

A LECTURE COMPANION

**"Diverse Intelligence: new frontier for
Origins, Possibilities, and Diseases of
Minds by" M. Levin**

Michael Levin

Recorded on July 10, 2026

About this document

This document is a companion to the recorded lecture *"Diverse Intelligence: new frontier for Origins, Possibilities, and Diseases of Minds by" M. Levin*, recorded on July 10, 2026.

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 ~54 minute talk titled "Diverse Intelligence: a new frontier for Origins, Possibilities, and Diseases of the Mind(s)" which I gave for an audience composed of mental health professionals and neuroscientists.

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.



And thank you for that extremely kind introduction and for giving me an opportunity to talk to this audience. If anybody wants to chase down some of the things that I'm only going to briefly mention, this is our website that has all the primary papers, the data sets, the software, everything is here. And then this is my personal blog where I write around what I think some of these things mean.

So what I would like to do today is four things. First of all, just briefly introduce myself and why I think some of these things are of relevance to you. And then I'm going to go over a framework that we've been developing that is designed to really advance a diverse intelligence research and applications. And I'm going to show numerous examples that are directly grounded in neuroscience and capabilities of brains and so on, but I'm going to take it in a much, much farther direction. I will show you one model system that I think is a good stepping stone for communicating with semi-alien minds, and that is morphogenesis, the ability of individual cells to work together as a collective intelligence that solves problems in anatomical space. And then at the end, I'll say some of the most speculative things about what I think is the future of the sciences of minds.

This is Our Future:



Non-neurotypical doesn't even begin to cover it

These issues are not about "them" - they're about us (broader)

So the first thing I would point out is that both in basic science but also for the mental health profession, I think your client list is about to get very, very strange. And I think that this notion of non-neurotypical doesn't even begin to cover what's going to happen. Basically, as I'll point out, there are tremendous changes that are going to be made both to the physical embodiment and the cognitive capabilities of various beings who are going to need different degrees of help. And we really have to start to understand how we're going to relate to them. I'll talk more about this. And we're on both sides of this door, right? As for you as mental health professionals, you need to know how to relate to new clients that are basically not the same standard humans that have been here for thousands of years. And again, for all of us entering this new world, things are going to really change.

Background:

- The road to transformative regenerative medicine leads directly to freedom of embodiment. There is no timeline in which humanity keeps the same hominid form factor and psyche into the future.
 - This is now, not sci-fi
- The road to a mature species leads through overcoming our mind-blindness
 - This means understanding our nature, our origins, our potential, and the meaning of “our”.

So the background to these claims is basically the following: that the road to transformative regenerative medicine, I’m going to show you how we’re doing this, meaning that the ability to induce growth and form and repair on demand leads directly to something I call freedom of embodiment. In other words, I don’t think there’s any way to avoid it. I don’t think we want to avoid it, but it’s very clear that these kinds of modifications to our minds and bodies are not some separate thing that we can choose to address or not. Basically, our efforts at radical healing are going to lead to this. They’re going to lead to this capability of basically having whatever kind of body and possibly whatever kind of mind you want.

I don’t believe there is any timeline in which humanity simply keeps the same form factor of hominids that we were sort of left with by the various actions of cosmic rays hitting ourselves over the evolutionary timescale. This is the embodiment that we have. I don’t see any future in which we just leave it exactly as that, as humanity moves forward into a mature species.

So this is happening now. This is not some future sci-fi world. This is actually happening right now. And I think the other step to becoming a mature species is to overcome what I think of as significant mind blindness. We have tremendous difficulties visualizing the possibilities of minds and unconventional embodiments, never mind communicating with them or relating to them, but even to visualize the minds at other scales of time, of space, operating in other problem spaces.

So the thing to keep in mind is that all of us are part of a continuum. We started out as a single cell, both developmentally and evolutionarily. There was a slow, gradual set of transformations. There is no magic lightning bolt during development at which you used to be the province of chemistry and physics, but now you're the subject of behavioral science and psychiatry. This is a slow and gradual process. And not only did it happen on these timescales naturally, we now have another continuum that we are at the center of, where through technological change and biological change, all of these things are now up for grabs.

And so this is really fundamental to questions of the mind. Whatever you think it is that a modern adult human has in terms of hopes and dreams and personal responsibilities and all of those kinds of things, we have to ask ourselves where on this timeline did that appear. We typically don't talk that way about cells, but eventually we do. And so what is the story of transformation that we're telling here? And that transformation is going to have a huge opening in terms of capabilities.

Slide 4 of 48 · Watch at [5:16](#)

Novel Beings, Novel Minds: it's not about LLMs



PERSPECTIVE

Artificial Intelligences: A Bridge Toward Diverse Intelligence and Humanity's Future

Michael Levin

Most of the issues brought up by AI are reflections of fundamental unsolved questions about ourselves


So, you know, we've all heard issues with AI and language models and trying to compare them to humans and so on. But I think there are two things to know here.

One is that none of this is really about language models or AI sitting in a box somewhere. This is first going to come forward when your neighbor has some percentage of their brain or body replaced with novel technologies, with engineered systems, with smart implants and all of that. Again, these already exist. And people are already getting them for therapeutics, but there's also going to be a great variety of optional modifications that people are going to want to take. And so trying to understand whether somebody is a proper human or a quote unquote machine, I think that whole distinction is not going to survive the decade. These binary distinctions are simply not going to do. And I think that most of the issues brought up by AI are basically reflections of fundamental unsolved questions about ourselves, our own origins and our potential.

Slide 5 of 48 · Watch at [6:19](#)

Summary:

- Intelligence, and thus diseases of cognition (not well-captured by models of organic disease - hardware issues) are way older than neurons and brains
- Borrowing tools of somatic psychiatry, because "you" are not the only intelligence in your body, and 3D "behavior" is not the only kind of behavior

