A Lecture Companion

”Self-constructing bodies, collective minds: the intersection of CS, cognitive bio, and philosophy”

Michael Levin

Recorded on November 28, 2024

About this document

This document is a companion to the recorded lecture *"Self-constructing bodies, collective minds: the intersection of CS, cognitive bio, and philosophy"*, recorded on November 28, 2024. You can watch the original lecture or listen in your favorite podcast feeds — all links are on the page here.

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

Lecture description

This is a 1 hour 12 minute talk presented at the 2nd Symposium on the Philosophy of Computing at the Universidad Nacional Autónoma de México. It covers the ways in which biological computation is different from most of today's computer architectures, and concludes with a different way to think about emergence of cognition (not just complexity or unpredictability) in a way that informs how we think about biological and artificial intelligences.

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.

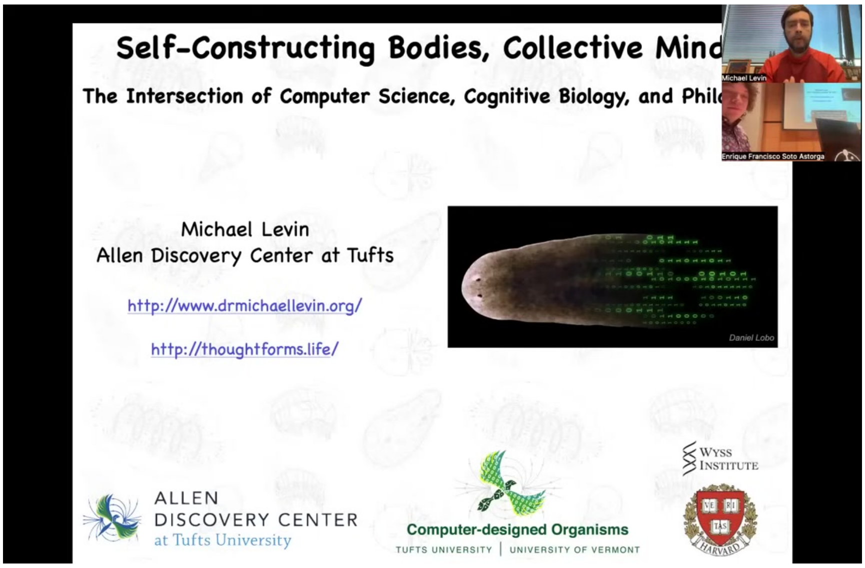

Thank you very much for having me here to speak to all of you. This is going to be a different talk than I usually give.

I normally talk about the applications in regenerative medicine of some of our ideas around biological computation and cellular learning. Today I'm going to give some exploratory ideas at the intersection of computer science, cognitive biology, and philosophy. All of the primary data, the software, the papers, everything is here if you want to follow up. This is my blog that has my personal thoughts on what some of this means.

Slide 2 of 54 · Watch at 0:40

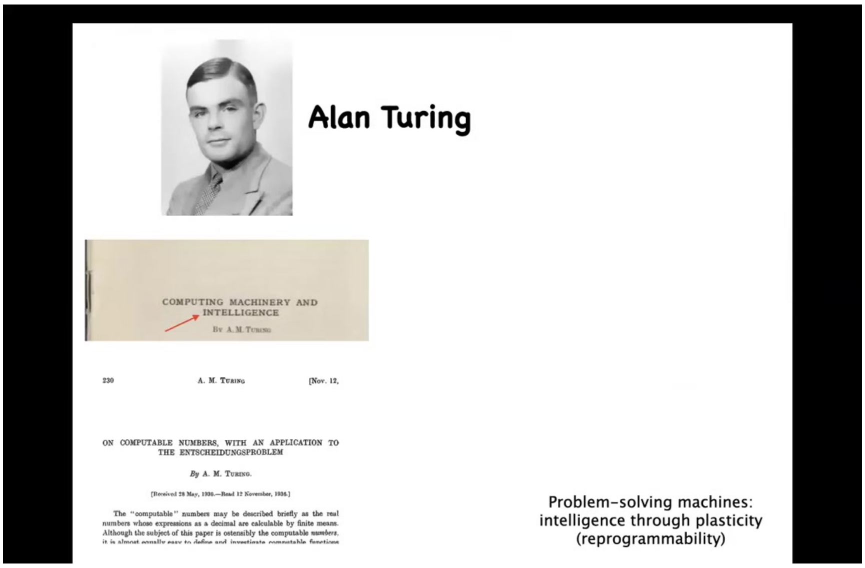

So Alan Turing needs no introduction. He was very interested in problem-solving machines and intelligence through plasticity and reprogrammability and computation and mathematics. But one of the more interesting things that he did was to write a paper on the chemical basis of morphogenesis, specifically trying to understand how chemicals may organize themselves during early embryonic development. Why would somebody who was all about intelligence and computation be thinking about chemicals in embryonic development? I think it's because he understood that there's a very profound symmetry between the self-creation of the body and the self-creation of the mind. This idea of problem solving in living machines, and the intelligence that we see in plasticity, I think is a very profound area of investigation. If he had lived, I think biology and medicine and also computer science would be much, much further along.

What I'm going to do today is talk about a few points. I will start out by looking under the hood of biology to show you some really unconventional aspects of living things that are not typically talked about. If we want to understand how to make truly bio-inspired engineering or computation, we need to understand this, going beyond neural architectures in particular. I'm going to then focus on trying to identify what is the same and what is different between biology and current computational approaches, because I think there are some commonalities, but biology does things that go very differently. Then, towards the end, I'm going to give some very speculative ideas about the space of unconventional beings and where their cognitive and other properties might actually come from and talk about the implications of that for our ethics.

Many people begin with this classic artwork of Adam naming the animals in the Garden of Eden. There's something we need to transcend here, something wrong, and then there's also something very profound.

The thing that's wrong here is that it gives this idea of discrete natural kinds, such that we know what all these different animals are. Adam is separate from them and even more different. As I'm going to show you during this talk, this idea is going to have to be revised.

But what is fundamentally deep here is that in this ancient story, it was Adam's job to name the animals. In particular, God couldn't do it, the angels couldn't do it. It was really up to Adam to name them, not just because he was going to have to live with them, but also because in those traditions, naming something means that you've discovered a very profound truth about their inner nature.

This was a call for humans to understand the true inner nature of the other beings with whom they share their world. We're going to have to go way beyond what Adam was called upon to do. I'll show you that.

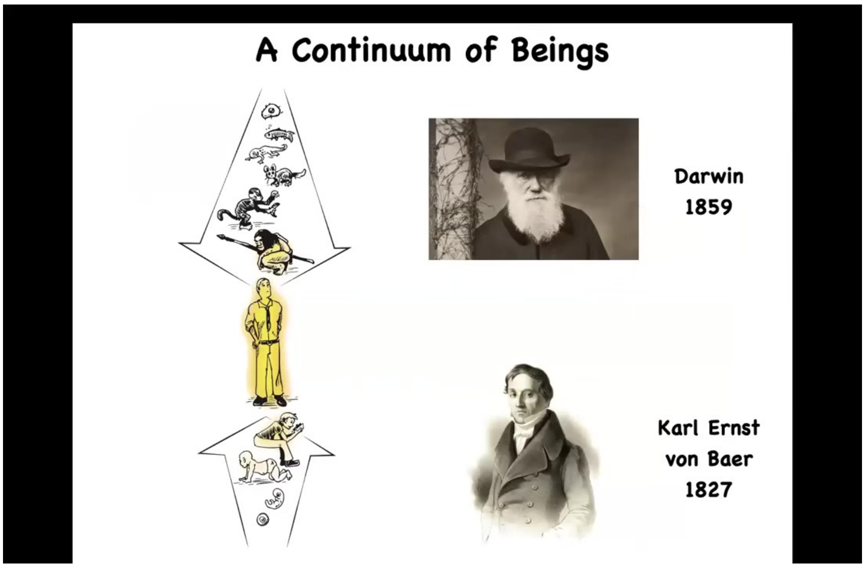

The first thing to realize is that, for example, in a lot of philosophy of mind, we talk about the human mind and the subject does this or that and what it feels like. But this adult modern human is really at the center of two very long and smooth continua.

During embryonic development we slowly develop from a single cell; in evolution the same thing happens. We really have to understand that whatever we apply to the human, whatever ideas about their cognition and their moral responsibilities and all these kinds of things, there is going to be some bleed over of these things down the developmental and evolutionary tree. We have to ask, what do those properties look like in these other beings?

It's important to start: the biology itself, developmental and evolutionary biology, tells you that these are all continuum problems, and we need to start to understand the scaling of it.

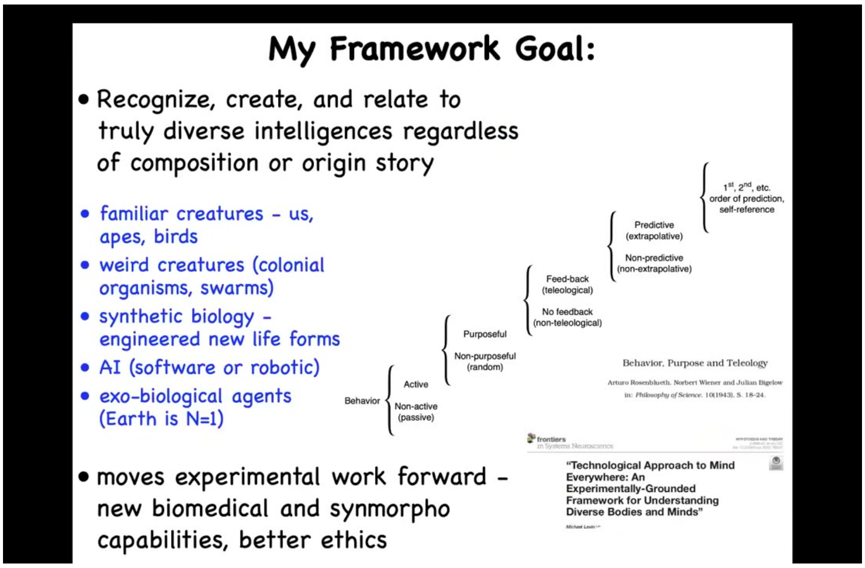

In my framework that people in my lab and I are working on, what we want is to be able to recognize, create, and relate to truly diverse intelligences, regardless of what they're made of or how they got here. We want to be able to think about familiar creatures like us, various other primates, birds, and so on, but also very unusual creatures such as colonial organisms and swarms, engineered new life forms, AIs, whether software or robotic, and someday even maybe alien organisms.

Of course, I'm not the first person to think about this. Here's, for example, Rosenbluth, Weiner, and Bigelow trying to make a ladder or a scale of how it is that you get from passive matter all the way up to human-level metacognition. My particular framework is described here in this paper. The goal of these frameworks is twofold. First, to move experimental work forward. It's not just philosophy. The idea is to have a framework that helps you make new discoveries and to actually enable ethics that are better grounded in the reality of what we now understand about bodies and minds.

We're going to do three things. I'm going to show you some unconventional biology, and in particular, the unique features of the biological substrate. I'm just going to give you a quick tour of a few unusual examples that you may or may not have seen before, and then we'll talk about why it works and what it means.

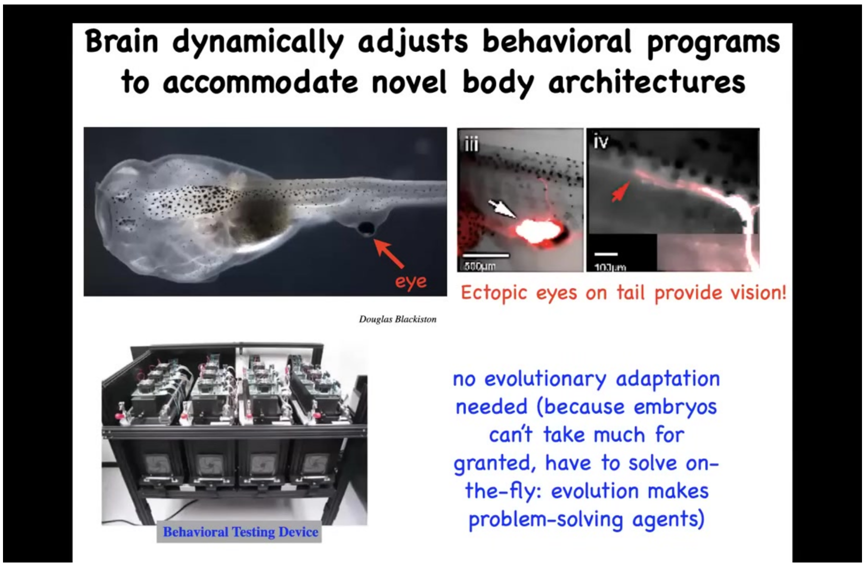

Here's one experiment. This is a tadpole. It's a larva of the frog, Xenopus laevis. Here's the mouth, the nostrils, the brain, the gut, and the tail. What you notice here is that there aren't any primary eyes. We've prevented those from forming, but we did put an eye on its tail. When you do this, the cells can successfully make an eye. This eye makes an optic nerve. That optic nerve does not go to the brain. It might synapse on the spinal cord or go to the gut, or in fact nowhere at all. Then we built this machine that automates the training and testing of these animals on visual cues. What we learned is that these animals can see perfectly well. Even though they have a radically different sensory motor architecture, the eye is now on its tail, not connected to the brain. They can still do visual tasks and they can see just fine. This is weird. Why do we not need additional rounds of evolutionary adaptation? Why does this work right out-of-the-box? I'm going to claim it's because embryos in general, the self-construction of the body and mind does not take priors too seriously. It doesn't over-train on priors. What they do instead is solve problems on the fly. What evolution makes is problem-solving agents, and this will become important shortly.

