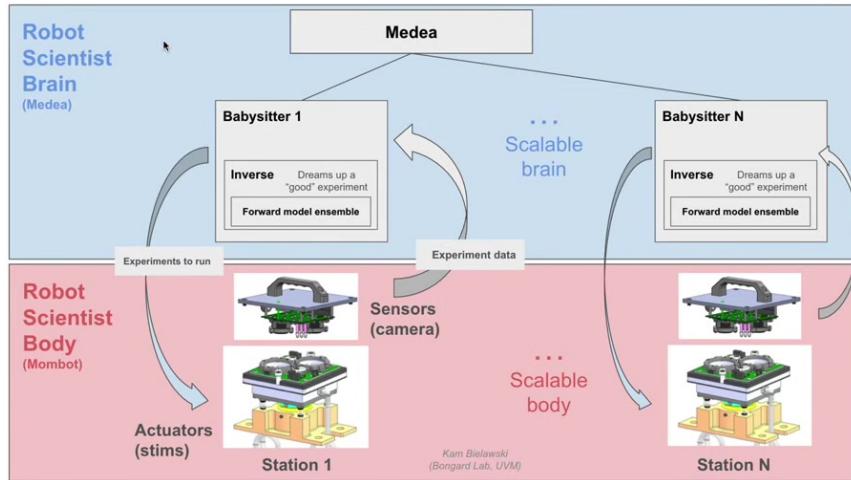
So it makes hypotheses about anatomical amorphous space. It then carries out the experiments, observes what happened, learns and changes its hypotheses, and it does this loop. The idea is for it to learn to make Xenobots, anthrobots, and all kinds of synthetic bioengineered constructs to specific goals.

We've had this diagram as a hopeful vision of what we want for the longest time. But what I can tell you now is that we've got the thing running in our lab.



It's exploring the space of prompts for cellular collectives because it has components by which it can produce chemical stimuli, electrical stimuli, temperature stimuli, vibration, other ways to communicate with the material. It has both a brain and a body.

## Introducing a new collaborator: an embodied AI scientist working in synmorpho



Like us, it has multiple concurrent cognitive sub-modules in its mind, and parallel, nested effectors in the body

The brain is a set of AI agents that try to compete for better ways to understand that space. It carries out experiments. This is what the thing looks like standing in our lab.

## Introducing a new collaborator: an embodied AI scientist working in synmorpho



**“MomBot”**

embodied AI for  
navigating  
morphospace via a  
living front end  
interface

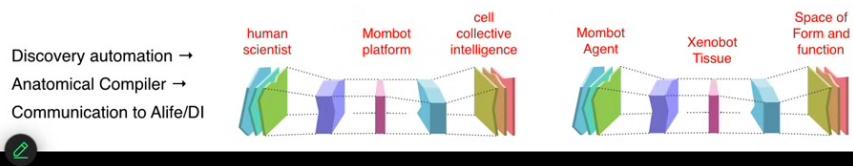


It's only been operational for a few weeks. These are very early days. This is an extremely early version of this. This is, to my knowledge, the first robot scientist that operates in morphospace. It's not doing experiments on metabolism. It's building new biological structures by exploring that space. It is an embodied AI. The AI has a body. It doesn't move through a three-dimensional space. This thing sits still. What it does is navigate morphospace. It navigates the space of anatomical possibilities via a living front-end interface, frog cells.

The future impacts are, first, the obvious thing.

## Future Impacts of this Approach:

- Useful synthetic living machines
  - In-the-body repair
  - Environmental cleanup
  - Exploration
- Robot science to crack morphogenetic code
  - Regenerative medicine - injury, birth defects, aging, cancer
  - Bioengineering, synmorpho -> anatomical compiler
- Tools to communicate with unconventional minds
  - Recognize diverse intelligences
  - Communicate with novel embodied minds
  - Be part of composite new systems



Useful synthetic living machines. If we can get that cycle of discovery sped up, that means we can do in-the-body repair with anthrobots made of your own cells, environmental cleanup, perhaps with xenobots, exploration of space. Going one step deeper, this is a robot scientist that can help us crack the morphogenetic code. If we understand how to communicate new goals to groups of cells, that leads to new regenerative medicine, and of course, bioengineering towards the notion of an anatomical compiler, being able to specify arbitrary forms. We're hoping that the next generation of this kind of system will help discover that.