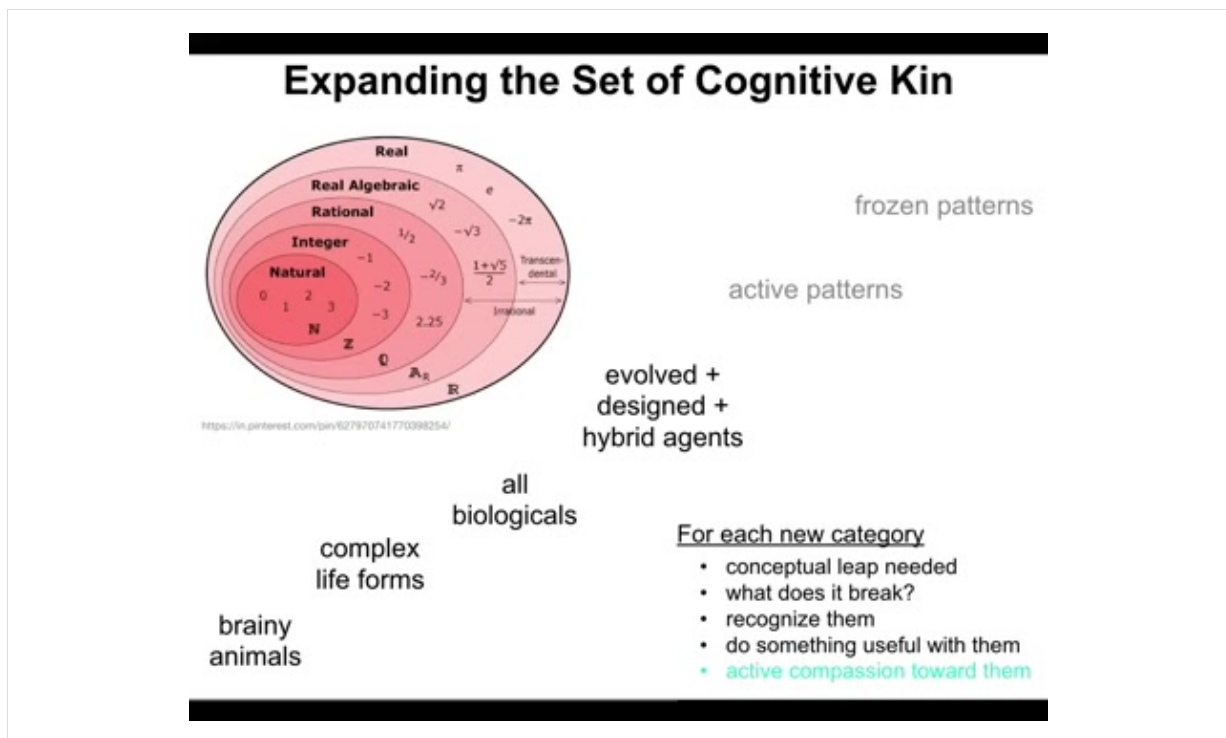


So what I'm going to tell you today is basically this: intelligence and diseases of cognition that are not well captured by models of organic disease, they are not all hardware issues, are way, way older than neurons and brains. The kinds of things that you all study are of intense and profound relevance to problems in developmental biology and evolution. None of this is specific to neurons or to brains. And what we've

been doing is borrowing tools of behavioral science, of cognitive science, and so on. We've been borrowing these tools and really applying them to all kinds of other problems, because the you that I'm speaking to right now that exists in your body is not the only intelligence that's present in your body. And the 3D behavior of you moving around in the outside world is not the only kind of behavior. And so what I think there is the possibility of now is a very rich kind of virtuous cycle, a positive feedback loop between the progress in cognitive science and what I can only summarize as multi-scale biology. So I'll show you momentarily, but basically this idea that cognitive and behavioral sciences have a lot to say about things that are not neurons or brains. And conversely, we can, on the biology and bioengineering side, the things that we discover can actually feed back and maybe help your community understand the origins of the kind of cognitive systems that you're all dealing with. Okay, so this is why I think these things kind of come together. And so now let's take a look at a framework for trying to understand what I mean by diverse intelligence.

So what I would like to do is to create a way to think about all kinds of different beings on the same sort of scale, regardless of their origin story or their composition. It shouldn't matter what you're made of or how you got here, whether it was natural evolution or engineering or some combination thereof. We should be able to think about what all significant minds have in common. And this means being able to think about not just familiar creatures, such as primates, birds, maybe an octopus or a whale, but also weird kinds of things like beehives and synthetic biology, which I'm going to show you some examples of, engineered new life forms, AIs, whether software or robotic, maybe someday exobiological agents as well. I'll say that there are a lot of people thinking about connecting to life off planet. And I'll just say that if we can't wrap our heads around what's going on with the cognition inside our body organs and our cells, there's no way we're going to be able to connect to actual aliens. This is a stepping stone.

And so this, of course, this idea is not new to me. So here's Rosenblueth, Wiener, and Bigelow in 1943, trying to paint a picture of the spectrum, how you get from passive matter all the way up to human-level metacognition and those kinds of things, and what the steps are in between. This is a very sort of cybernetic view, and I think a good one. And the framework that I'm interested in, it's laid out here in detail, has to move experimental work forward. It cannot just be philosophy or linguistics. It has to move experimental work forward. It has to have implications that drive new discoveries in biomedicine and synthetic morphology. But also I think it has massive implications for ethics and for how we relate to these beings.



And I see it kind of like what's happened in math. So we started out with natural numbers here, and then progressively by breaking rules, breaking arbitrary limitations on what we thought these objects were, we discovered yet more and more different kinds of important objects out there. And that always happens when you examine your assumptions and you examine your categories and you say, okay, but what happens if I give up this axiom or that axiom? What can I discover? And so this is kind of the trajectory that we've been on. I'm not going to get a chance to talk about any of this, but basically, looking at beyond just brainy animals, what does the material of life actually offer? And then some other things that are kind of combinations of evolved, designed, and hybrid agents. And for each of these, we have to think about what is the conceptual leap that's needed, and what can it do for us? And then how do we actually relate to these novel beings that we're discovering? And I think the set of other minds all around us that are difficult to notice right now is going to grow enormously. I think it's going to expand greatly.

Easy to Detect Mind in Brainy Mammals




https://www.youtube.com/watch?v=f7SVet_sJNo

Why so obvious? Same spatiotemporal scale, same space, similar goals

So let's just look. When we look at things like this, right, so brainy mammal, so here he is, he's setting up a little accident scene. He's got a very good theory of mind about his owners. He knows what's going to happen when they see this, right? Like, very clever. He's even going to look to see, make sure that they're watching and things like that. These kinds of things are so obvious to us because it occurs on the same spatial temporal scale as we do, so the behavior in the 3D world, same scale of space and time as us, say similar goals of safety and all of that. But even with brains, things are not simple.

Even with brains, things are not simple



Minimal brain structure
or function (Savant syndrome)
cases of high performance

Mind & Matter 20(1), 15-49 doi: 10.1017/med.2015.11

Cases of Unconventional Information Flow
Across the Mind-Body Interface

Karina Kaufman
Faculty of Dentistry, University of Toronto
Toronto, Canada
and
Michael Levin
Alpha Discovery Center
Tufts University, Boston, USA
and
Wynn Institute for Biologically Inspired Engineering
Harvard University, Cambridge, USA

Figure 2. Select cases of reductions in brain matter with normal function. **[A]** Image from (Feuillet et al. 2007) showing a white collar worker case of extreme hydrocephalus; he led a normal life as a civil servant, who possessed an average IQ of 75. During his neurological assessment at age 44, his (i) CT scan and (ii) T1 weighted MRI scans with contrast showed extreme ventricular enlargement. LV indicates lateral ventricle, III and IV indicate the third and fourth ventricles, respectively. **[B]** Image from (Alders et al. 2018), showing the case of a 60-year-old with a bad mood with massive ventriculomegaly and severely reduced cerebral mantle and corpus callosum, that went largely unnoticed. The left column is T1 weighted MRI images taken in the transverse, coronal, and sagittal planes of the patient. The right column represents T1 weighted MRI scans of a healthy control. **[C]** Image from (Persad et al. 2021), imaging of a Canadian living a normal, independent life with massive hydrocephaly. MRI scan taken from the axial view (plane parallel to the ground) at the level of the lateral ventricles (arrow points to extremely thin layer of cortical mantle, LV stands for Lateral Ventricle). **[D]** Image from (Asaridou et al. 2020), showing the T1 Weighted MRI scans of a child born without left hemisphere (i) taken in the coronal plane, (ii) taken in the axial plane. The child had normal cognitive development and language skills despite hemispherectomy of the left hemisphere and near-absence of the corpus callosum. All images re-used with permission.


For example, Karina Kaufman and I reviewed a number of cases where there's a really radical mismatch. These are rare clinical cases where there's a radical mismatch between the amount of human brain tissue that you would expect to be present and the cognitive performance of that individual. So already we know there's, I mean, this is not common, but it does occur, and so we know there's something going on here. And you can try to patch it up with the talk of redundancies and things like that. But fundamentally, the models we have do not predict that things like this will happen. They can be maybe shoehorned into it, but they don't predict it.

Sensory-Motor Architecture Highly Reconfigurable




no eyes

eye



iv

100µm



Configuration	Percent of tadpoles (%)
Wild type	~40
No eyes	0
Tail eye	~20



Behavioral Testing Device

Douglas Blackiston



RESEARCH ARTICLE
Ectopic eyes outside the head in *Xenopus laevis* provide sensory data for light-modulated learning
(Douglas J. Blackiston and Michael Levin)


Brain dynamically adjusts behavioral programs to accommodate different body architectures, no lengthy adaptation needed!