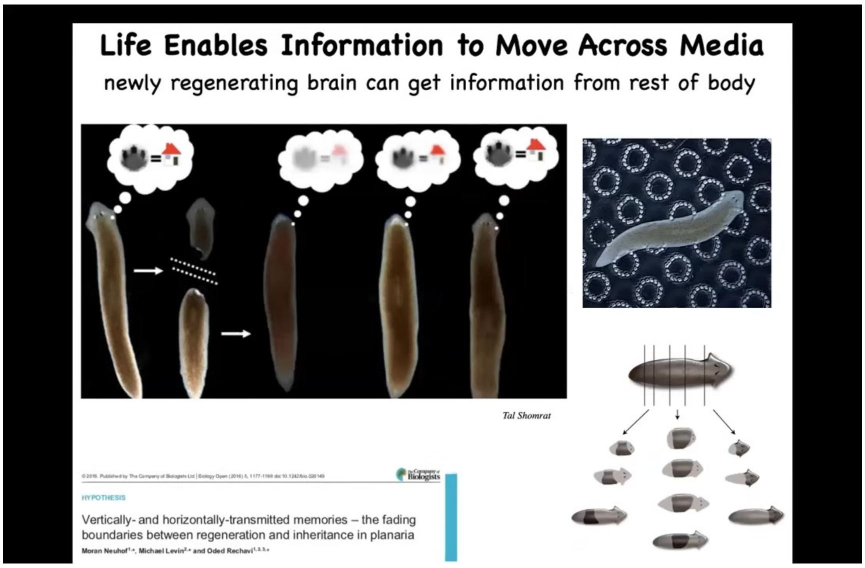

Another interesting phenomenon is this. These are planaria, flatworms. You can chop them up into pieces. Every piece regenerates a perfect little worm. So they have this anatomical memory. But they also have a behavioral memory. So you can train them to look for food on these little bumpy disks. If you do that, and then you cut off their head, which contains their centralized brain, the tail will sit there doing not much until it grows back a new brain. So you need a brain for behavior until you grow back a new brain. They don't do anything. Once it grows back a new brain, you find out that they remember the original information. It's called place conditioning. They remember that this is a place to eat. That means that, first of all, the information was stored somewhere else in the body. Second, it was imprinted onto the new brain as the brain regenerated. This material not only is plastic with respect to the behavior in novel configurations, not only does it remember how to rebuild the organism upon a very drastic injury, but it also is able to move information in and out of the brain and imprint it onto new tissue. That's remarkable.

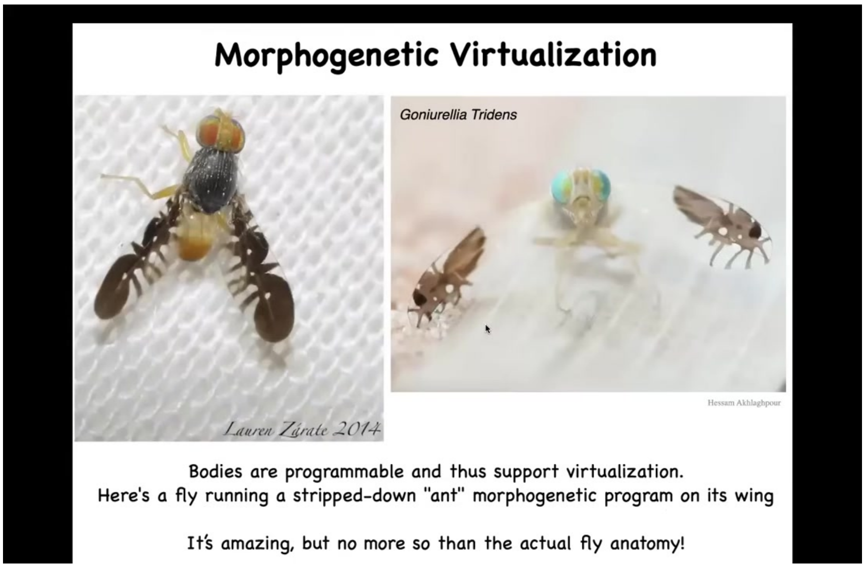

Here you see an interesting example of virtualization. This fly is running a very stripped-down, two-dimensional morphogenetic program of some ants. And it does this to scare off predators. It waves their wings and it looks like the ants are running around and then the predators don't want to deal with the ants, so they keep away from it. It has the remarkable three-dimensional morphogenetic program here, but it's able to run a very simplified virtual one on its wings for a completely different creature.

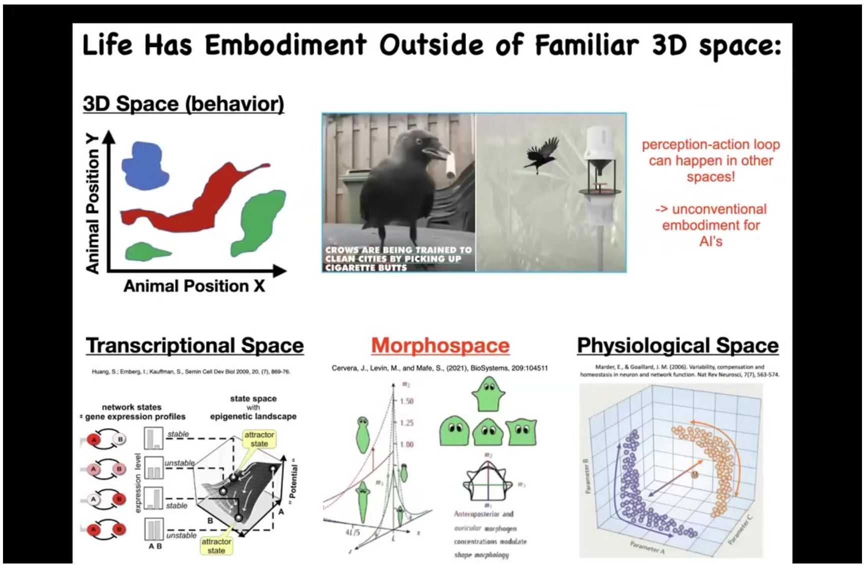

What we're seeing in many of these examples is this idea that we need to start to understand problem-solving competencies. If we're going to have a system that self-assembles in novel ways and deals with novel circumstances and potentially supports virtualization, and as I'll show you momentarily, is reprogrammable and so on. We need to be able to understand the intelligence that these systems have.

We as humans are pretty good at noticing intelligence of medium-sized objects and medium speeds in the three-dimensional world. You've got crows, dogs, and octopuses. But biology also operates in many other spaces. For example, the space of possible gene expression, the space of possible anatomical states, and the space of possible physiological states. There are many others.

This means that because we as human observers are not very good at noticing the perception-action loop in these spaces, if we could directly sense our blood chemistry, I think we would have no trouble seeing our liver and our kidneys as these kinds of autonomous symbionts that navigate space all day to keep us healthy. Because we're not very good at noticing this, we have to be careful when we talk about AIs having or not having an embodiment. People talk about purely software versus embodied agents like robots. The three-dimensional space in which conventional robots move is not the only space that matters.

There are other spaces in which intelligences can operate. That's largely what we study in my group, and I won't have time to show you too much of that. Cells and tissues, in fact, do navigate these other spaces and do so with great competency.

The other thing to keep in mind is that when we interact with systems, for example, language models, the system itself may be feed-forward, like a traditional artificial neural net. With the human in the loop and the interaction that goes back and forth, there is a cycle. The cycle as a whole is not purely feed-forward because depending on what the network does, the human will give different inputs.

We can even think of extremely unconventional agents such as this one, where you've got biologicals being read out by various types of AI systems implemented by robots, with humans in the loop reacting to the movement of the robots. This whole thing gets back and does rewards and punishments for the living tissue; you can concoct very, very unusual composite agents.

This plasticity goes beyond just the plasticity of the biology. We have to now understand what the plasticity looks like, both of bodies and of minds, in these hybrid scenarios.

Slide 10 of 54 · Watch at [12:20]

Let's talk about one specific kind of navigation that biological systems do, which is the space of anatomical, possible anatomical outcomes. This is known as anatomical morphospace. The production of a normal human body from an egg is extremely reliable. I'm going to claim that this process is intelligent. It's intelligent not because it's reliable or because it increases in complexity. Both of those things are easy to get. There's something else going on here, which is the creative problem solving. This is important for understanding some points about the philosophy of computation and biology and computationalism. This is a new kind of active matter. In fact, it's an agential material that has problem-solving competencies. What are those? First, for example, if you split early embryos into pieces, you don't get half bodies. You get perfectly normal monozygotic twins, triplets. From different starting configurations and starting positions in this anatomical morphospace, which is quite high dimensional, but I've cut it down to two, this system has the ability to get where it's going, this ensemble of goal states, despite various starting positions and despite various local maxima. We know we can get to where it's going despite different starting positions.

Slide 11 of 54 · Watch at 13:50

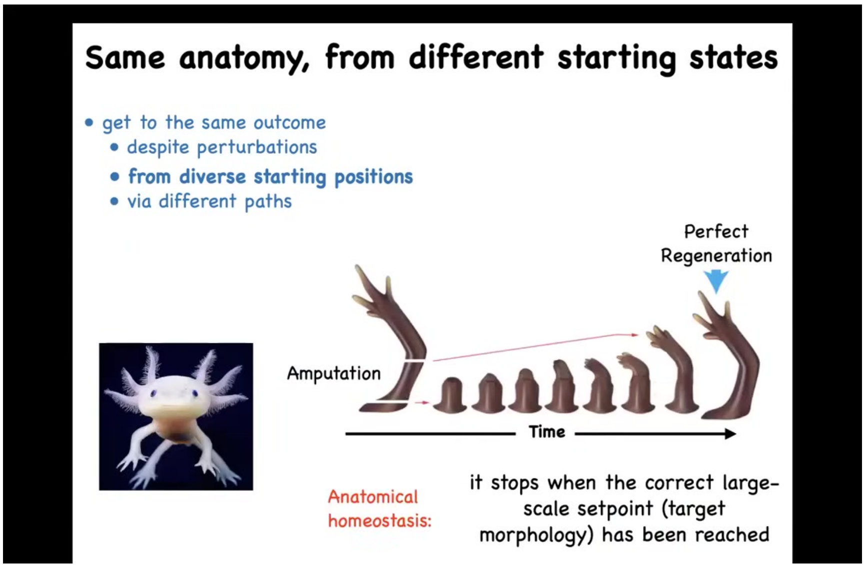

Some animals can do this throughout their lifespan. This little guy is an axolotl. They regenerate their limbs, their eyes, their jaws, their spinal cords, their ovaries, portions of the brain and heart.

When they inevitably bite each other's legs off, anywhere that it happens, the system will very rapidly grow and rebuild the limb. Then the most amazing thing happens: it stops. Why does it stop? It stops when the correct structure has been formed.

What you're really looking at is a kind of anatomical homeostasis. The system can sense when it's been deviated from the correct region of anatomical space. It undergoes rapid action to try to navigate back to that position. When the error is below some tolerance level, it stops.

Slide 12 of 54 · Watch at [14:40]

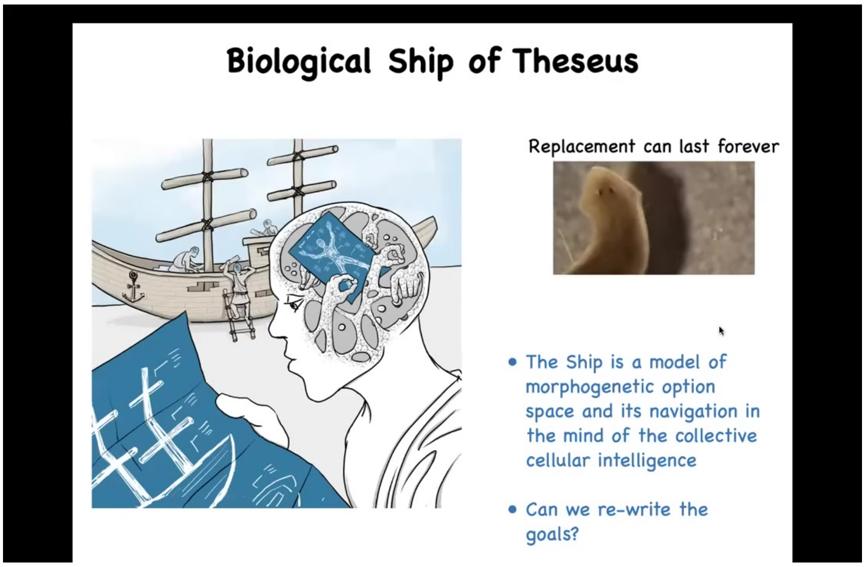

Now, you can think about that whole process as a kind of ship of Theseus in the sense that during, even without drastic injury, cells and tissues are constantly aging, degenerating, being removed from the body, and then being replaced. And yet the whole organism remains. So what does that mean? We need to understand that the actual ship of Theseus is not the physical body. It's the pattern in the proto-cognitive medium of the cells that are doing the replacement that allows this thing to remain and to conserve its identity. That is, the cells themselves are a ship that literally replaces itself, and it's the information structures within the components of the ship, the replacement machinery that are actually responsible for this.

Now, notice that this can last forever. In other words, this planarian, which is incredibly good at regenerating, so much so that when individual cells senesce and die, they are regenerated. The animal as a whole has no obvious lifespan limit. The asexual form of these flatworms can live forever. There's no sign of aging in these guys. They continuously regenerate themselves.

What this is is a model of that kind of morphogenetic option space of different kinds of shapes and the continuous navigation of the collective cellular intelligence, whose job it is to build and repair and maintain this thing, despite all sorts of weird perturbations that happen both naturally and the kinds of things that I'm going to show you we do experimentally. Given that you have this very active system that continuously maintains the form, can we rewrite the goals? That is, under standard circumstances they always try to maintain this. Could we get them to build something else?

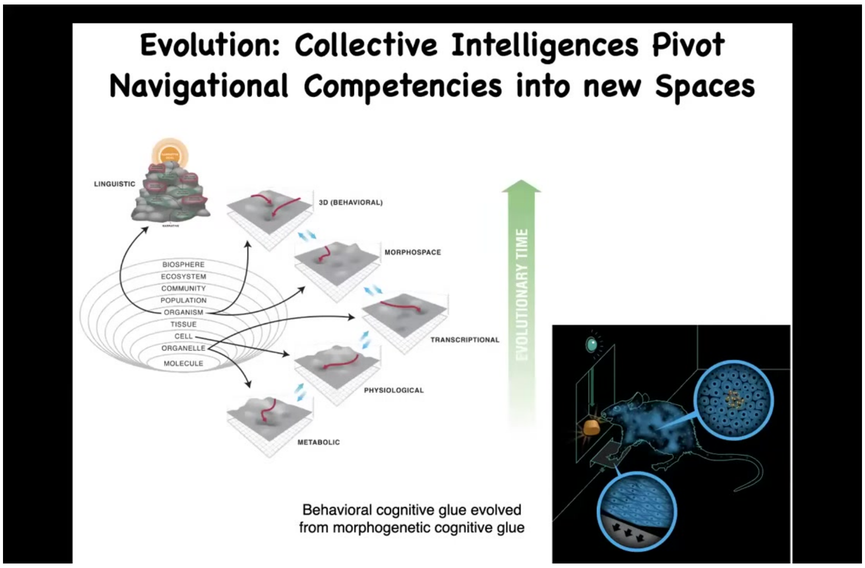

And what we're seeing is that the competencies that this material has have been rotated through different spaces by evolution. Originally, in very simple life forms, they had some metabolic set points and then physiological and then the genes came along and transcriptional kinds of goals. Multicellularity came along and they were able to pursue anatomical goals and morphospace. Eventually, nerve and muscle developed and they were able to run around in three-dimensional space and have behavioral goals, conventional behavioral goals. Maybe linguistic goals as humans navigate land, and now other language models navigate linguistic space, and so on.