But even going one step deeper is this idea that what this really is is a tool to communicate with unconventional minds. You can think about it two ways. You can think about it as a tool where the human scientist, through this MomBot interface, can actually interface to the collective intelligence of cells. We specify our goal states. The MomBot figures out what stimuli those have to be and gets them to execute exactly as your high-level cognitive processes issue commands that, through your body, have to get molecules to act in your muscle cells so you can carry them out. But another way of thinking about it, and I think these are both useful, is that the MomBot is the active agent. The frog tissue is the translator interface. What it's exploring is the space of possible forms and functions. There are different perspectives on this. You can take the perspective of different parts of the system. I think all of these are important because in the future, it isn't only language models and standard AIs that you have to think about.

## Plasticity of Life: beyond standard forms

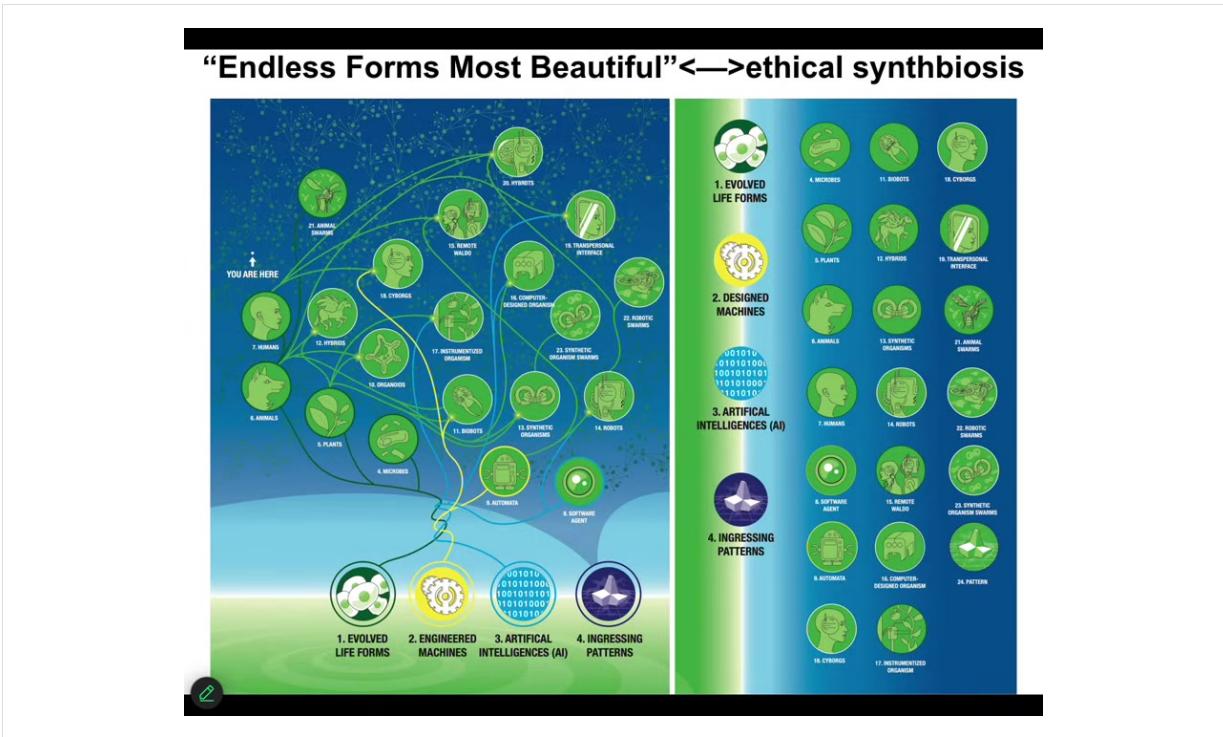
The diagram illustrates the plasticity of life across five levels of organization: CELL, TISSUE, ORGANISM, SWARM, and ECOYSTEM. Each level is represented by a pyramid containing icons of life forms. The diagram shows how various synthetic components can be integrated into these levels, blurring the lines between biological and technological life.

- CELL:** Exogenous DNA, Nanomaterial, Scaffold, Smart Material, Bacteria.
- TISSUE:** Inorganic Components, Sensors, All Instance Effects.
- ORGANISM:** Sensors, All Instance Effects.
- SWARM:** Robots, Soft, Hard.
- ECOSYSTEM:** Sensors, All Instance Effects.



This kind of plasticity, the fact that we are all going to change both biologically and technologically, slowly and gradually, because life is interoperable, every layer can be interacted with by synthetic engineered components.

And eventually, this is going to be your neighbor. It's not going to be as easy as I'm a real biological and this is just a machine. It's a language model. What we're going to be talking about is really stretching the notion of what an active agent is and the fact that you do or do not have components that were evolved by trial and error over millions of years or components that were designed by engineers. Those things are going to mix in a way that makes ancient categories completely irrelevant.