Ectopic eyes on tail provide vision!

And then there are many other really interesting things that we don't understand. For example, this is a tadpole of the frog, *Xenopus laevis*. Here's the mouth, here are the nostrils, here's the brain, the spinal cord. And so what you'll notice we've done is we've made sure the primary eyes don't form, but we put an eye on its tail. And this eye puts out an optic nerve. Here's the optic nerve. It does not connect to the brain. It synapses sometimes on the spinal cord, sometimes on the gut, sometimes nowhere at all. And then we built a machine, this device, to automate the behavioral training of these animals for visual cues. And what we discovered is that they can see. They can learn in visual cues. Even though this thing does not connect to the brain, even though they have a completely different sensory motor architecture, immediately, out of the box, this works. No new rounds of mutation, selection, adaptation to this completely different form factor of the sensory and the processing systems. No problem. Right out of the box, it works. Where does this incredible plasticity come from? Why is it so highly reconfigurable? Why don't you need rounds of selection to make this thing work? Starting out with 100% normal, genetically normal. We don't touch the genome, by the way. So 100% genetically normal animal. Why does this work?

Memories Regenerate, Accessible to New Tissue

implications for brain repair, longevity, and augmentation



Tail Shows...

If hearing algorithms can take over visual processing areas, what can your thought patterns do with an extra hemisphere?

© 2016. Published by The Company of Biologists Ltd. Biology Open (2016) 5, 1171-1180 doi:10.1098/bioopen.2016.00140

HYPOTHESIS

Vertically- and horizontally-transmitted memories – the fading boundaries between regeneration and inheritance in planaria

Marion Neuhoff^{1,2}, Michael Levin^{1,2} and David Reiche^{1,2,3,4}

We should also talk about the ability of memories to move through tissues. So these are planaria. You'll see some more of them in a minute. These are flatworms. They are similar to our direct ancestor. They have a true centralized brain, same neurotransmitters that you and I do. They're highly regenerative, so you can chop them into pieces. Each piece regenerates. But they're also pretty smart. You can train them, for example, place conditioning and feed them in these bumpy little circles. And what you find is that if you chop off their heads, the tail sits there doing nothing. Eventually it regrows a brand new brain, and now behavior kicks in, and you can, by testing them in that same machine that I just showed you, find out that these animals remember the original training.

So what's happened here is that the memory was somewhere. It was not just in the brain, that's clear. It was imprinted onto the new brain. The new brain has access to it, okay, after it appears, so this new tissue. So the new tissue has access to these memories formed by the old tissue. So this has implications for understanding what's going to happen when we trigger regenerative responses in brains. So for patients, longevity, even augmentation. If, for example, the algorithms that power hearing can take over visual processing areas in the blind, what else can your thought patterns inhabit? What other real estate can they take over? If we grew you a new hemisphere, would they take over? Would they move into that space, right? So the ability, this intersection between behavior and morphogenesis, when novel tissue, be it brain or not, grows is extremely interesting.

Plasticity of Subject Matter Requires Plasticity of our Perspectives

Douglas Blackiston

Memories Persist Through Drastic Refactoring of Substrate & Re-map onto a new embodiment

3 Perspectives

- Of the caterpillar facing impending singularity
- Of the butterfly, saddled with mysterious behavioral traits from a past life
- Of the memory itself, facing the paradox of change

entropy

Research
Self-Improving Memory: A Perspective on Memories as Agential, Dynamically Reinterpreting Cognitive Glue
Michael Lewis

The stability of memories during brain remodeling: A perspective
Douglas I Blackiston¹, Ted Brown^{2,3}, and Michael Lewis^{1*}

© 2015 Entropy Research, Inc. All rights reserved. Entropy Research, Inc. is a registered trademark of Entropy Research, Inc. in the United States and other countries.

We can also think about examples like this, where caterpillars turn into butterflies. In order to do this, they rip up their brain, basically dissolve most of the connections, kill off most of the cells, build a completely new brain suitable for a very different kind of lifestyle. But one thing that has been found is that if you train the caterpillars, the memories are recovered in the butterfly, or moth. Not only do they persist, I mean, that's a big enough issue right there, is how do memories persist when all the synaptic connections are broken, right? So that has implications for theories of where memory is and things like that. But the bigger issue is that the actual memories of the caterpillar are of no use to the butterfly. In other words, butterflies don't eat the leaves that the caterpillars got as a reward. They don't crawl the way that the caterpillars behaved. It has to be remapped. The memories have to be remapped onto completely new architecture, and so again, that plasticity that's baked into living tissue is there for us to think about and perhaps take advantage of in terms of how do you remap the kinds of cognitive capabilities that you have, including your memories and your personality and so on, onto new hardware, whether that be new hardware of the body or of the brain. And there are lots of interesting things we can talk about in terms of taking the perspective of the caterpillar facing the signal, the perspective of the butterfly that's saddled with some memories that has no idea where they came from because it wasn't part of the training in this embodiment, and yet it has those memories, and also perhaps of the memory itself, which has to be remapped. We can talk about that. So this plasticity, I think, is very interesting, and it's telling us something important. Here's another example that I like a lot. If I said to you, I would like to take a reptile, a

shy, slow-moving turtle, I want it to work at a cat-like speed of life. I want it to be playful. I want it to move at the speed of a cat. What would you have to do? And you might think that, wow, millions of years of evolution, or perhaps some sort of neuro-engineering that we have no idea how to do.

Slide 12 of 48 · Watch at [17:20](#)



Not much.

Even small adjustment to physical embodiment unlocks a new cognitive domain.

What was the latent space of its possibilities? These engineering changes are the periscope to find them.

(Thanks to Marsa Hickey for finding this)

<https://x.com/AMAZINGNATURE/status/184532258879217914>

Actually, it turns out you don't have to do much. So here's this guy. He put his turtle on the little skateboard, and immediately unlocked this incredible change in behavior. The cat looks puzzled. That makes sense. We should all be puzzled at this. The turtle has no problem moving at the relevant speed. He clearly wants to play. He's keeping up with the animal. What was the latent space of possibilities? What else could you unlock with very minimal changes to the embodiment? Have these turtles been around for millions of years, just capable of all of this and not being able to do it because their body didn't last? And what is going to be unlocked in humans when we have things that are much more capable than a skateboard for us to connect to? So that kind of plasticity of form, of function, of in fact of integrated form and function, asks us to really understand what's happening when cells build and accommodate new structures.

And so what I'm going to do for the next, let's say, 15 minutes, is to talk about morphogenesis as a model system. So let's talk about the cells that we're all made of to understand what I think of as the agential material of life, this multi-scale intelligence that underlies our behavior and our cognition.

Slide 13 of 48 · Watch at [18:37](#)

All Living Embodiments are Collective Intelligences




Lacrymaria = 1 cell
no brain
no nervous system

high competency
at cell-level
agendas

Can you reward or punish it?

Comment
<https://doi.org/10.1038/42208-024-00000-0>
nature machine intelligence
Discussions of machine versus living intelligence need more clarity
Nicolas Beaudou & Michael Levin

So first, just to remind us that not just beehives and ant colonies, but all of us are collective intelligences. We are all made of parts. Everything that we are able to do is underwritten by the cognitive glue, by the policies and mechanisms that keep our parts aligned towards specific goals. This is the kind of thing we're made of. Now, this is a free-living organism called the lacrimaria, but you get the point. It's got no brain, no nervous system, extremely competent in its own little environment, doing the things that it needs to do to survive. And these are the kinds of things that work together to build our bodies. Now, when I ask people, could you reward or punish something like this? Most people say yes. They feel you could. If I ask you, could you reward or punish a chemical network? Most people say no, chemicals don't care what happens. You can't reward or punish them. But of course, what is this made of, right? That's exactly what this is made of, is a chemical network. So what happened? What's happening from here to there?