The idea here is that there's a strong relationship between the policies that allow the subunits of biological tissue to navigate these various spaces and the kind of behavioral cognitive glues such as the electric signaling within our brains that allow individual neurons to work together to form some sort of emergent higher-order agents, such as a rat that learns to press a lever to get a reward. No individual cell has both experiences. The cells at the foot touch the lever, the cells in the gut got the reward. But there exists a rat that owns the associative memory because of this bioelectric cognitive glue that binds the individual neurons together into one higher-level agent. This kind of thing evolved from being pivoted into various problem spaces in evolutionary time. Could we take advantage of this information to reassign some of these goals?

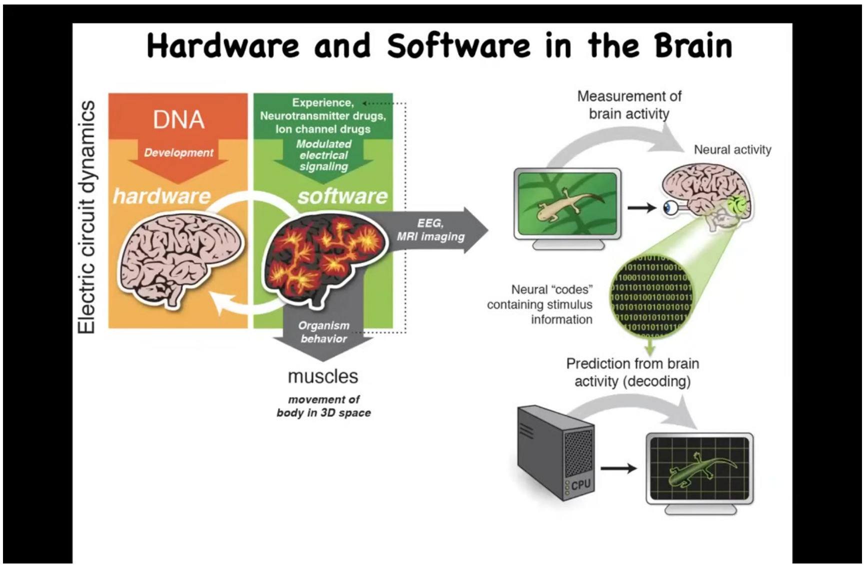

The amazing thing is that this electrical system runs a kind of software. The physiological activity and the information processing in this network are what store the memories, the goals, the preferences, and so on.

In the case of behavior, it operates the muscles to move you through three-dimensional space. Neuroscientists try to read this electrical information and decode it so that they understand what the goals are. But it turns out that system evolved from a much more ancient system that was the same.

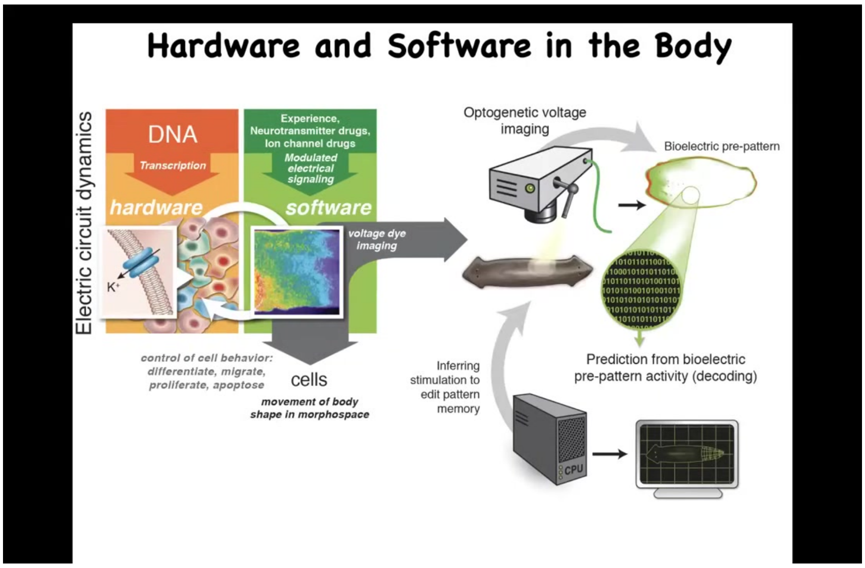

In terms of molecular mechanisms, it's the same. It uses ion channels and gap junctions, electrical synapses, neurotransmitters, all the machinery is the same, but it's extremely ancient. It arose around the time of bacterial biofilms.

What it was doing before nerve and muscle came on the scene was, instead of controlling muscles to move you through three-dimensional space, controlling all cell behavior to move the configuration of the body through anatomical morphospace. This kind of model uses what we know about evolution of the nervous system to understand that morphogenesis, the creation of the body and the maintenance and repair and so on, is a kind of behavior.

It's a behavior of a collective intelligence, the same way that your behavior is the result of a collective intelligence of your neurons. While your brain thinks about all kinds of goals in three-dimensional and other spaces, this system thought about shape. It thought and solved problems in anatomical morphospace.

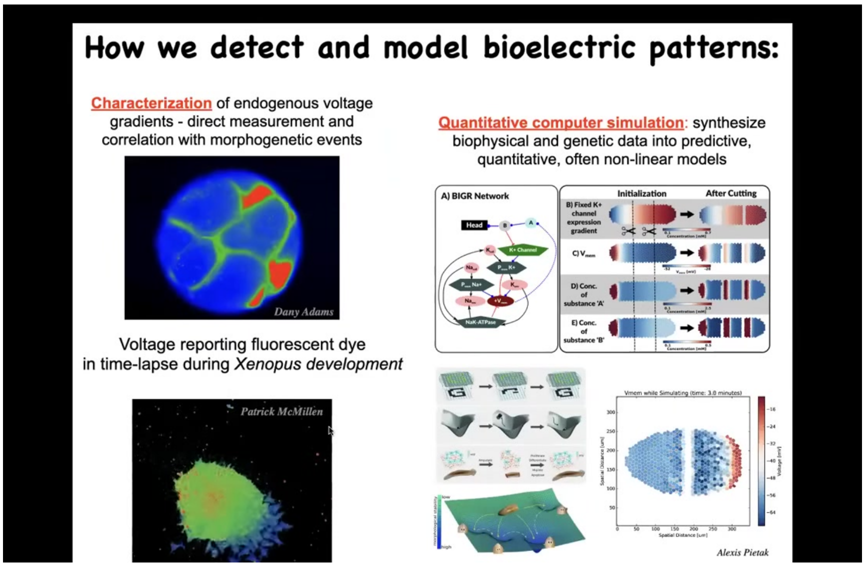

We developed some of the first tools to read and write these electrical pattern memories so that we can start to decode the content of this intelligence.

For example, here we're using a voltage-sensitive fluorescent dye to look at an early frog embryo or here in some explanted cells. The colors indicate different voltage levels. Think of this as a brain scan, except these are not brains. These are cells figuring out what a frog embryo should look like and solving various other problems.

We collect this kind of data, then we do a lot of quantitative simulations all the way from the molecular level up into the tissue and whole-organism level to try to understand how they navigate that space, why the particular electrical gradients are what they are, how the network can do things like pattern completion, which you see during regeneration when a piece of it is missing.

We have a simulator, this beautiful simulator made by Alexis Pytek that allows us to simulate some of these processes. I'm going to show you what that looks like.

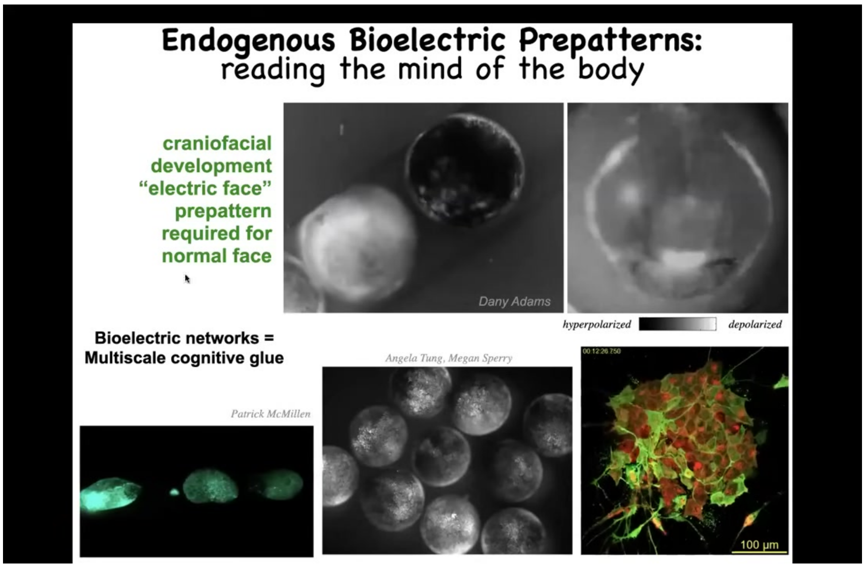

Here is a time lapse of an early frog embryo putting its face back together from scratch.

You can see the voltage map. Here's where the eye is going to be. Here's where the mouth is going to be. This is one frame from this video. We call this the electric face for obvious reasons. It's the bioelectric pre-pattern that tells the cells which genes to turn on and what the face is going to look like. These bioelectrical networks are not only the cognitive glue for these cells, meaning that they bind all these cells into the same journey in physiological state space, but they do the same across embryos. It's not only within an embryo, but across embryos.

You can see what happens. This embryo triggers a wave that then is picked up by all these other embryos. Here's one. This one is going to get poked, and you can see within some number of seconds, these two find out about it. This is very rich electrical dynamics. We now have the ability to look into this tissue and try to understand what the electrophysiology is that is responsible for some of these morphogenetic events.

Slide 18 of 54 · Watch at [22:20]

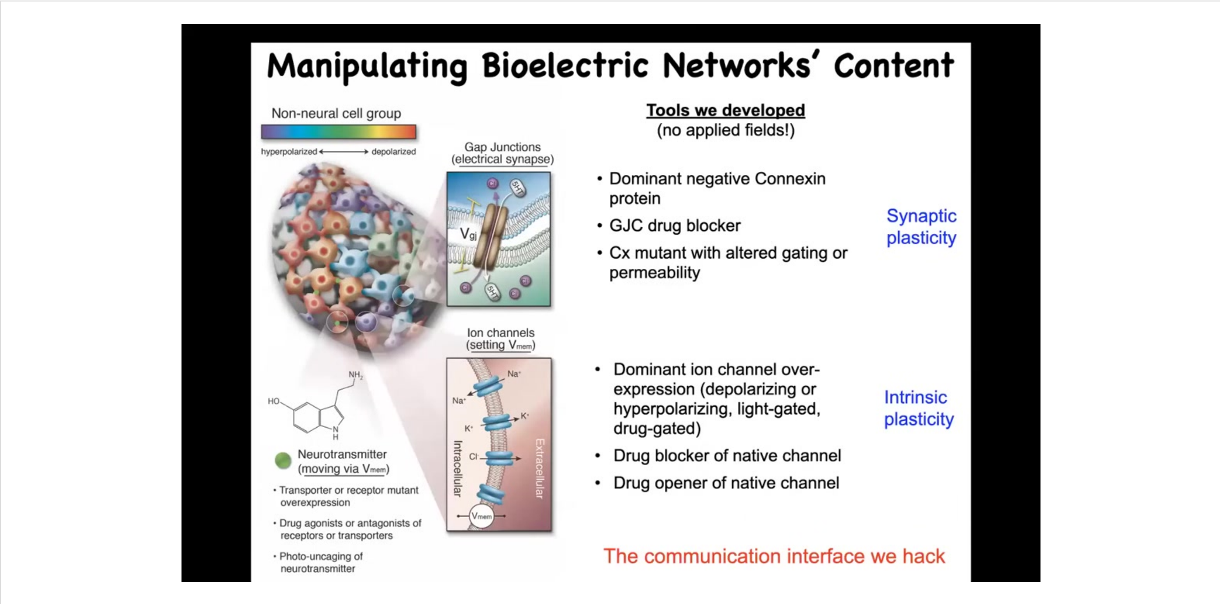

Now it's important to be able to manipulate those and rewrite those patterns so that we can make sure that these patterns are functional and so we can do experiments on the plasticity of that biological layer.

What we've developed are tools to write that information. We don't use any fields, electrodes, magnets, waves, or radiation frequencies. What we're doing here is hacking the normal interface that these cells use to shape each other's behavior. So that means targeting the ion channels, gap junctions, or the electrical synapses.

So just appropriating the tools of neuroscience to other cell types. The tools can't tell the difference. It's entirely likely that the whole distinction between neural and non-neural cells is in many ways artificial. It helps us maintain departments and universities, but the tools work perfectly well on all cells because most cells in the body do the same kinds of things that neuronal cells do, albeit much slower and perhaps in different spaces.

This is the communication interface that we control. We use optogenetics and pharmacology to open and close channels. That's how we rewrite patterns into these tissues.

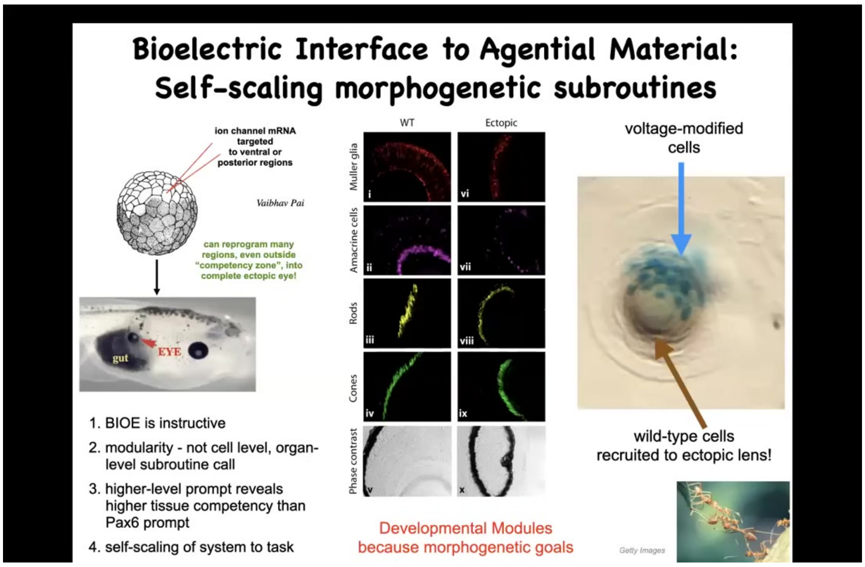

I'll show you a couple of examples. We could do a couple hours on this, but as one example, I showed you a little voltage spot that says where the eye is going to be. You could try to reproduce that voltage spot somewhere else. Let's say we inject some potassium channel RNA to reproduce that voltage spot here. The cells make an eye, and here are the sections through that eye: lens, retina, optic nerve.