This is the space of possible embodied minds. Any kind of living material, engineered material, software, and patterns from that platonic space are some kind of embodied agent.

Cyborgs and hybrids and chimeras of all kinds, this is Darwin's "endless forms most beautiful"; all the variety of life on Earth is a tiny little corner of the space. These things already exist to some extent. We will be living with all of them. We have to learn to understand beings that are not like us so that we can enter an ethical synthesis with these beings. This is why developmental biology and bioengineering are so critical. It's not just the useful living machines, it's not just evolution that we learn to understand, but we can use that as a model system to start to communicate with minds that are not like ours. I'll end here by saying that improved communication, not micromanagement, but communication and collaboration with living matter is going to lead to definitive regenerative medicine and bioengineering.

## Summary:

- Improved communication, collaboration with the agential material of life -> definitive regenerative medicine + bioengineering
  - not just a genetics, cell biology, or biophysics problem - it's part of cognitive science (and CS)
- Bioelectricity = interface to
  - rewritable goal states, borders of decision-making agents
  - problem-solving competencies of cells and tissues
  - cognitive glue up-scaling intelligence
- Bioengineering as a model system to answer Big Questions
  - how do embodied minds ingress and scale
  - how to recognize unconventional cognitive kin
- It's Behavioral Science all the way down (and up)

Matter To Life

Matter To Mind To Life

Mind to Matter To Life



So everything that you're working on is not a genetic, cell biology, or biophysics problem. It's part of cognitive science and computer science and vice versa.

Bioelectricity is a convenient interface to that collective intelligence, but only because evolution loves it. It's not magic by itself and there are other kinds of cognitive glue as well. Bioengineering is an amazing model system for big questions. How do embodied minds ingress into the physical world? How do they scale and transform? How can we recognize unconventional minds that are all around us? I tend to think that the basic diagram of the different sciences is wrong and that it's behavioral science all the way down and all the way up. We can think about different ways to think about it. Matter to life, yes, but also matter to mind first and then life. We can talk about that if anybody has questions. I think cognition is wider; it doesn't come late in life. I think it's a much earlier and larger-scale phenomenon and only a subset of that is life. We can even go further. That's it.

## Thank you to:



**Postdocs and Staff Scientists:**  
Patrick McMillen - bioelectric basis of cell collective intelligence  
Vaibhav Pai - voltage gradients in eye/brain induction and repair  
Tal Shomrat - planarian memory  
Juanita Mathews - bioelectronics of cancer in human tissues  
Nestor Oviedo, Junji Morokuma - planarian bioelectric pattern memory  
Dany Adams - bioelectric face prepattern

**Ph.D. Students:**  
Brook Chernet - bioelectronics of cancer  
Sherry Aw - bioelectric eye induction  
Fallon Duran - planarian bioelectric pattern memory  
Gizem Gumuskaya, Nic Davey - Anthrobots

**Undergraduate Students:**  
Pranjal Srivastava, Ben G. Cooper, Hannah Lesser, Ben Semegran - Anthrobots  
Maya Emmons-Bell - planarian barium adaptation

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

**Bongard Lab at UVM:** Josh Bongard, Kam Bielawski  
Krishna Srinivasan, Shawn Beaulieu, Piper Welch, Thomas Varley, Jeantyne Lunshof

**Blackiston Lab at Tufts:** Douglas Blackiston and Tomas Gonzalez-Zugasti

**Illustrations:** Jeremy Guay @ Peregrine Creative

**Model systems:** frog, fish, human cells, animals, mice, and many others

**Collaborators:** Richard Watson, Chris Fields, Anil Seth, Robert Chis-Ciure, Giovanni Pezzullo

**Funding support:** CRREL, JTF, John Abele, DARPA, Paul G. Allen Frontiers Group, NIH, NSF

This material is based in part upon work supported by the Broad Agency Announcement Program and the Cold Regions Research and Engineering Laboratory (ERDC-CRREL) under Contract No. W913E524C0012. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Broad Agency Announcement Program and ERDC-CRREL.

**Disclosures:** Fauna Systems, Astonishing Labs



**BOSTON ENGINEERING™**  
Imagine the Impact™  
Heather Moyer, Mark Smithers & team

**softmax**



I'm going to stop here, thank the postdocs and the grad students who did all the work, and Josh's team that did all the coding for the Mombot, and the team in Boston Engineering who built the thing for us. Disclosures: there are three companies that have licensed some of the work that I've shown you today and that provide sponsored research agreements. Thank you so much 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](mailto:adi@aipodcast.ing).