And what I want to point out is that even chemical networks, in fact, small chemical networks, the smallest one is just four nodes, can do Pavlovian conditioning. Chemical networks, as we've now discovered, can do at least five or six different kinds of learning. They can do habituation, sensitization, associative conditioning. They can count to small numbers. Walter Fontana showed they can do probabilistic inference. All of this, not brains, not neurons, not even cells, small chemical networks. The material that we are all made of, that our cells are made of, is already cognitive right at the beginning. It can do primitive kinds of operations that are easily recognizable to any behavioral scientist.

Now, there's a couple of interesting things that I'll just point out about this work. One is that this is a kind of molecular placebo, because when you train this network in an associative conditioning context, when I hit it with a drug that causes some sort of a strong outcome, and then another drug in paired presentation where the neutral stimulus normally doesn't do anything, if they're applied together, the system learns, and then the neutral stimulus becomes the conditioned stimulus and activates the response. It's a kind of placebo effect, basically. At that point, the system reacts to the inert drug as if it were something else, because its past history of experience has convinced that that's what this is. So you can already start to see a kind of the scale down of the things we're all used to in human medicine that scales down to these kinds of things.

The other interesting thing about this is that we found that if you look at causal emergence, so these are metrics that basically tell you the degree to which the whole is more than the sum of its parts. So people use this to distinguish between locked-in patients versus comatose patients versus just a pile of neurons. You know, the question of is somebody home? So Giulio Tononi and others use these kind of metrics. What we've shown is that if you use them to look at these chemical networks, you will find that quite often with training, the causal emergence goes up, so their integrated agency goes up the more you train them, but also the more it goes up, the better they get at learning, and there's this incredible feedback loop, which actually is asymmetric; it points in one direction, because... If you force them to forget, you do not lose the gains that you made in causal emergence. That's very interesting. It means that what happens in these networks is they're primed for higher intelligence and higher agency. It's not a symmetrical thing where it's just as easy to fall back down. And if you ask, where does that come from? It doesn't come from physics. It doesn't come from selection. This is a property of mathematics, actually. It's a free gift for mathematics, that the properties of these kinds of things are such that they're aimed upwards in terms of intelligence and integrated agency.

Slide 15 of 48 · Watch at [22:30](#)

Larger Selves Emerge from an Excitable Medium of agential components: how many?




There but for the grace of electrical synapses go "I" rather than "We"

Beings are dynamic, self-telling, variably convincing stories of greater potential

Agential material: how many agents per mm³?

Where is my **border** from "environment"?
every cell is some other cell's environment


Issue of individuation in cognition:
split brain patients, dissociative disorders, etc.



So we looked at the molecular networks that are inside cells. Now we go to the cellular level, and we look at an early embryo. So here's a blastoderm, maybe a few 100,000 cells at this point. We look at that and we say there's an embryo. What is there one of? When we say there's one embryo, what is there one of? Well, what there's one of is a kind of collective model. All of the cells are connected and are aligned, both physically and physiologically. They're aligned toward the same journey in anatomical space. They're going to undertake actions that bring them closer to a specific species-specific target morphology.

Now, if you take a little needle and put some scratches into this blastoderm, each one of these little islands before they heal up is going to make its own embryo. And so that's how you get conjoined twins and triplets and so on. And so now some very interesting questions arise. The number of beings, the number of individuals inside this embryo is not known in advance. It is not fixed by the genetics. It is an outcome of the physiological events that take place. We can ask how many agents per cubic millimeter of embryonic tissue, like what's the carrying capacity? For computer science, we think we know how to calculate those kinds of things based on the number of logic gates you have and so on. But for biology, we really don't. And of course, you all deal with these kinds of things all the time, split-brain patients, dissociative disorders, how many personalities can you actually fit in a standard amount of brain? These are all open questions.


Life Has Nested Intelligence, not Merely Structure



And so the bodies of which we are made are a kind of multi-scale competency architecture. Each level of organization from the molecular networks on up has its own agendas. They solve problems in various spaces, and they have learning capacity. They have different ways of sensing and acting and so on at every level. This is one difference between how we currently make robotics and so on. We are a multi-scale intelligence that is collective at every step.

Life Has Embodiment Outside of Familiar 3D space:


3D Space (behavior)



perception-action loop
can happen in other
spaces!
-> unconventional
embodiment for
AI's


Transcriptional Space

Huang, S., Erbe, C., Kaufman, S., Simon, C. (2019). *Cell*, 175, 100-115




Morphospace

Gershenk, J., Levin, M., and Mello, S. (2021). *BioSystems*, 209, 104511



Physiological Space

Marder, E., & Goaillard, A. M. (2006). Variability, compensation and homeostasis in neuron and network function. *Neural Systems*, 317, 563-574



Now, I will also point out that what's really important to know is that these kinds of systems have embodiments way outside of familiar three-dimensional space. We are pretty good at recognizing intelligence as movement in 3D space, but biology has been solving problems long before nerve and muscle showed up. So cells have been navigating the space of gene expression, high-dimensional, difficult space, like 20,000 dimensions in some cases: anatomical space, physiological state space. So when you're looking at an organoid that's just sitting there not moving around and you say, well, it has no embodiment, or, in fact, a language model, you say it has no embodiment, maybe it doesn't have embodiment as movement in 3D space, but there are many other spaces in which agents do this perception, decision-making, action loop, active inference, all these things. They're doing this in all kinds of other spaces.

Symmetry between Behavioral Science and Developmental Biology

"Intelligence is the ability to reach the same goal by different means."
- William James

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


Cellular Swarms vs. Brains (or, what is neuroscience about?)

- 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)

Neuroscience - Detects **Translate from one domain to another** Developmental Biology

So what we've been studying is the symmetry between behavioral science and developmental biology. I like James's definition of intelligence, the ability to navigate spaces towards the same goals. And what we've been studying is the hypothesis that morphogenesis, the creation and repair of bodies, is a collective intelligence that exerts its behavioral competencies in a different space, in anatomical space. Whereas you and I are collective intelligences in the sense that our neurons and other cell types work together to help us navigate 3D space. There are some pivots that you can make around time scale and so on. And actually the symmetry is incredibly strong. And we've actually made tools where you can paste in a neuroscience paper and it swaps the word milliseconds for minutes and neurons for the word cells. And then you have yourself a developmental biology paper. So there are really, really strong similarities. And that's because, as I'll point out in a minute, the underlying evolutionary origin is common. These are not really, really distinct things that I'm trying to pull together. They actually have the same underlying molecular mechanisms.


Top-Down Control: how does mental intent make the chemistry dance?



COGNITION AND MUSCLE MOTION



Tim Flach



Tail grafted onto flank

Time →

Franziska Fortmeyer, Experiment, 1956 (11)

Anatomical homeostasis:

local order obeys global plan

So I'll show you a couple of interesting examples. I could talk for an hour just on that subject alone. But let's talk about the kind of top-down control. So here's an experiment. We take an amphibian like this, and this is not our work. This is back from the 50s. They would graft a tail onto the side here, the flank, and this thing slowly transforms into a limb. Why? Because the system as a whole has a stored representation, as it turns out, of what a correct system is supposed to look like. And when there are errors, it tries to repair them and to move towards the right location. Now look, the cells at the tip of this tail are tail tip cells, where they belong. There's nothing wrong with them. They don't have any damage. There's no injury. And yet they turn into fingers. Why is this happening? Because the system is primed for a transduction from a large-scale, highly abstract piece of information about what the layout of a typical amphibian like that looks like. That has to filter down to the molecular events of cells. Cells don't have any idea what a limb is or how many fingers you're supposed to have. The collective does, but the individual cells have to be instructed as the system transduces down into the chemical signals that actually are required to make this happen. Where have we seen this before? Well, basically, voluntary motion. When you get up in the morning and you have these incredibly abstract goals around, financial goals, research goals, whatever, in order for you to actually move on, to act on any of those goals, the chemistry has to change. Ions have to cross your muscle membranes for this to happen. And our bodies are basically an amazing system for transducing very abstract kinds of goals in very high-level spaces

down into actually making the chemistry and physics act in accordance, and of course, vice versa.