You learn a couple of things from this. First of all, these bioelectric patterns are instructive. We've learned to induce predictable complex changes in anatomy.

The next thing you see is that it's extremely modular. We provided a very simple trigger, and what it did downstream was respond with a very complex morphogenetic behavior. This is a high-level subroutine where all we say is "build an eye here," and then the material does the rest. We don't have to tell it how to build an eye, how big it should be, or what the components should be. We're not micromanaging it. We are programming at a high level of abstraction. We're calling up specific organs.

The material has some other interesting features. If you only target a few cells and there's not enough to build an eye (the blue ones are the ones we targeted), they recruit other cells to help complete the task, much like other collective intelligences, ants and termites. We didn't have to do any of that. The material is already competent to do this.

But it's up to us to discover what all the subroutine calls are. We now have the API. The question is, what are all the subroutine calls that we can make? What is the material capable of? As I'm going to show you, it's quite a lot.

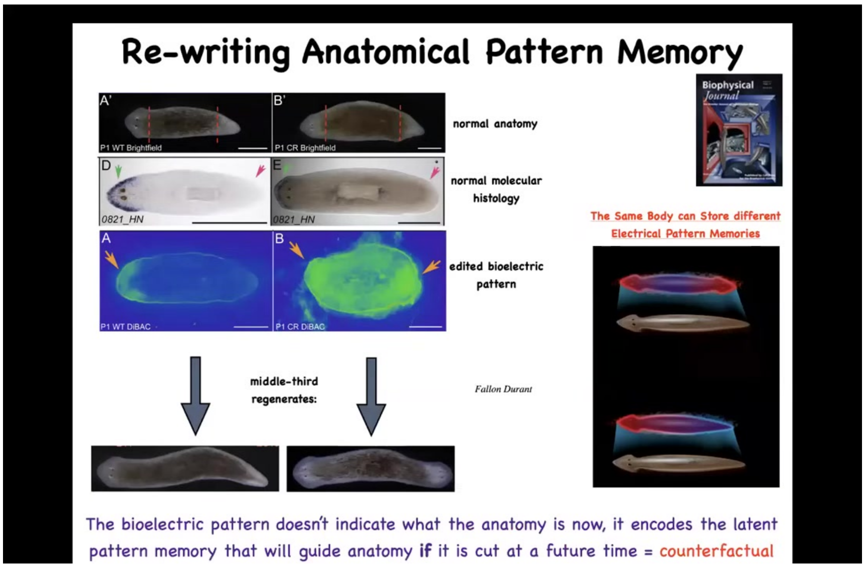

Another thing we can do is we can ask, how does this planarian know how many heads it's supposed to have? Now you might think it's genetics, but what does that really mean? When you amputate the head and the tail, this little fragment has an interesting electrical gradient. It turns out that the fact that it has one depolarized region says build one head. So we can change that. We can make it like this. The technology is still being worked out, but we can provide a different pattern. If you do that and then cut this animal, it will in fact generate two heads. This is not Photoshop, it's not AI, these are real animals.

Interestingly, this voltage map is not a map of this animal. This voltage map is a map of this perfectly normal one-headed animal. It's going to stay that way because this memory is latent until it gets injured. What this is is a representation of the pattern to which I will regenerate if I get injured in the future. It's a counterfactual memory. It's a very primitive form of that time travel that brains are able to do in thinking forward and back about scenarios that are not true right now. We can rewrite it and then these cells will build whatever it says here.

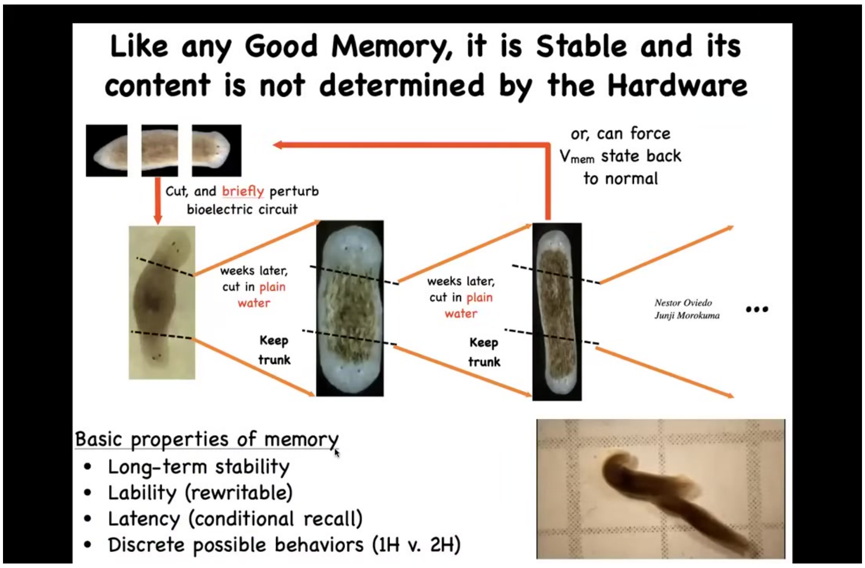

If you do this, it's permanent because the electric circuit that holds that information has a memory property. So once you have these two-headed worms, you can continue to cut them. As you cut them as many times as you want, they will continue to build two-headed worms. This is what's wrong with saying that the information is in the genome, because there's nothing wrong with the genome here. We haven't touched the genome, we haven't edited the genome. If you were to sequence these two-headed animals, you would see nothing different from the one-headed animals. The information about head number is not directly in the genome. And you can see here that these two-headed guys are doing what they do.

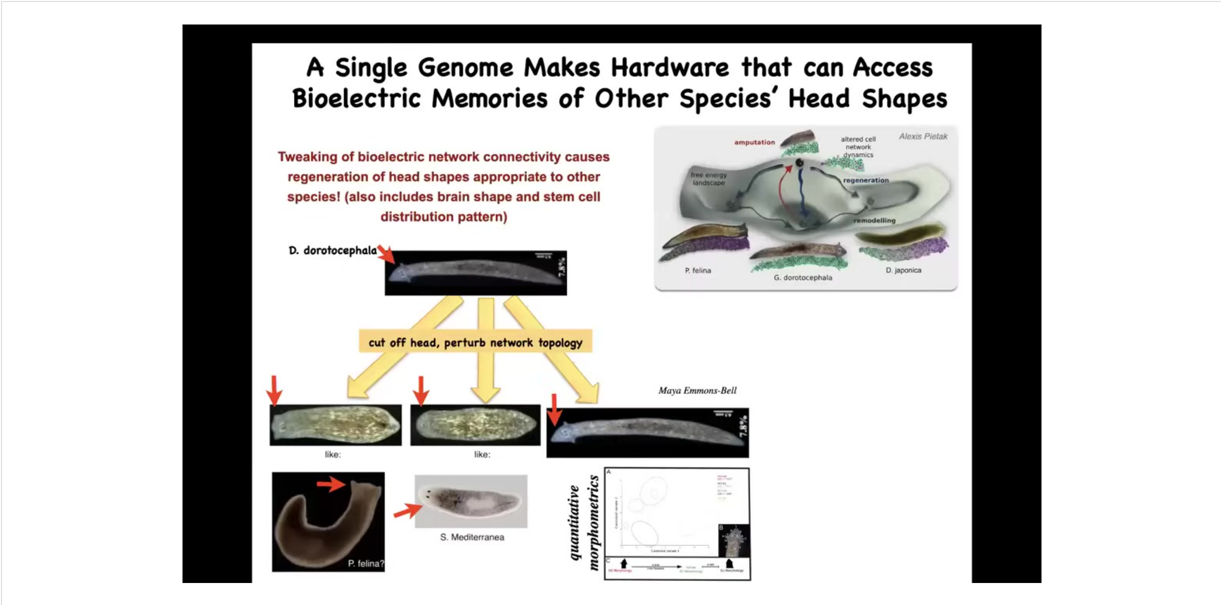

You can also ask them to, instead of making two heads, make heads of different species. This perfectly normal, genetically normal hardware is able to visit other attractors in anatomical state space corresponding to other species. It can make a flathead like a P. felina; it can make a round head like an S. Mediterranean. They're about 100 to 150 million years in distance, and they're perfectly happy to do that.

The shape of the brain and the distribution of stem cells are just like these other creatures. We're starting to see the incredible plasticity, the use of bioelectrical networks to navigate the anatomical space the way that we conventionally use bioelectrical networks to navigate three-dimensional space.

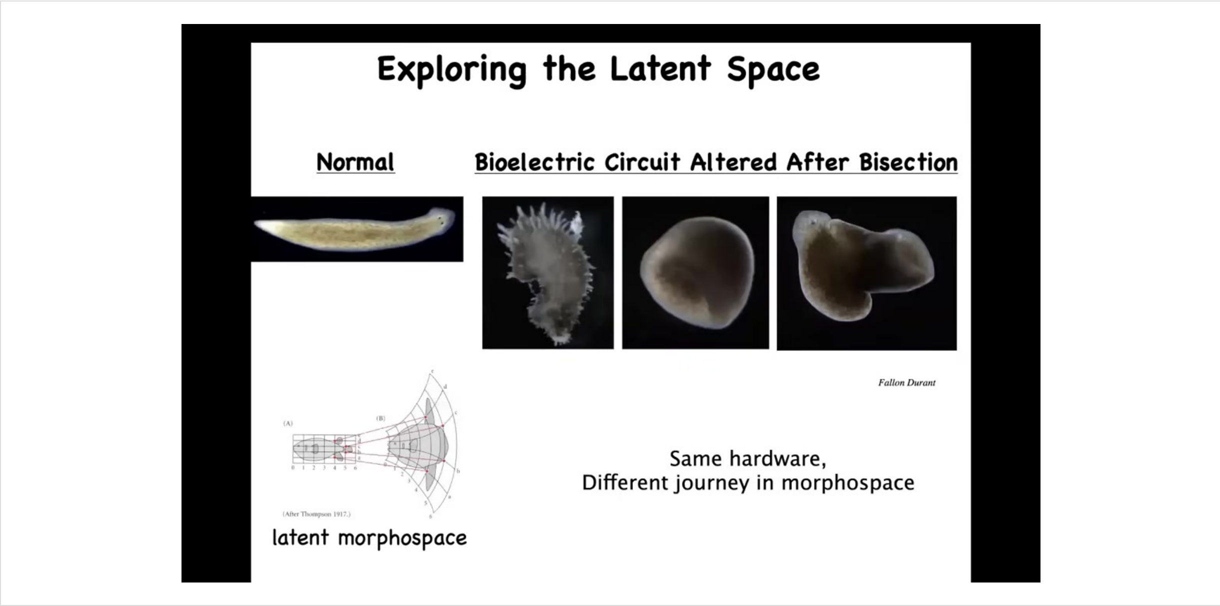

And we're starting to see now that by rewriting these pattern memories, we can start to get an idea of how we can control this navigation in the morphospace. We can also reach regions of morphospace that don't look like planaria at all. This crazy spiky-shaped thing, a nice round cylindrical thing, a hybrid form — all of this is available. Same hardware, different journey in anatomical morphospace. These electrical patterns are a way to explore the latent morphospace of what that genetically specified hardware is capable of.

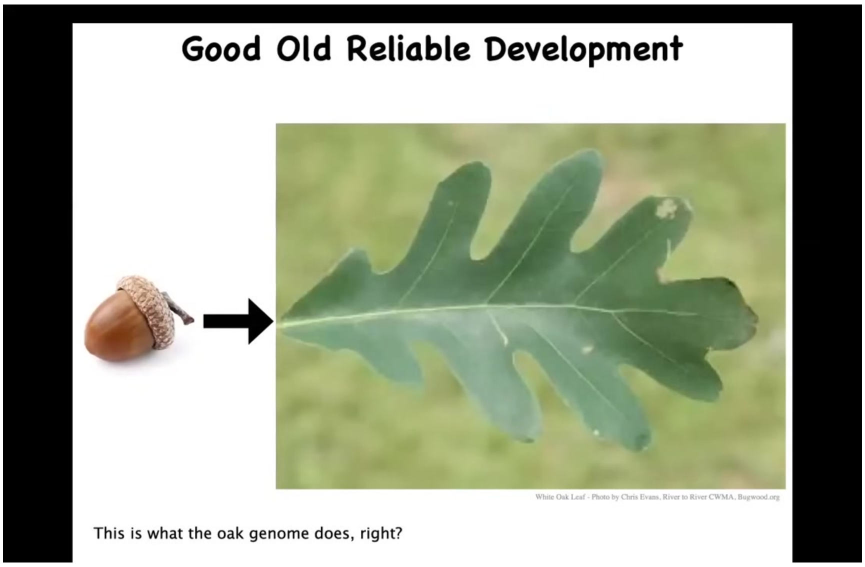

Now, we are not the first to do this. I'll show you an example of a non-human bioengineer. Here's an acorn. Acorns make oak trees, which have these leaves 100% of the time. Billions and billions of these leaves in any given area every year. You start to get the idea that this is what the oak genome is capable of. This is what the oak genome encodes.

Slide 25 of 54 · Watch at [29:12]

But along comes a non-human bioengineer, this little guy, it's a wasp. And what the wasp does is interesting. It prompts the cells to build what it likes, which is this. It drops off some chemicals with its embryo, and it hacks the cells to build this kind of gall.

Now, if this wasp hadn't done that, you would have had no idea that these cells are actually capable of it. We don't know what the material is capable of until it is prompted with the right signals. And this wasp is not there building it the way that it builds its actual hive, which is much more like a 3D printing process where they build a thing step by step. This is not like that at all. These are all plant tissues. All it did was convince the cells to build this kind of structure.

And so the biology is certainly exploiting the reprogrammable hardware, not just for the plasticity that I showed you before, but actually they're using it to hack each other. All living things are constantly hacking each other.

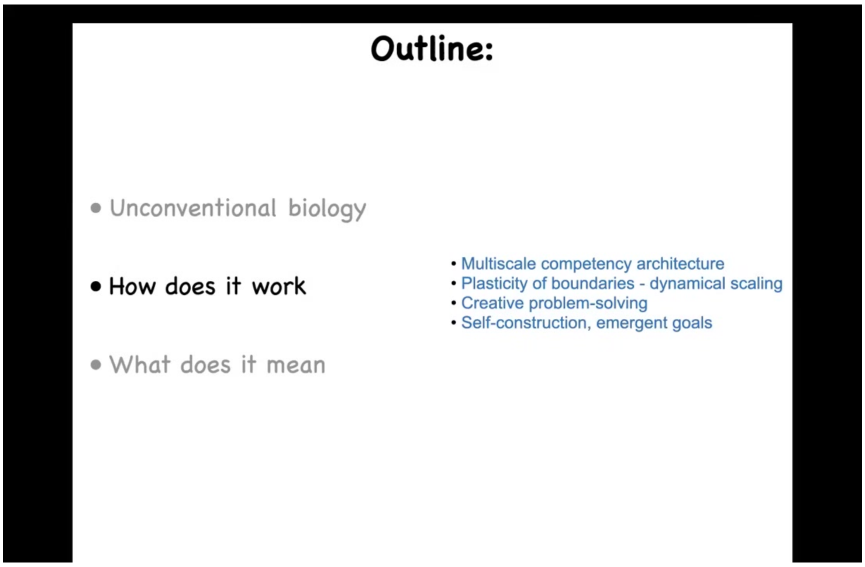

What I want to do now is talk about specifically ways in which this biological architecture is different from how we do computing. And that's going to become important to start to answer this question of whether our computational paradigm is sufficient for understanding biology and what we can learn from the biology. We're going to run through a couple of examples involving the multi-scale competency architecture, the incredible plasticity and the dynamical scaling of boundaries, and some problem solving.

Let's remind ourselves of how we got here. We all started life as a single cell. This is a quiescent unfertilized oocyte. It's a little BLOB of chemistry. And through the slow and gradual process of development, we become one of these things or even one of these things. The process is continuous. There is no point at which a magic lightning flash converts what used to be just chemistry and physics into something that has mind. This is conventionally thought to be amenable to chemistry and physics. This is thought to be. We have real minds that are best dealt with by communication, psychoanalysis, and behavioral science.

What happened? How did we get from here to here? One thing we clearly understand is that we need to develop the models of the scaling. There was no special point at which these mental powers kick in. We need to understand the scaling from these systems to these systems.

This is not even the end of the story. Some of the cells that are participating in this — and they did an amazing job creating all this stuff — can actually defect. This is a process known as cancer. We'll talk about that. They can disconnect from that electrical network. As I'll show you momentarily, some cells from the body can have an entirely new life in a new configuration. These are called anthrobots, and I'll show you those. In fact, after the patient donates the cells, they may or may not be alive. In the case that they're not, these are having an interesting kind of life after death of the original organism.

Slide 27 of 54 · Watch at [32:27]

That's the first unsettling point: we are all slowly and gradually scaled up. The second point is that we consist of this; this is the material we're made of. This is a single cell right here. This happens to be a free-living being known as a lacrimaria, but notice what it's doing. It's a single cell. There's no brain. There's no nervous system. There aren't any stem cells. It's an incredibly competent creature that's hunting in its area, maintaining its single-cell goals, its metabolic and proliferative and morphogenetic goals. This is the agential material of which we are made. The plasticity, everything I've shown you in the first part of the talk, the amazing reprogrammability and plasticity of the body, is because we are made of a material that is already intelligent at the very bottom.

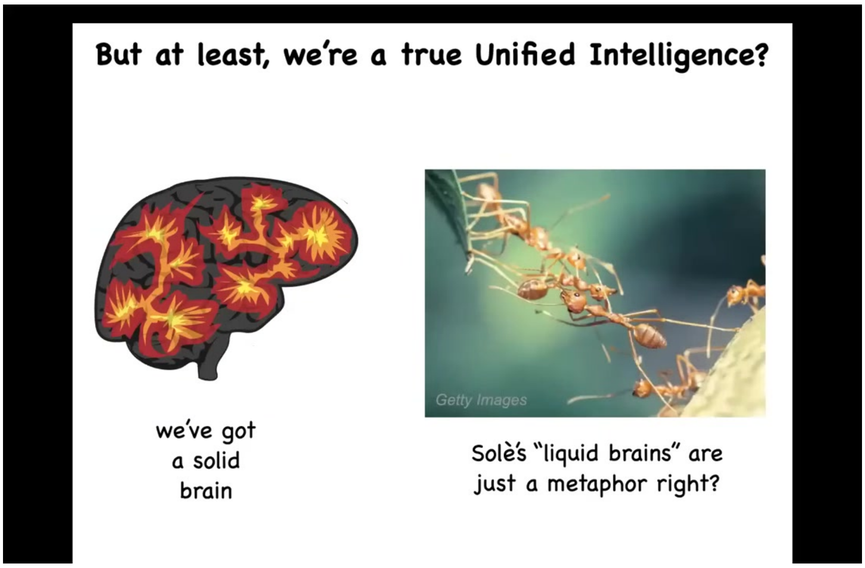

We're made of this agential material. We scale up from a single cell, but at least we are true unified intelligences, right? Ricard Soleil talks about ants and termite colonies as liquid brains, but surely that's just a metaphor. We have this nice, solid brain. We are a singular intelligence, not like this distributed thing, which is maybe just metaphorical.

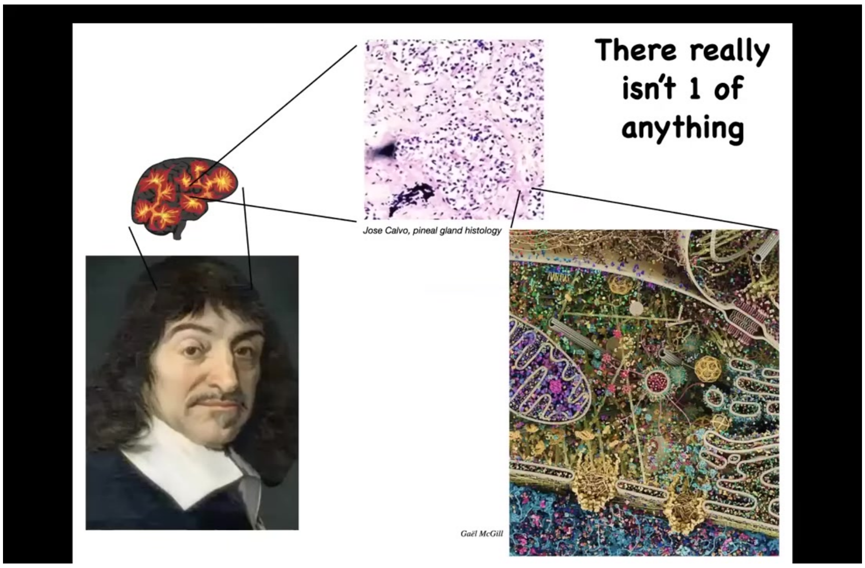

René Descartes thought so, and he liked the pineal gland because there's only one of those in the brain. He thought that the unified experience and the cognition that humans have deserves a single point, a single unified structure in the brain that would be the center of this. That was his idea, but he didn't have access to microscopy. If he had a good microscope, he would have looked inside the pineal gland and discovered that there isn't one of anything. There are tons of cells inside of it. Inside each one of those cells is all of this stuff.

All the way from human bodies and brains made of tissues and cells, and within the cells, there isn't one of anything. We are collective intelligences all the way through.

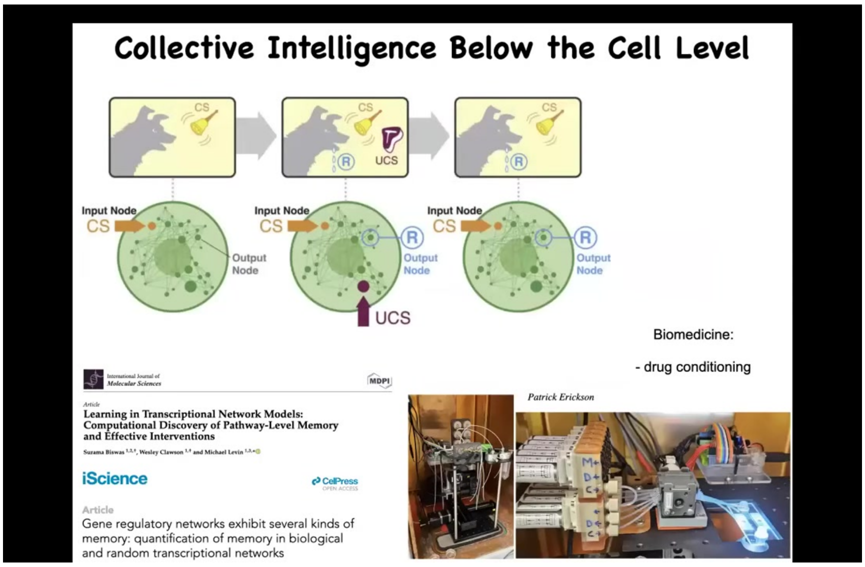

And even below the single cell level, it turns out that very simple things like gene regulatory networks, just a small set of chemicals turning each other up or down, already is sufficient for six different kinds of learning, including Pavlovian conditioning. So in these papers, we describe how you can train gene regulatory networks just by experience. The biological networks have learning ability out-of-the-box. And here we are exploiting this feature to try to train cells with drug stimuli so that we can condition to various drugs. So it's all kinds of biomedical implications of the fact that your molecular networks inside of your cells can learn.

So what we're dealing with is this multi-scale competency architecture where every layer is solving problems in its different spaces.

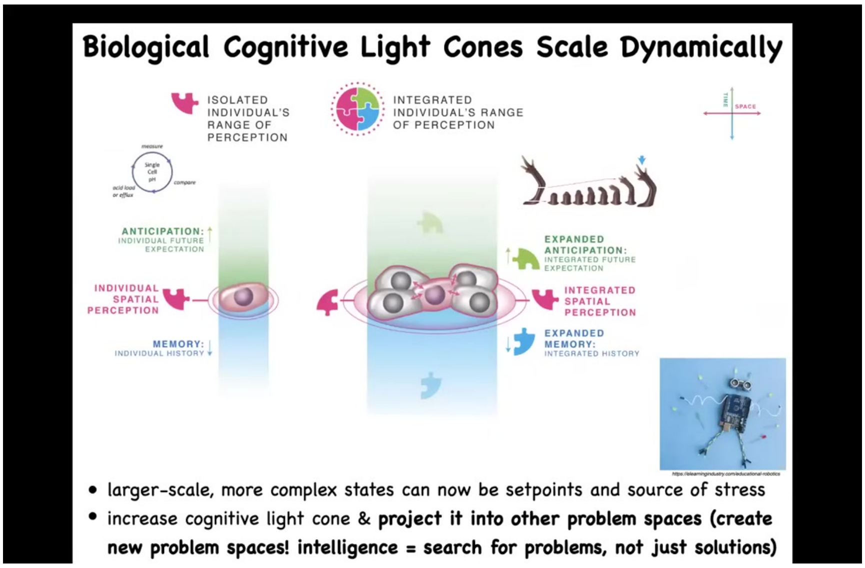

And that means that when we are putting these cells together, what's enlarging is something we call the cognitive light cone. The cognitive light cone is simply the scale of the biggest goal that the system can pursue. So single cells have little tiny goals in space and time. They can do things like keep their cell pH or hunger level or something within the cell. So spatially, they care about a very small region. Temporally, they have some memory going back. They have a little bit of anticipatory potential, but it's very small. But when you put them together, they form electrical networks that have a much larger cognitive light cone. The radius of their sensing and memory goes up, and they are able to work on these grandiose construction projects like building and maintaining large-scale structures. So these networks allow collections of cells to have larger memories and have more complex states be the set points and the source of stress. What do you care about? What kinds of things are you stressed out about when the conditions aren't met? They're able to care about much larger things. And they can project this into other problem spaces.

And this is one place where it's very different from how we built our technology, because mostly if we have a robot, we hope the complete construction is intelligent and can do things, but it's made of dumb parts. The parts that it's made of are themselves not agential. They don't have agendas, which is why I also sometimes give a talk on cancer called Why Robots Don't Get Cancer. They don't get cancer because their individual parts have no agendas. Unlike biology, where you self-assemble and you consist of lots of parts that have their own learning ability and their own goal-seeking competencies of various types. And so what happens during evolution is the increase of these glycons as they project into new problem spaces from metabolic and physiological and so on all the way up. And in fact, they can create new problem spaces. That kind of intelligence is not just the search for solutions, it's actually an exploration for new problems to solve.

Slide 32 of 54 · Watch at 37:45

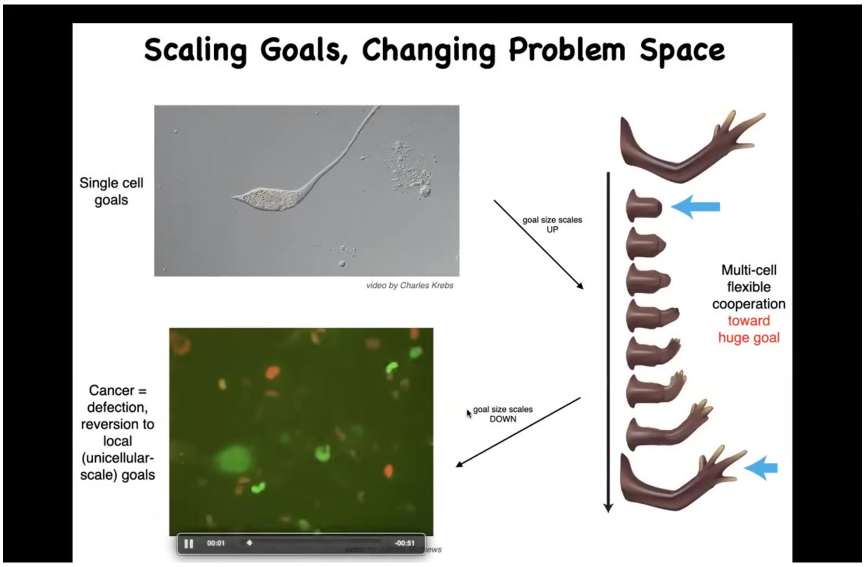

Here's what happens during evolution or development. Individual subunits like this work together to scale up their goals incredibly, so that they're working on these giant things. Then there's a failure mode where the individual cells can detach from the electrical network. At that point, they can't remember or think about anything like this. They go back to their ancient unicellular goals. This is cancer. This is a human glioblastoma here. They just go where life is good. As far as they're concerned, the rest of the body is just external environment. They're not more selfish than other cells. They just have smaller cells.

That boundary between self and world shrinks. Now instead of this being one giant self in anatomical space, it's just back down out to single cells. So you see the automatic dynamic scaling up and down of the border between self and world. It can shrink and grow. This is also not something we currently have in most of our computational models.