Slide 20 of 48 · Watch at [28:33](#)

Genome and Beginner's Mind



newt kidney tubule cross-section

Fankhauser, 1945, J. Exp. Zool., 59(3): 445-455

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

you can't even count on your parts!

- Creative, intelligent problem-solving - repurpose available tools to new circumstances

INTERFACE

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

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

Integrative Biology

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


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

So how do they know? they don't "really" know, do they?

And then we have these creative problem-solving capacities where if you take a normal newt that has 8 to 10 cells that normally make up its kidney tubules, we can artificially make newts where the cells are gigantic, and so Fankhauser did this years ago. If you do that, the cells automatically adjust, fewer cells, but the exact same large-scale structure. If you keep pushing it and make truly enormous cells, just one cell will bend around itself, leaving a space in the middle, giving you the exact same structure. So two things happening here. One is the ability to call up different molecular mechanisms. This is cell-to-cell communication. This is cytoskeletal bending. Call up different mechanisms to solve a problem, right? They have those on every IQ test. Here's a set of objects, use them to solve this problem. Here are the molecular affordances you have. Use them to solve a problem you've never seen before. So creative problem solving in anatomical space. And of course, because you can't even count on your parts as a new embryo coming into the world, not only do you not know what the environment is like, you don't know how many copies of your genome you're

going to have, you don't know what size your cells are going to be. You have to get the job done under these kind of very unreliable medium. Biology is an unreliable medium, and it drives a feedback loop of creative intelligence that was there long before neurons appeared.

Slide 21 of 48 · Watch at [30:01](#)



And so, what I've shown you is these ideas about top-down control and problem solving, and one other thing to sort of cement together the cognitive science and morphogenesis and evolution is the following. None of us have access to our past. What you have access to are the memory engrams that past versions of you have created based on experiences that they had. So basically the memories that you have are messages. They are communications from a past self that has left a specific, highly compressed kind of representation somewhere in your brain or body. And then, so that's what happens in the past. Here's our now moment. Every cognitive system is in charge of interpreting these steps into this information into the future. Now, because this compression process and the compression and generalization inevitably throws away information, you don't have all the details, nor could you function if you had all the details. But what it means is that this is a creative process. As you know, you guys

know better than anybody, recall of memories has a lot of reconstruction associated with it. Your job at any given moment is to figure out what your own memories mean. And you don't have necessarily an allegiance to how your past self interpreted. You are trying to tell the most adaptive story going forward into the future. So this is a creative improvisational process.



Slide 22 of 48 · Watch at [31:33](#)

Genomic Memories are Also Improvised Upon

CellPress **Trends in Genetics**

Opinion
What does evolution make? Learning in living lineages and machines

Benedikt Hart ^{1,2,*} and Michael Levin ^{1,2,3,*}





Biology assumes the hardware is unreliable
 Environment and your own parts will change, **don't over-train**
 Little allegiance to past Self's meaning of engrams
 Re-interpret on-the-fly - present/future is all that matters
 Engram is highly compressed - creative remembering, not deduction
 Unreliable hardware and improvisation are a feature, not a bug


And that is exactly what evolution does. So you get handed a set of genetic affordances, which are compressed representations of the experience of your ancestors. But now, as an embryo undergoing morphogenesis, your goal is to reconstruct the most adaptive meaning, not what it used to mean, because you have actually no idea what it used to mean. What you have to do now is to use them as best as you can to construct a good story going forward. And this is why the plasticity is such, all these examples, this is why this thing works, the eye on the tail, and why these guys can do it, and many other kinds of examples of amazing plasticity, because the whole process of building a body is not the thing that we were told for decades, which is you have the DNA and the DNA tells you what you're going to be. That isn't it at all. The genome is basically a set of engrams or prompts, and it is an active, intelligent process by groups of cells to

interpret it and to do something coherent. And I'm going to show you some very unusual things that they do in a moment.


Slide 23 of 48 · Watch at [32:39](#)



So how does all this work? How does all this intelligence work? Well, in the brain, we think what's happening is that you have the hardware. I don't have to tell this group what the hardware looks like. It's an electrical network employing both chemical and electrical synapses. And what it allows that network to do is information processing that can do things like move you through three-dimensional space. And then people who work on neural decoding will try to read these electrophysiological states and infer what the subject was thinking about, memories, goals, that kind of stuff. So it turns out that this system here is incredibly ancient.




It goes back to the time of bacterial biofilms. This is what evolution co-opted and sped up to make nervous systems. Every cell in your body has ion channels. Most cells in your body have electrical synapses, known as gap junctions, to their neighboring cells. They make electrical networks. And basically, we can take all the tools of neuroscience and behavioral science, and if you just make a few pivots, if you're willing to talk about morphospace instead of 3D space, and you slow the scale down to minutes and hours instead of milliseconds, everything works, because that is basically what happened. All of the same kinds of mechanisms, the ability to create voltage potentials, to propagate them through paths, through tissue, to use them to compute, to integrate information, to store memories and counterfactuals and so on, all of this is extremely ancient.



We developed the first tools to read and write these electrical pattern memories, because we have some idea of what the electrical networks in your brain like to think about. What did these electrical networks think about before there were neurons and muscles able to move you around? Well, it turns out that what they thought about was shape. That is what these networks originally were for, to process movement through anatomical space, to move your configuration from a single cell to that of a human or whatever else. That is the navigation that they used to think about. And so we developed some tools using voltage-sensitive fluorescent dyes here, for example, to see all of the electrical conversations that these cells are having in an early frog embryo. You can track it at the individual cell level here. We do a lot of computational modeling to do a multi-scale model all the way down to the transcription of the channel genes and all the way up to network properties such as in-painting, out-painting, the kinds of things that you see in regeneration and so on. And so now for the next few minutes, I want to show you what does this give us? This figuring out this symmetry between the algorithms that are carried out by cognitive systems and the ones that build your body in addition to the ones that build your mind. What does that give us?


Communicating with Somatic Intelligence: prompt = make an eye



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

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

Developmental Modules
because morphogenetic goals



voltage-modified cells


wild-type cells recruited to ectopic lens!

Gertly Images


Well, one of the things we can do is we can start to communicate with this somatic intelligence, the problem-solving system that I've been showing you up till now that has this incredible plasticity and the ability to solve various problems. What we can do now that we understand the medium in which it records its goals, it's basically the same medium as in the brain, so these electrophysiological networks. Once we can read and write information there, we can use that as an interface to communicate novel goals, not to micromanage them, not to, the most interesting thing I think about neuroscience is that it helps us to understand this kind of multiscale phenomenon where when I'm talking to you now, I don't need to worry about what your synaptic proteins are doing. You will take care of all of that. I'm giving you high-level prompts through a thin language interface, and your system will then transduce this down to the chemistry that's needed to make things happen. That's the reason humans were able to train dogs and horses for thousands of years before we knew any neuroscience at all. But this works all the way down. And so we can make a very simple prompt to a bunch of cells here in the early embryo and say, make an eye. Okay, we don't have to tell it how to make an eye. We have no idea how to make an eye. But what we do know is a little bit of the electrical language. And this is what we're trying to do, is to crack this code such that we can inject ion channels that induce a particular voltage state in these cells. And we now know that the cells interpret that particular voltage pattern. It's a multicellular voltage pattern that the cells interpret as, oh, we should build an eye here. And they build an eye. Here's one sitting in the gut. All lens, retina, optic nerve, all the stuff it's supposed to have. If there's too few of them, they actually, so the blue

cells are the ones we injected with the channels. They will actually recruit their neighbors to help them build the eye. They know there's not enough of them, so they recruit their neighbors. Here's ants and termites and things do exactly the same thing. This is a common feature of collective intelligences.