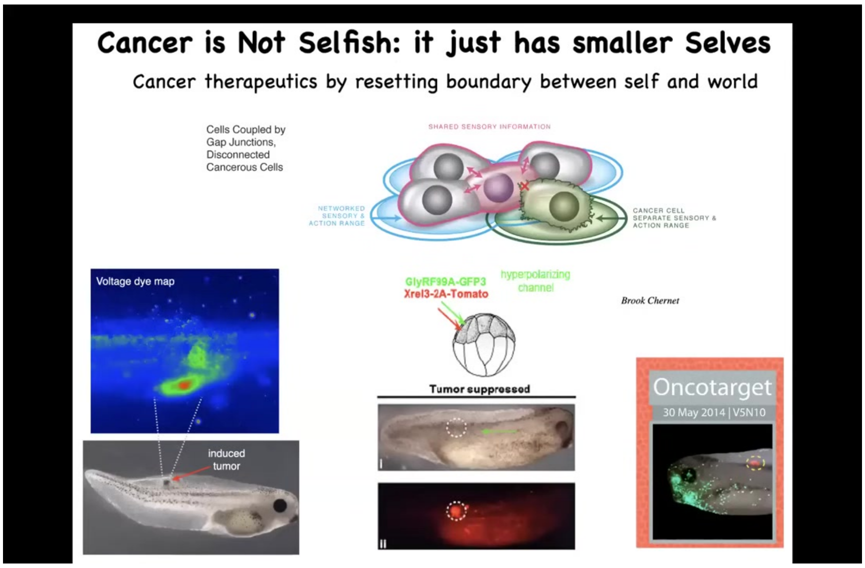

So those kind of ideas suggest therapeutics. I haven't talked about any of our regenerative medicine approaches, but this is just one simple thing. You inject a human oncogene here that will eventually induce a tumor, and you can see how electrically these cells are already disconnecting from the network. What you can do is instead of, as currently done, try to kill those cells with chemotherapy, instead of killing them, you could forcibly reconnect them to the network. And that's what we've done here by injecting an ion channel that keeps them in the right electrical state. And you see, this is the same animal. Even though the oncoprotein is very strong. There's a serious hardware malfunction. Actually, the protein is mutated, but there is no tumor because some hardware defects are actually fixable in software. If you keep the cells connected to the network so that they can all remember that they're supposed to be working on nice skin and muscle and various other things, then you don't see this kind of metastatic behavior, even though the hardware defect is there. I'm not suggesting this will be true for all hardware defects, but we're working on a wide class of birth defects, traumatic injuries, and other things that we can treat this way. And so what we did here is we artificially reset the boundary between the self and the world. We've made it larger.

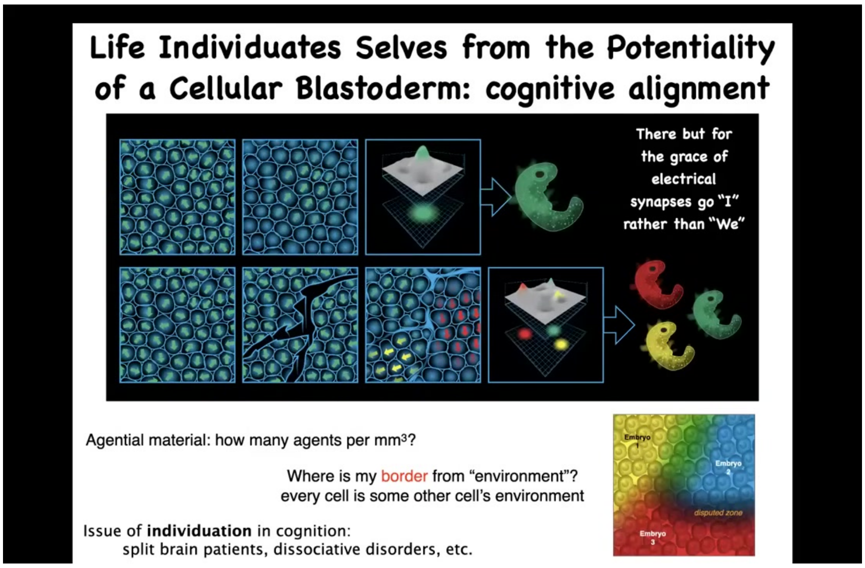

Another thing that's interesting about biology is that it isn't even clear how many intelligences are within any given piece of material.

This is an early embryonic blastoderm, and there are 50,000 or 100,000 cells. Why do we say that's one embryo when we look at this? There's many cells. What is there one of? What there's one of is alignment. It's alignment of all of these cells towards a specific journey in the anatomical space where they're going to build a specific thing.

What you can do, and I used to do this as a grad student with duck embryos, is take a little needle and put some scratches into that blastoderm. If you do this before they heal up, each of these independent areas can't feel the others, and it decides that it's the embryo; it organizes itself. When they heal, you have triplets, twins, or quadruplets.

When you look at this material, how many individuals there are is not set by the genetics. It could be anywhere from zero to half a dozen, depending on how the physiology goes.

This kind of uncertainty is not something we have in computer science, where if you have some number of transistors on a chip, you can determine exactly what the memory capacity is and what the material is capable of. But here, much like with neuroscience: if I showed you a human brain and you didn't already know what a human was, it wouldn't be obvious how many individuals are in there. The question of how many cells per cubic millimeter of substrate we actually don't know, because it self-assembles.

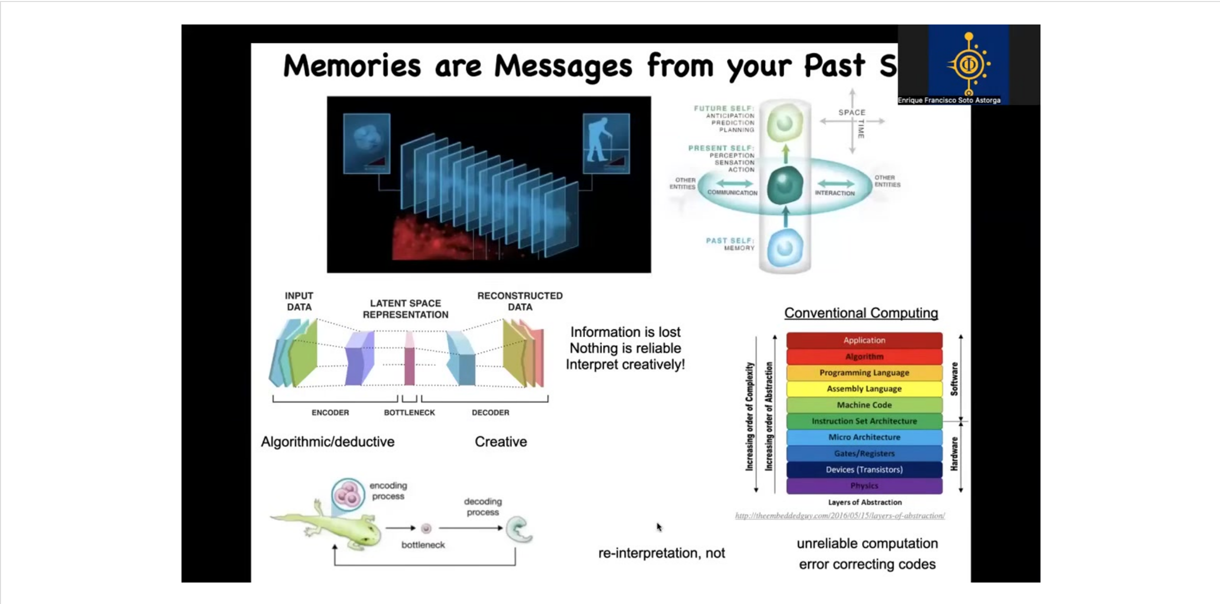

This is another fundamental way in which computing is different from what biology does. What we have in conventional computing is a clear set of abstraction layers where if you're working up here, you don't need to be thinking about how warm your copper is. You don't think that your registers are going to float off because something's happening to the silicon. It's completely isolated. You can trust that the hardware is behaving as appropriate. In the field of unreliable computation, what we do is we try to maintain this kind of system with a bunch of error-correcting codes. The idea that, using some kind of redundancy or various other kinds of systems, we can actually pretend that this stuff is completely reliable from the perspective of the higher levels.

In biology, at any given point a living organism doesn't have access to the past. What you have access to are the engrams, the memory traces that the past has left in your brain and body. What you then need to do is interpret those traces in the appropriate context for behavior now. Because these memories are memories from a past self—it's your past self that's passing on these memories to you—Sam Gershman has a funny line where he says, "your most important collaborator is you six months ago, and he's not answering emails." At any given moment, all you have are the messages that were left to you by your past self. As with any messages, you're free to interpret them however you want.

What happens in biology is this really interesting architecture that you may recognize as a kind of autoencoder, where there's a bottleneck in the middle and the experiences that you've had before get compressed and squeezed down into a much smaller representation that generalizes from the particulars into the memory. When it's time to decode all of that, there's a lot you have to add because information has been lost; all the correlations have been lost. The biological medium is entirely unreliable.

As a biological organism, everything is going to change. You will be mutated. The environment will change. Where the encoding process is algorithmic and deductive, for a living organism this part is creative. You have to interpret as best as you can the information that you have. You cannot take it literally.

Here you see the same thing happening with evolution. Animals typically don't give rise to other animals. They give rise to an egg. The egg has to reinflate; it has to be decoded into an organism. As I showed you from the beginning, that is not a hardwired process. You can make changes, do all kinds of perturbations. That information will not be taken literally, but it will be used to construct something as best as it can.

This is a huge difference. Whereas in computing we are committed to the fidelity of the information, here we are committed to the creative interpretation of the information as messages. It's not anything you can take literally.

Slide 36 of 54 · Watch at [45:12]

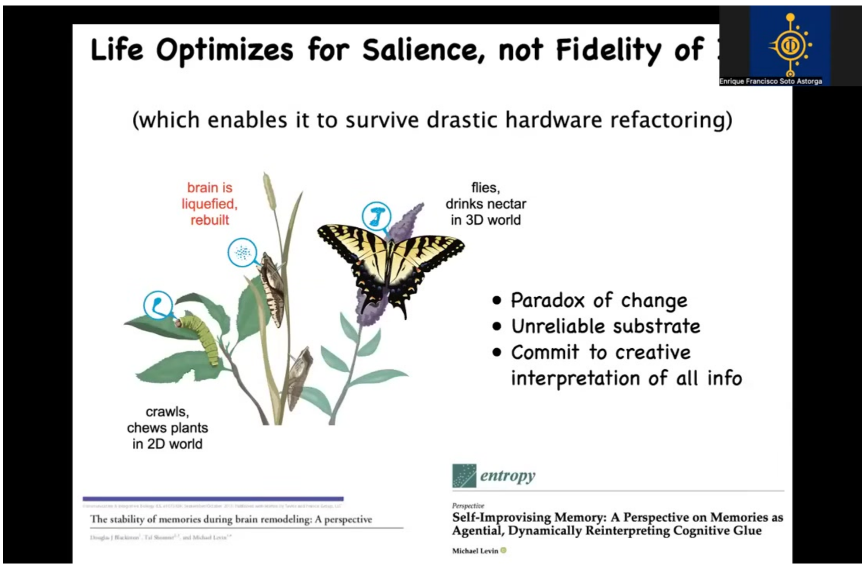

This is the paradox of change. If you're a species, if you don't change, then you will surely die out as the environment and other things will change. But if you do change, then you're not the same anymore. And again, you're gone. This is the idea that you can only persist not as a fixed being, but as a process, as a creative self-interpreting process, such as the caterpillar-to-butterfly transition.

All of this is in these papers.

What happens, for example, is that caterpillars can be trained. They dissolve their brain and build a completely new brain that's suitable for a hard-bodied creature that flies in three dimensions. This is a two-dimensional being. Butterflies remember the original information, but what's critical is that the actual memories of the caterpillar are of no use to the butterfly because it doesn't like the same food, it doesn't want leaves, it wants nectar, and it doesn't move the same way. Knowing how to move in order to get leaves is of no relevance to the butterfly, but what it does have to do is reinterpret that information in a new way that makes sense for it.

This unreliable substrate is the rule in biology. Biology from day one commits to creative interpretation, not persistence of information.

Slide 37 of 54 · Watch at [46:35]

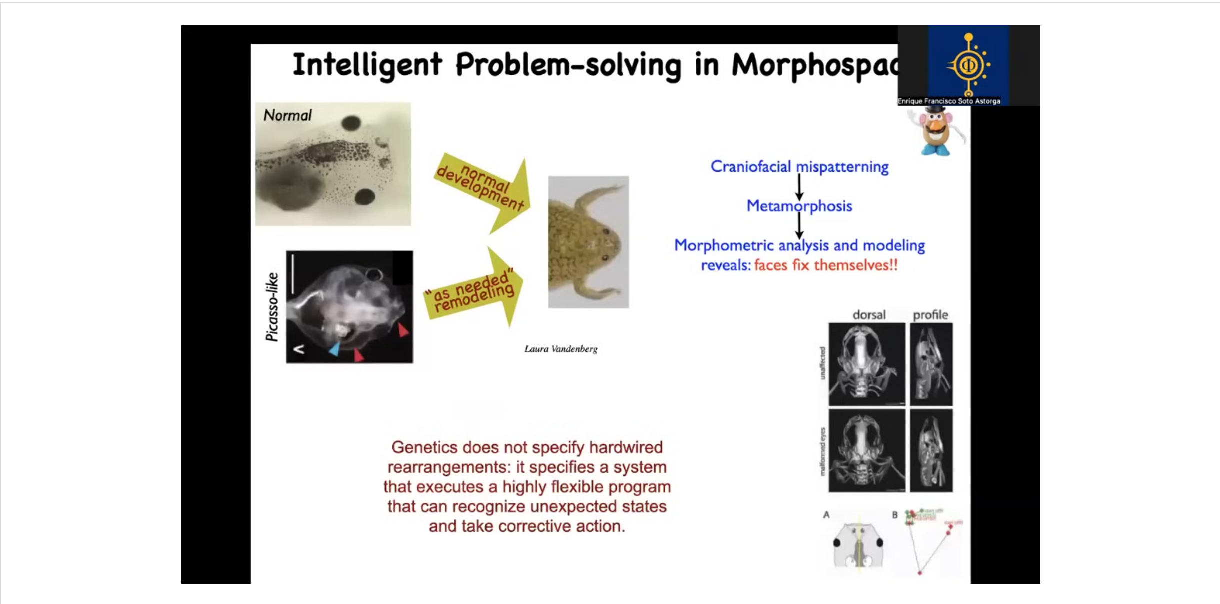

That's what allows it to do things like this.

If we take a tadpole, which normally turns into a frog, and all the parts of the face have to move in a particular way to do this, we can make what we call a Picasso tadpole where all the organs are starting off in weird positions. The eyes are on the back of the head, the mouth is off to the side. This does not make a weird-looking frog. It makes a perfectly normal frog because all the components move around in novel paths to get to where they're going.

Because the genetics doesn't specify hardware rearrangements. It specifies a problem-solving agent that was not taking the past literally, and it can't.

Slide 38 of 54 · Watch at [47:15]

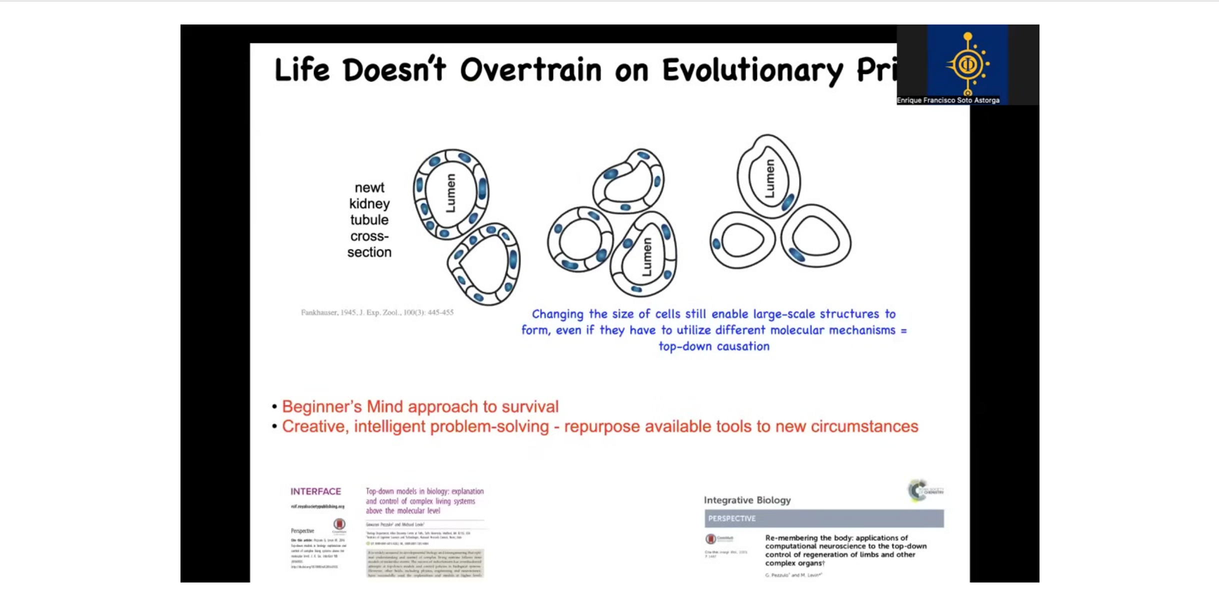

Here's one of my favorite examples. This is a cross-section through a kidney tubule in a newt. You can see about 8 to 10 cells working together. What you can do is create newts that have multiple copies of their genomic material. When you do that, the cells get bigger. Remarkably, the newt stays the same size. How is that possible? Fewer cells are used to build that same structure. If you make a newt that's 6N, high copy number, the cells get so big that one cell can wrap around itself and create the normal structure.

What's interesting about this is that these are completely different molecular mechanisms. This is cell-to-cell communication. This is cytoskeletal bending. This is a clear example of intelligence where you solve a problem you haven't seen before by using different tools at your disposal, the molecular mechanisms that the system calls up to get the job done.

Think about what this means when you're a newt. You're coming into this world. You don't know how many copies of your genetic material you're going to have. You don't know how big your cells are. You don't know how many cells you're going to have, and you still have to get the job done. You can use the wisdom of past generations that have given you all this material, but you can't use it in a hardwired way. You have to adapt on the fly. It's a beginner's mind idea. Creative problem solving by repurposing the tools you have.

Slide 39 of 54 · Watch at [48:50]

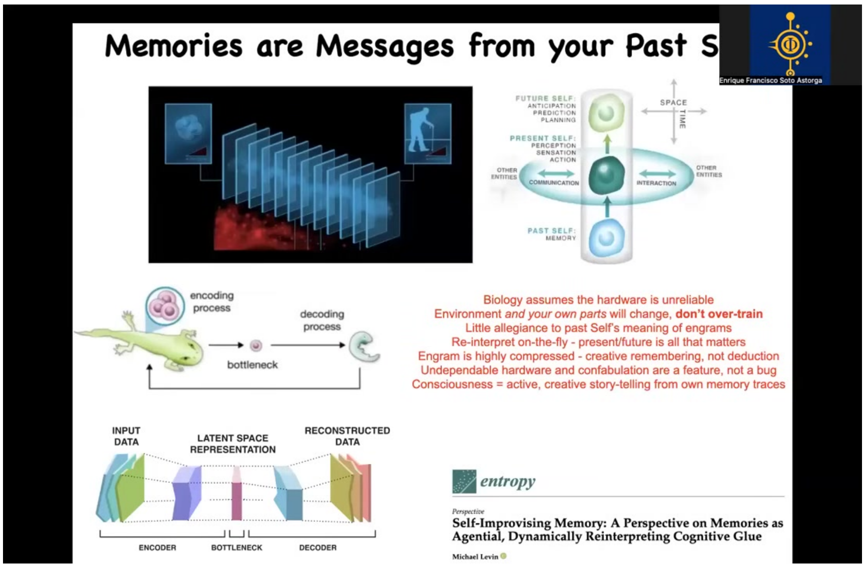

Here's what's going on with the biology. Biology assumes that the hardware is unreliable, that both the environment and your own parts will change. You can't overtrain on priors. You can't have allegiance to your past self's meaning of your memory traces. You have to reinterpret them. What did they mean for me now? That might have been left by a caterpillar. I am now a butterfly. I have to reinterpret these things in new ways. My body is different. My goals are different. Everything is different.

Because the n-gram is highly compressed, memory is a constructive, creative process. It's not a deduction. This is actually a feature. The fact that you have undependable hardware forced evolution to produce these creative agents that are intelligent all the way from the beginning of solving anatomical problems to the much higher degrees of intelligence of us in different spaces, such as linguistic space. That's the feature of the biological system: it is creative at every step.

You might imagine, this is speculative, but one way to think about consciousness is as the active, creative storytelling from your own memory traces. It's what it feels like to be continuously telling an ever-shifting, adaptive story about what your own memories mean to you in your current environment.

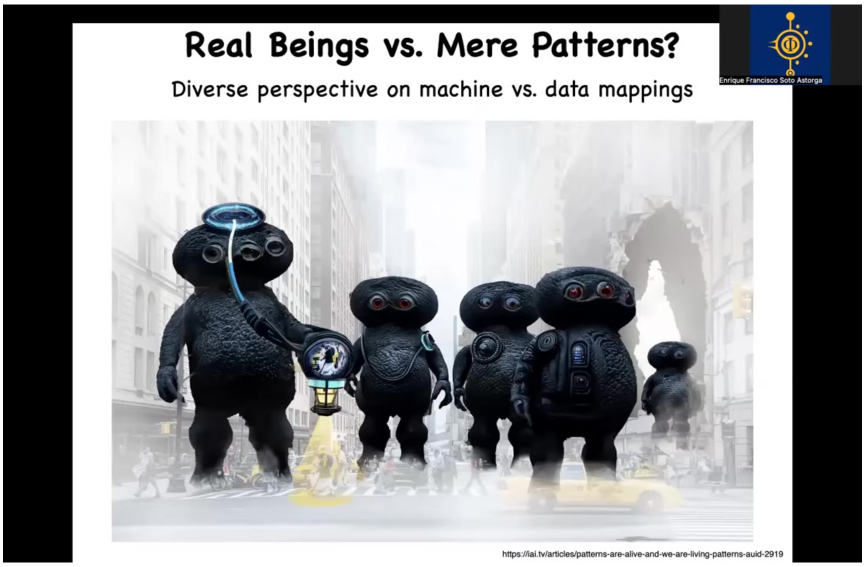

If all of that wasn't weird enough, I want to get a couple of degrees even more unconventional. I want to start by telling you a story. This comes from an old science fiction kind of story from some decades back.

Imagine that these creatures came from the center of the earth. They are incredibly dense because they live in the core. They come up to the surface. What do they see? They certainly don't see any of our objects. What we think of as objects is a fine, tenuous kind of a plasma around the surface. It's even less dense than a gas to them. And so they're walking around, crashing through our stuff. They don't even see it the same way that when you walk through a garden, you disturb all of the fine patterns of scents and pollen and everything in the air. They're not aware of it. They're walking around.

One of them is a scientist and he's been watching this gas and he says to the others, "I almost think that there are some kind of patterns in this gas that seem to be doing things. It's like they're agential. They hang together for a while and they seem to have some kind of goals that they pursue." The others say, "That's crazy. Patterns in the gas can't have goals. We're real physical beings. We are thinkers. We are real cognitive systems. How can patterns in a gas do anything? How long do these patterns hang together?" "About 100 years." "Nothing interesting can happen in 100 years. That's crazy. These are just patterns in a medium."

What the story is reminding us of is that this question of what is the machine and what are the data patterns that move through the machine, what are the thoughts and what are the thinkers? This is William James's claim that thoughts are thinkers themselves, little tiny thinkers. This distinction between the machine and data is observer relative. Depending on your perspective, you might see this quite differently.

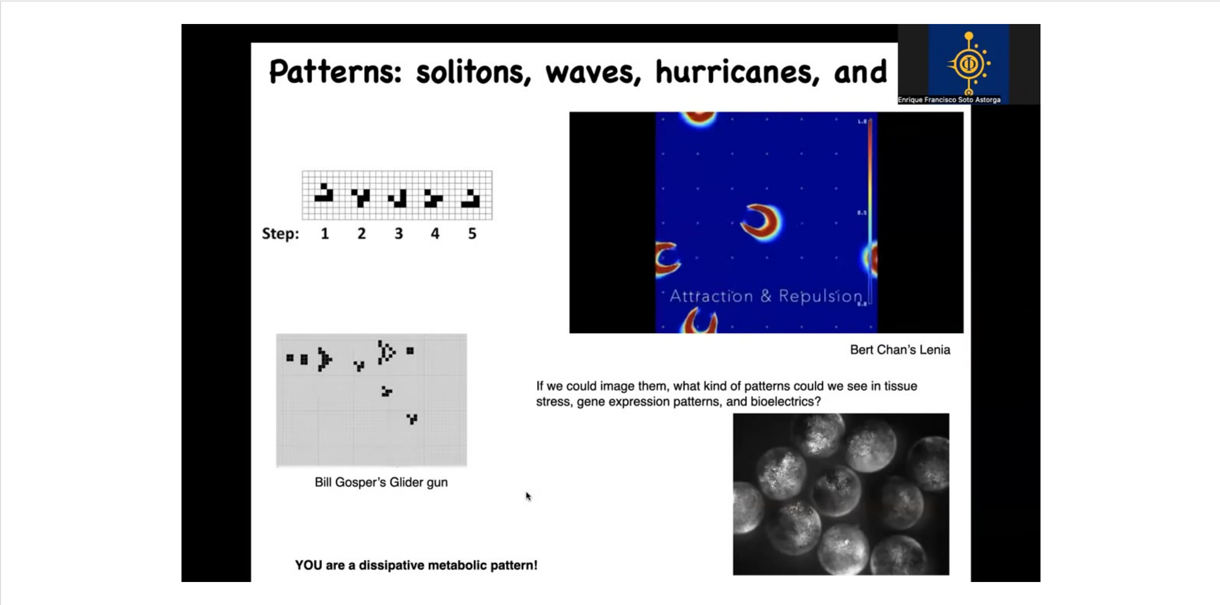

And I want to remind us all that we ourselves are a dissipative metabolic pattern, like hurricanes and gliders in the game of life and these patterns. And this is Lenia, which is this amazing system. And it has these, none of these things are objects in the permanent sense of the word. In fact, there are no permanent biological objects. All living things are patterns within some excitable medium that hang together for some amount of time.

And we have many patterns within our bodies, and we have many patterns across our bodies, as you can see here. Now that raises an interesting question. If we are patterns too, so we as objects that are thinkers, that are an excitable medium that can have patterns in it that serve as thoughts, as cognitive content, what other patterns that we normally don't recognize might also be agents in that sense. They might also have some degree of goal-directed behavior by being patterns in a different excitable medium.

All of this is designed to widen our view. I think the biology is really good at reminding us how our views on these kinds of things are extremely narrow.

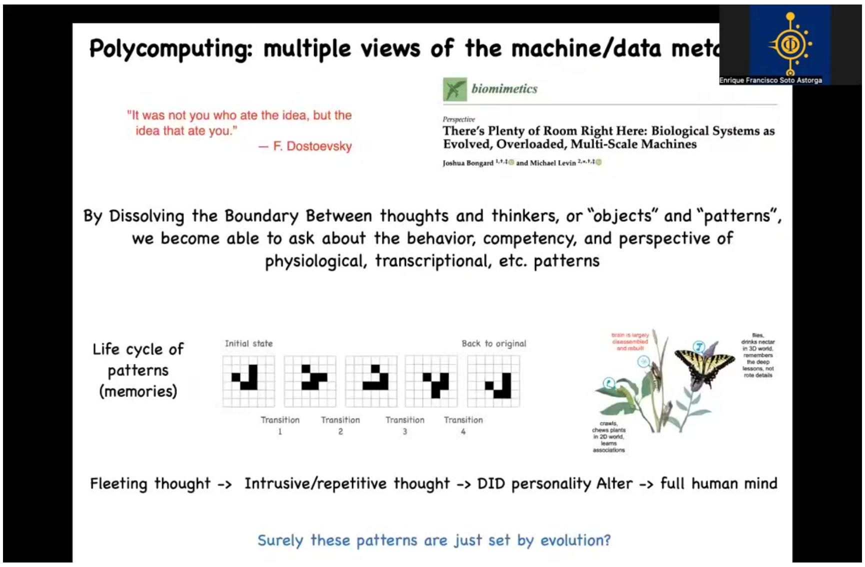

This idea that's developed here with Josh Bongaard of polycomputing, that there are multiple observers that interpret the same physical events in different ways, is really interesting. Some of these patterns have life cycles. Here's the life cycle of a glider in the Game of Life. Other patterns such as the memories of this caterpillar have a different life cycle. It has to adapt in order to survive this process. It has to change, much like this has to change in order to persist. You can think of the butterfly as the agent that has to remap the memories of the caterpillar, or you can think of the memories of the caterpillar as adapting and changing themselves in order to survive within the butterfly. By dissolving this fixed boundary between thoughts and thinkers, or objects and patterns, or data and machines, we become able to ask some really interesting questions about some unconventional agents.