Slide 27 of 48 · Watch at [37:38](#)



What happens then is that you can induce these ectopic foci, but there's a battle that goes on because the surrounding cells, they know perfectly well there should not be an eye there. And as a cancer suppression mechanism, they try to remove, they try to normalize the voltage. So you may not get three or four eyes here, you may only get one or none, because there's a battle for the future of the morphogenetic journey. Some of them say go to the eye fate, some of them say no, you should stay skin or whatever you are. And so we can actually watch these conversations take place.




In those flatworms that I showed you, if you cut them into pieces, every piece makes one head and one tail. And you might think that it's the genetics that tells every piece how many heads to have. But of course, remember that the genome doesn't say anything about heads. The genome, if you actually read it, refers to proteins. And so we still have this question about how does the piece know how many heads it's supposed to have. Turns out there's a voltage gradient that you can read. And the voltage gradient says one head, one tail. And if we rewrite that gradient and say, no, you should have two heads, and then after the fact, cut the animal. So here's a perfectly normal body with a perfectly normal molecular marker expression, but it has an aberrant memory of what a correct planarian is supposed to look like. Now that memory is latent, doesn't do anything. It just sits there until we cut. And at that point, all the cells consult that memory and they say, two heads, fine. And they build this two-headed worm. This is not AI or Photoshop or anything like that. These are actual animals. A single body can hold up to, probably more, but at least two different representations of what a planarian should look like. This, I think, is the origin of mental time travel that we can all do and counterfactual thought, because this creature is holding a pattern that is not correct right now. I only have one head right now, but my memory of what to do if I get injured in the future, potentially, and so that's the counterfactual part, this is what I would build in the future.

Same Genome, Different Outcome




2 brains in 1 body


And that's in fact what they do. So lots of interesting neuroscience to be done on these animals with two heads. Again, reminder, we don't touch the genome. This is not about the hardware whatsoever. The planarian genome builds a system that by default makes a single head, but it is highly reprogrammable, that is not nailed down.



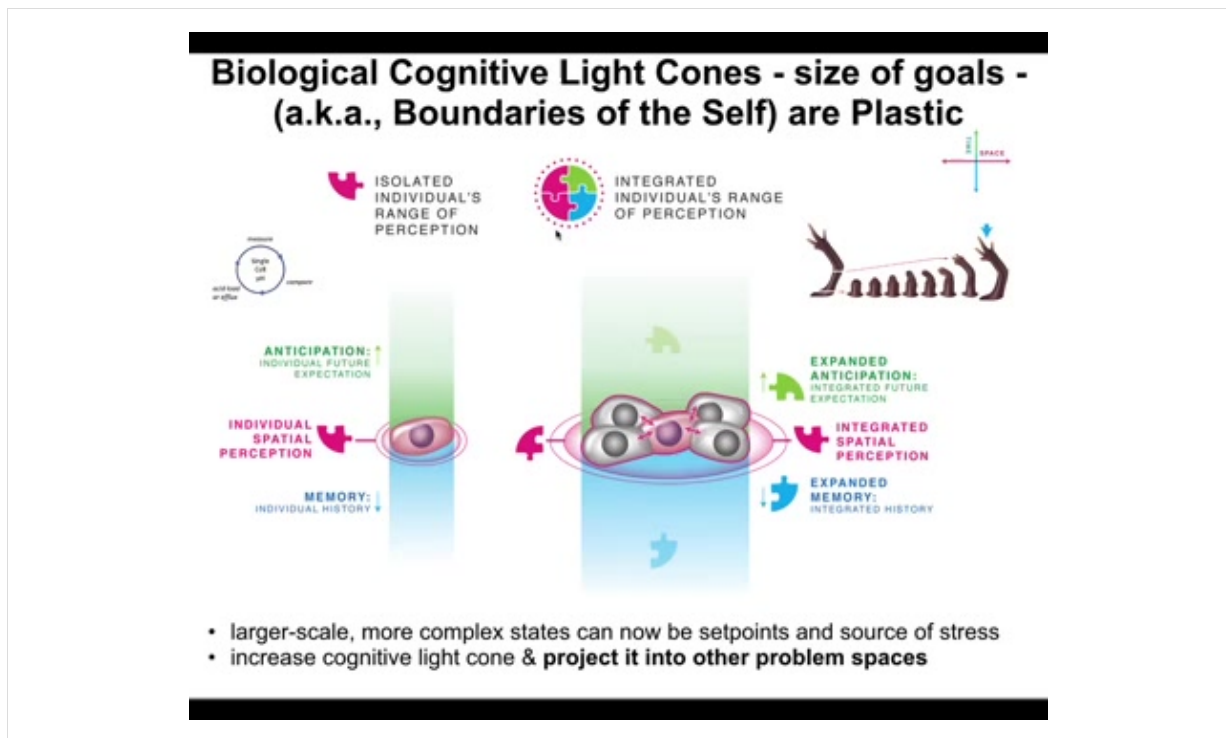
You can also ask them to make heads of other species. So here's a *Drosophila* with a nice triangular head. They can make worms of flatheads, round heads, and so on. Not just the head shape, but the shape of the brain and the distribution of stem cells, just like these other species, 100 to 150 million years of evolutionary distance. No genetic change. The same hardware is perfectly happy to visit the attractors in morphospace that are normally occupied by these other systems.




We can use the same thing to repair. I won't go into the details, but basically this is our program in birth defect repair. So if you do a dominant mutation in a gene called Notch, you probably all know what this is, very important neurogenesis gene, the brain is completely wrecked. No forebrain, midbrain, and hindbrain are a bubble, no behavior, profound damage. With a very simple modification of an HCN2 ion channel, we get back to a correct bioelectric pattern for the brain, and then the cells are built to suit. So you get a normal brain, normal gene expression, normal IQs, their learning rates, which we measure, go back to those of controls, even though they have this profound dominant Notch mutation. So at least some hardware defects can be fixed quote unquote in software at the level, not at the level of the genetics, but at the level of the physiology that drives.



We have a similar program in limb regeneration where we can take frogs, which normally don't regenerate their legs as adults, and give a simple 24-hour stimulus that then leads to a year and a half of leg growth. Just a very simple trigger at the very beginning, say go down the leg-building path, not the scarring path, and then by 45 days you've already got some toes and a toenail and then this thing grows for a year and a half. So these are the consequences of using this, of pursuing this hypothesis that this kind of multiscale problem solving and communication with these systems is not just for brains in 3D space, but is actually fundamentally baked into biology all the way to the bottom. And we can now take advantage of this.




And the final example I'm going to show you has to do with a cancer problem. And what happens in evolution and development is that you start with single cells, which have extremely tiny goals. So a single cell is only concerned about managing its own pH, hunger level, those kinds of things, right? So very tiny cognitive light cone, very tiny scalar kinds of goals. When you make a network, this network is able to store much larger set points for its homeostatic activity. So here's an amphibian limb. If you amputate anywhere along here, the cells very quickly go back and they grow what's needed, then they stop. Why do they stop? Because this is a homeostatic process. They note the error, they try to reduce the error, and when the error is within acceptable limits, then they stop. So I've just been showing you in the case of the worm and so on, how do they know what the pattern is? It's stored bioelectrically. There's a bioelectrical pre-pattern that tells them what the set point for this homeostatic activity is. But note that collections, collectives of cells can store enormous grandiose goals. Their cognitive light cones are huge. They can, the collective can store a memory of something like this, this large-scale shape. But that system, of course, has a failure mode, and that failure mode is cancer.




Cancer cells aren't more selfish than normal cells. They just have smaller cells. Because what happens when a single cell disconnects from this electrical collective, it no longer has access to this enormous memory of what they were building. Now the border between it and the outside world shrinks, right? It's back to being a unicellular organism as far as it's concerned. The rest of the body is just external environment. And it pursues its own tiny little goals, which are eat as much as you want, reproduce as much as you can, go where life is good, and that's a metastasis.

So we've been using this idea of cancer as basically a somatic dissociative identity disorder, right? All of these things together have the identity that they're building a limb and they have ways of doing it, but individual cells, once they disconnect, have no access to that. So there's a diagnostic here. You can use these voltage dyes to see when it's about to happen, because what we do is we inject human oncogenes into these tadpoles. Eventually they make tumors, but you can see when the cells are about to disconnect. And then, most importantly, when you reconnect them, we don't fix the oncogene, we don't kill the cells, this isn't chemotherapy, we leave the cells alone, but we reconnect them, we forcibly reconnect them to the electrical network, then even though the oncoprotein is blazingly expressed here, there's no tumor, because it's not the genetics that drives, it's the physiology that drives. And as long as the network is thinking about making large muscle and large-scale structures like muscle and skin and so on, then they're kept away from these kinds of single-cell, kind of ancient unicellular behaviors.



So the bottom line to all of this is that I really think that the control of growth and form for regenerative therapeutics, these are issues that affect most of the problems we have in medicine, complementing the bottom-up kind of stuff that everybody's been focused on for hundreds of years is this kind of top-down control that's taken straight from concepts in behavioral and cognitive sciences. And I actually think that future medicine is going to look a lot more like a kind of somatic psychiatry than it's going to look like chemistry. It's not going to be this kind of bottom-up intervention that we try to do with drugs today. And the bioelectric layer is a major, it's not the only interface, but it's a massively useful interface.

Biomedical Endgame: Anatomical Compiler



It's a communication device, not a 3D printer

Problems of memory, allostasis, incorrect priors, decision-making, perception

Why we need it:

- Birth defects
- Traumatic injury
- Cancer
- Aging
- Degenerative disease

What's left?


So just the last couple of things I want to point out. First of all, we've been talking about this, addressing things like birth defects, traumatic injury, cancer, aging. And what we want is complete control. We want an anatomical compiler, right? All of these things would go away if we were able to communicate, so this is not a 3D printer of some sort, it's a way to communicate your goals as a bioengineer to the goals of the cellular collective, right? So maybe you draw this like three-headed flatworm, and that is what the compiler will do, is give you the exact stimuli that have to be given to cells to get them to build whatever you want them to build. So we are very far away from a complete version of this, but we've taken the first steps. So that's fine. We can rebuild, the near future will be to rebuild the organs, standard structures that belong in your body. What's left? What goes beyond that?

Overview:

- Introduction: why a talk from a computer scientist and developmental biologist?
- TAME: a framework for diverse intelligence research and applications
- Morphogenesis: a model system for communicating with semi-alien minds
- The future of the sciences of ~~the mind~~ minds


where do mental architectures come from?
what will be the goals and preferences of novel beings?
(and what does this say about our own?)

And so here's the last part, just for a couple of minutes. What's left is this interesting question of where does your mental architecture come from? Well, mostly people would say evolution, right? It's been shaped by your evolutionary history. But what are going to be the goals, preferences, and capabilities of entirely novel beings that have never been here before, that don't have an evolutionary history that selected them for specific cognitive properties? What is the content of their mind going to be like?



And so I want to show you two quick examples. The first we call xenobots. These are made of frog cells. We dissociate some epithelial cells from a frog embryo. This is just skin. There's no neurons, there's nothing. What they do is they combine into this little thing that swims.


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



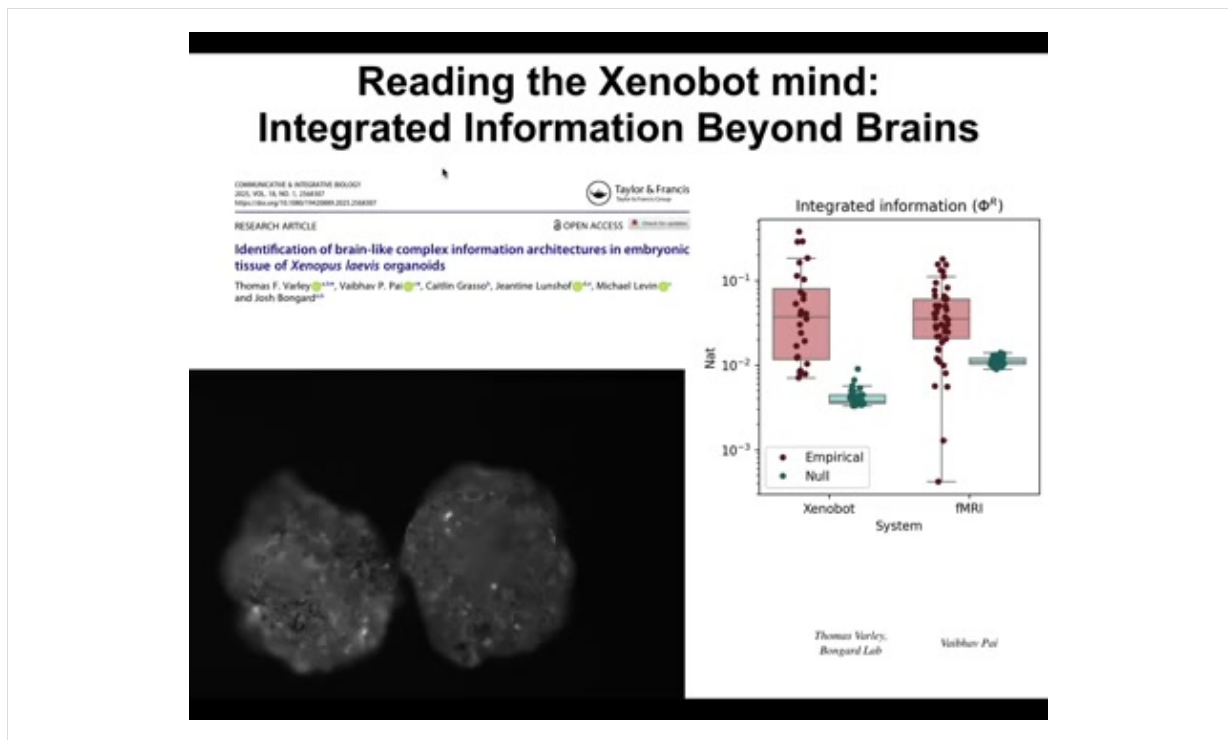
- 1) it traverses maze,
- 2) rounds the corners without bumping into walls, and
- 3) it makes a spontaneous decision to turn around without hitting anything.

Douglas Blackiston

It uses the cilia on the epithelial cells to actually coordinate swimming activity. Here it's running this, traversing this kind of maze. It turns around whenever it wants to.




It has these amazing properties that if you give them loose epithelial cells, it will collect them into little piles and compact them. And guess what the little piles turn into? They turn into the next generation of xenobot, which then goes on to do the same thing and the same thing. So this is kinematic replication. We've made it impossible for them to reproduce in the normal fashion. They figured out another way of doing it. No other creature on earth that we know of does kinematic self-replication. They make copies of themselves from materials they find in the environment.



If you track the calcium signaling, remember there's no neurons here, this is just epithelial cells. If you track this thing and you apply some metrics, so this is Thomas Farley's work from Josh Bongard's lab, basically compared to null models, they do about as well as we see in fMRI. So this is for brains. So I'm not saying that these are brain-like. What I'm saying is that both brains and these kinds of simple systems are exploiting dynamics that are incredibly fundamental. And we can catch this with metrics like causal emergence metrics. Now, that might seem weird and maybe specific to amphibians.

What would *your* cells do if liberated?



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?

Genetics & psychiatry

Glynn Gammazza

I would ask you, what would your cells do if I liberated them from your body? So you can look at something like this. This looks like we got this off the bottom of a pond somewhere. Actually, if you were to sequence it, 100% *Homo sapiens* genome, completely normal human genome. These are made from adult cells taken from patients that go in for tracheal biopsies. We buy the cells, we let them reboot their multicellularity, and they make these amazing little creatures known as anthroblasts. Anthroblasts have over 9,000 different gene expressions than they had when they were sitting in your trachea. Remember, we don't touch the genome, no synthetic biology, no scaffolds, no weird nanomaterials. These are just new ways, new lifestyles that your cells can adopt given the opportunity.



One of the things they like to do is heal neural wounds. So, for example, if we make a scratch here, throw a bunch of human neurons in the dish, these guys get together and over the next four days they will heal. They'll take the two edges of the wound and actually heal it together. We did not teach them to do that. This is not anything that's in their history. And so this is just the first thing we looked for.

Evolution was Supposed to Explain Complexity, with High Specificity for Selection history...

The diagram illustrates the evolutionary path from the *Xenopus laevis* genome to embryos and then to Xenobots. It shows a genome map, a cell division diagram, and images of embryos and Xenobots. Labels include 'Xenopus laevis genome', 'Path A: embryos', 'Path B: Xenobots', 'Developmental Time', and 'Behavior'.

Xenobot bodies and minds have no straightforward evolutionary back story;

When was the computational cost of Xenobot features paid?!

Whence the specificity of evolutionary explanations?

And so now we have a simple question with a very difficult answer, which is, we know when we paid the computational cost to make a good frog and to have these nice developmental stages and some tadpoles with behaviors, all the years, the eons of this genome bashing against the environment, that's when we pay that computational cost. When did we pay the computational cost to make a good Zenobot or a good Anthrobot? There's never been any Zenobots or Anthrobots. There's never been selection for those features. And we have to think really carefully about what it means if we were going to say that, well, at the same time that you selected for frogs, you also got kinematic replication and hearing because actually Zenobots can hear in ways that tadpoles don't and things like that. Yeah, you just sort of got them at the same time. That basically rips up the whole point of specificity between the history of environments and selection and what we get at the end, right? That was kind of the whole point of evolutionary theory. So that kind of plasticity and the ability to just out of the box have these capacities really asks us to think about where do they come from in the first place. Selection is only part of the story. We need to be able to understand them because we make novel things all the time. Internet of Things, swarm robotics, social and financial structures, we have no idea how to predict.


Patterns Come From Genetics, Environment, and ??

$z = z^3 + 7$

<https://thoughtforms.life/halley's-method-fractal-art/>

No aspect of physics or history is responsible;
Could such patterns in-form physiology, behavior?

And so when we think about where do these things come from, as biologists, we like to think that they come from genetics and environment. But I just want to point out that there is another source of information. So this pattern, for example, is what happens when you plot, when you do a Halley plot of this tiny little equation, and complex numbers. All of this is hidden in there. Where is this specified? Like where does this come from? There is no aspect of physics, there is no aspect of history or environment that tells you that this has the specificity of exactly this pattern. Where does it come from? And so I've actually been working on a model in which these kinds of truths of mathematics are not just relevant to mathematical objects, but actually to patterns of physiology, patterns of behavior, aka kinds of minds, so we can talk about that.



So, the bottom line is this: pretty much any combination of evolved cellular material, engineered material, and software is some kind of viable agent because of the plasticity of life. This tiny little corner of the space is everything that includes what Darwin called the endless forms most beautiful. Actually, the space is enormous. Cyborgs and hybrot and combinatorial explosion of different kinds of embodied minds. And I really do think we have to work very hard to be able to enter some kind of ethical symbiosis with the beings with which we're going to be sharing our world very, very, very soon.

Summary:

- Intelligence is much older, much broader, than brains
- Morphogenesis = model system for semi-alien intelligence
- Agential material of life - it's behavioral science at the bottom
- Conserved mechanisms and algorithms -> behavioral science tools are helping us communicate with cell groups toward transformative applications in regenerative medicine.
- Conversely, diverse intelligence research helps us understand the
 - Mental health diseases of novel beings
 - Patients with alien and highly diverse Umwelts, sensory motor capabilities, connections, IQs, goals, opportunities
 - Dream analysis, collective unconscious, and symbology of beings who do not share our evolutionary history
 - Novel sources of mental architectures and trauma
 - Testifying in court cases re. enhanced capacity and novel challenges to the free will question
 - Beyond disease: a good and meaningful alien life
- Possible topics for discussion:
 - Thoughts are thinkers continuum (patterns as agents)
 - Memory, genetic and behavioral, as improvisation

And so I'll just summarize here. What I've told you today is that intelligence, aka problem solving in different spaces, is much older and much broader than brains. Morphogenesis is a model system for beginning to communicate with these alien intelligences, and the applications in regenerative medicine are basically the outcomes that are telling us we're on the right track. I think that this standard picture where it's physics and then chemistry and then behavior science is somewhere at the top. I think it's upside down. I think for many reasons, it's basically behavior science all the way down and certainly by the time you have molecular networks, all of your behavior textbooks are relevant. And what's coming is that the field of diverse intelligence research is really going to inform the fact that there will be mental health diseases of novel beings that have completely different life experiences. You'll have patients with highly diverse umwelts, with new sensory motor capabilities, with connections to others. We're going to have to figure out what we do. For example, if you're analyzing dreams, what do the dreams of novel beings look like? What do they mean? How do we help beings that have never been here before, that don't share the same evolutionary stream? How do we help them have a good and meaningful life? So I will just stop here.

Thank you to:

Post-docs and staff scientists in the Levin lab:
Nestor Oviedo, Junji Morokuma - planarian pattern memory rewriting
Douglas Blackiston - brain-body interface plasticity, synthetic living biobots
Vaibhav Pai - voltage gradients in eye/brain induction and repair
Federico Pigozzi - causal emergence in minimal models
Patrick McMillen - bioelectric imaging and embryogenesis

Graduate Students:
Gizem Gumuskaya, Nikolay Davey - Anthrobots
Sherry Aw - bioelectric eye induction
Fallon Durant - planarian pattern memories
Emily Ertle - free lunch controllers
Adam Goldstein, Taining Zhang - emergent competencies of algorithms
Angela Tung - inter-embryo communication

Undergraduate Students:
Pranjal Srivastava, Ben G. Cooper, Hannah Lesser, Ben Semegran - Anthrobots.
Maya Emmons-Bell - planarian head shifts to other species

Technical support:
Erin Switzer - vertebrate animal husbandry


Collaborators: Allen Center members +
Joshua Bongard - Xenobot simulations and AI
Thomas Pollak, Thomas Doctor, Olaf Witkowski, Bill Duane, Paul Colognese - Buddhist models of AI
Sebastian Risi - open-ended evolution
Simon Garnier - computational analysis of Anthrobot form and function
Chris Fields - physics of sentience and sentience of physics
Richard Watson - computational models of cognitive scaling and evolutionary learning
Giovanni Pezzulo - cognitive science applied to morphogenesis
Matthias Scheutz - robotics of free lunch controllers
Mark Solms, Marsa Hickey - psychology and psychiatry of diverse intelligence

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

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

Illustrations: Jeremy Guay @ Peregrine Creative

Disclosures: MorphoCeuticals, Fauna Systems, Astonishing Labs, SoftMax



and thank the people who did the empirical work that I showed you today and our many, many collaborators. And here are some disclosures. These are companies that have licensed some of the IP that we've had. So thank you so much. I'll stop here.

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.