You can think about a spectrum like this. There are patterns that come and go, like a fleeting thought through your mind. There are some that hang around longer. There are intrusive and repetitive thoughts that hang around. Not only do they hang around, they do niche construction. They alter your synapses to make it easier for you to have those same kinds of thoughts. Further down, you might have a dissociative personality alter, which is not quite a full human mind, but it's a lot more than a repetitive thought. It absolutely has its own goals and different kinds of intelligence and different agendas that may be quite distinct from other personalities in the same hardware of the brain.

I hope we're now starting to see that biology is showing us that by having this continuous self-interpretation process, it reveals an amazing kind of computational system where it's not clear who the data is, which is the machine, and where the agency lies.

Let’s shift to the final part of the talk, which has to do with where these patterns come from. This is maybe the weirdest part of the talk, depending on your philosophical leanings.

We have some patterns. We have morphogenetic patterns here. We have thought patterns and different types of cognitive patterns that different creatures are able to have. Where do they come from? Surely they’re set by evolution. They’re the product of mutation and selection for specific functionality. All of these patterns — physical, mental, physiological, transcriptional — must be set by evolution.

Let’s think about this. What is actually going on here? I’ll show you the implications of the fact that all that plasticity — the fact that embryos cannot do the same thing every single time — is that they must solve problems from scratch. They have to interpret both their genetic information and their behavioral information on the fly. That creativity — what does that mean? One thing it means is that it makes life extremely interoperable.

Slide 43 of 54 · Watch at [56:37]()

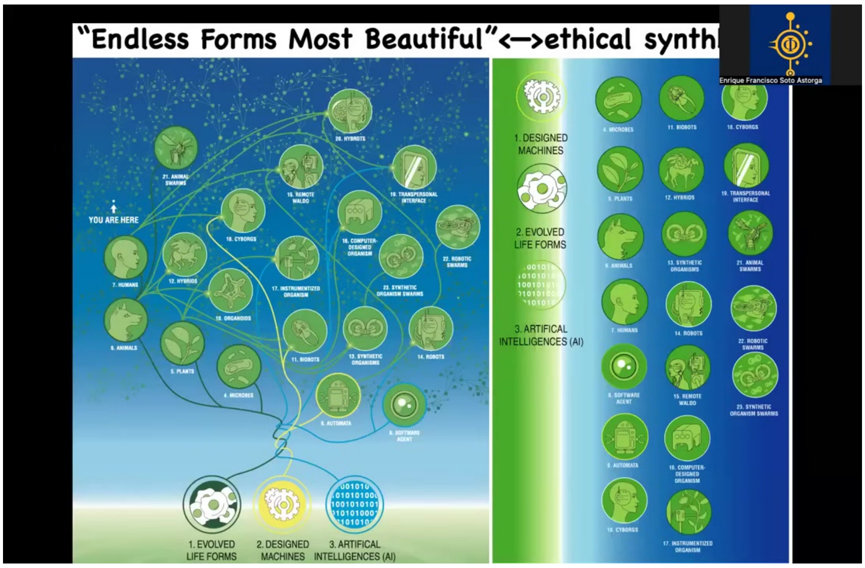

Almost any combination of evolved material, engineered material, and software is some kind of active agent. There are cyborgs and hybrots; many of these already exist.

In fact, when Darwin used the phrase ”endless forms most beautiful,” he was talking about the kind of variety of life and mind on Earth. Everything he was talking about is like a tiny dot in this incredibly large option space of bodies and minds because of that continuous self-interpretation aspect of life. The plasticity is massive, and you can make all these different combinations.

We are in the coming decades, and you, the young people in this audience especially, are going to be living with other beings that are nowhere on the tree of life with you. In other words, the old questions of what are you made of and how did you get here, meaning designed or evolved, are not going to be useful guides to ethical behavior.

We must understand how to enter a new kind of synthbiosis. That's a word made up for me by GPT. I asked it for a good term that encompasses the idea of living together in harmony and mutual benefit with these novel creatures, and the proposed synthbiosis is pretty good.

So all of these are now possible beings.

Slide 44 of 54 · Watch at [58:05]

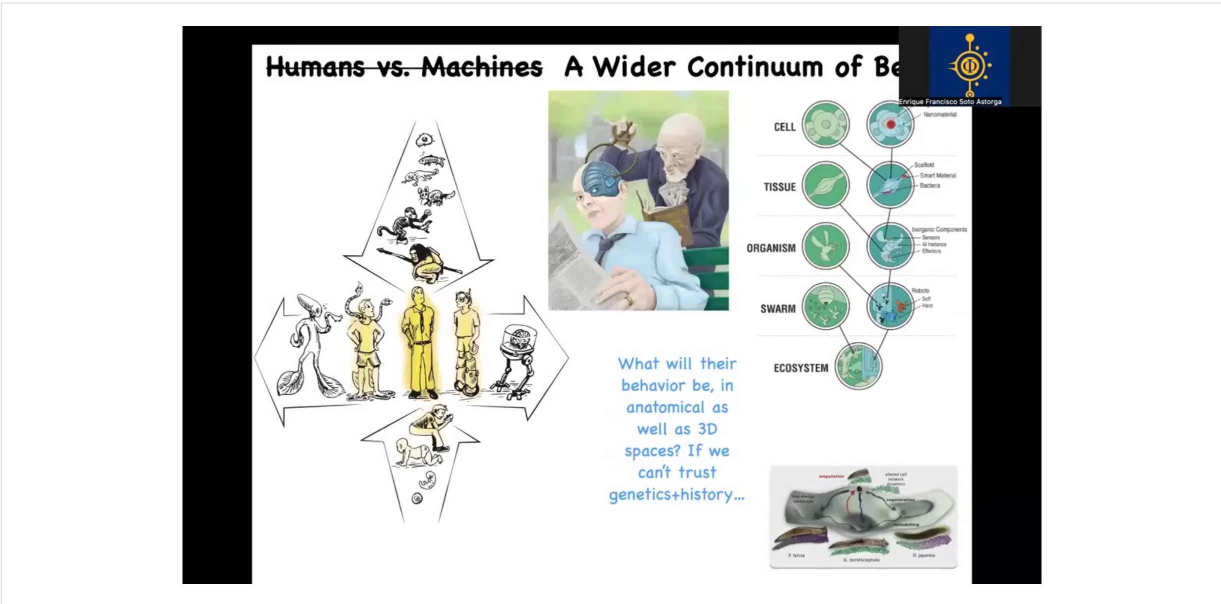

We are at the epicenter of two continua, not only the biological one here, but this modification where technologically and biologically, this form can slowly change into all kinds of different things. The definitions of humans and organisms and "quote unquote" machines and all of this are not natural kinds. These categories are going to dissolve. They were created at a time when our imagination and technology was insufficient to understand what's going on here. But we now understand that these are all continua and that it doesn't make sense to try to figure out if somebody's 51% biological or to categorize what kind of being they are. It's not going to work.

If we don't share a history, an evolutionary history with these beings, and we can't use genetics and history as a guide, how are we going to figure out what their properties are going to be. What are their behaviors in anatomical and conventional behavioral spaces? If they don't come from evolution, where do they come from?

I want to take a step back, because this is a question that people always ask: where do these patterns come from? I want to think about what we mean by that question.

This is the pattern right here, and these are just some videos made by modifying this equation. This is a Halley plot, a kind of fractal that's produced by plotting this very simple equation in complex numbers. That's it. Just six or seven characters.

Within this tiny little seed hides this incredibly rich pattern. It doesn't hurt that it looks very biological, but my point mainly is that if you were to ask, where does this come from? There is no genetics, there is no evolutionary selection. There is no law of physics. There's nothing about the physical universe that we know of that determines this. It comes purely from the laws of mathematics governing the behavior of complex numbers in the plane. It's remarkable.

We can't pin this on anything in the physical world. We're starting to understand that when we ask where patterns come from, some kind of story about physical history or physical objects may not necessarily be the entire story.

This suggests a prediction. It suggests that we ought to be able to make interesting systems that have no history in the physical world as such. I'm going to show you two quick examples.

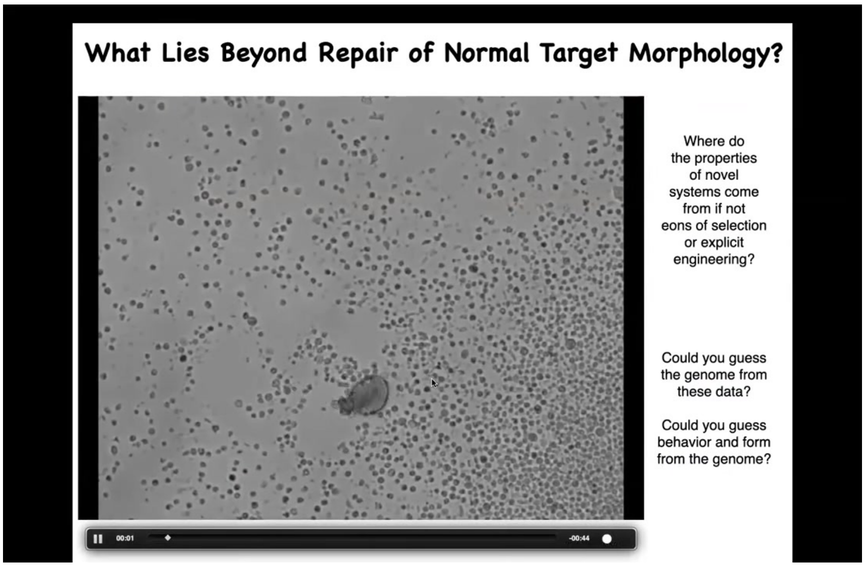

This is a little creature. If I showed you this, I would ask you to guess what this is. You might think that this is something we got out of the bottom of a pond somewhere. I would ask you to guess the genome and you might choose something that's appropriate to those kinds of primitive organisms. And then I would tell you that this is actually 100% Homo sapiens. This is 100% human genome. So not only can you not really guess the genome from the shape and behavior of an organism, but actually, if you have the human genome, I don't think you could ever say what would appear, and you would never predict this. This is not a stage of any kind of human development. This is not a human embryo. This comes from adult human cells, and it's called an anthrobot, which we allow to self-assemble from donated elderly patients' tracheal epithelium.

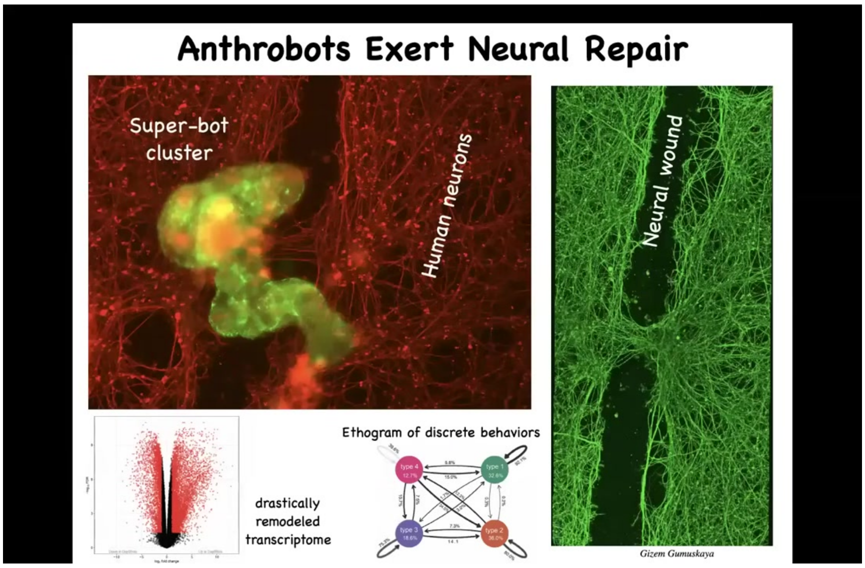

They have amazing properties. One property that they have is that even though their genome is completely normal, there are no synthetic biology circuits, there are no scaffolds, there are no weird drugs. All they have is a different life cycle. All of these dots are altered gene expressions. So about 9,000 genes that are altered, about half the genome is completely rewired. They have four different types of behaviors. Here's an ethogram of how the behaviors convert to each other when they're in a free environment. And they do this amazing thing where if you let them loose on a dish of human neurons in which you've put a big scratch through the middle here, they will assemble into this superbot looking thing and they will start to heal together here. Once you lift it up, they heal together the two sides of that neural wound. Now, who would have thought that your tracheal cells, which sit there quietly in your airway for long periods of time and basically just waft the particles up out of your lungs, are capable of self-assembling into a tiny little self-motile organism that has the ability to go around and heal neural wounds. These are all completely unpredicted forms and behaviors. There has never been any anthrobots in evolution. There has never been any selection to be a good anthrobot. These situations have never existed before, and yet out of the box, here they are.

The same is true for xenobots. So these are our frog biobots made of frog epithelial cells.

If you give them some construction materials, these are all skin cells, they implement von Neumann's dream, which is a machine that makes copies of itself from materials in the outside environment. So here they are. They form these little balls.

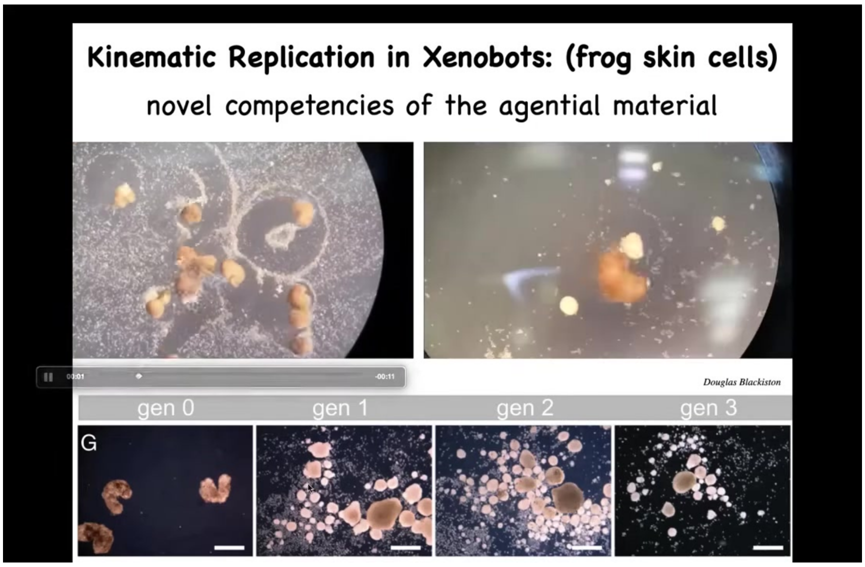

The little balls mature and become the next generation of xenobots. And guess what they do? They make the next generation. So this is called kinematic self-replication. Again, this never existed on Earth before.

Xenobots were never selected to be good xenobots.

This is something that happens.

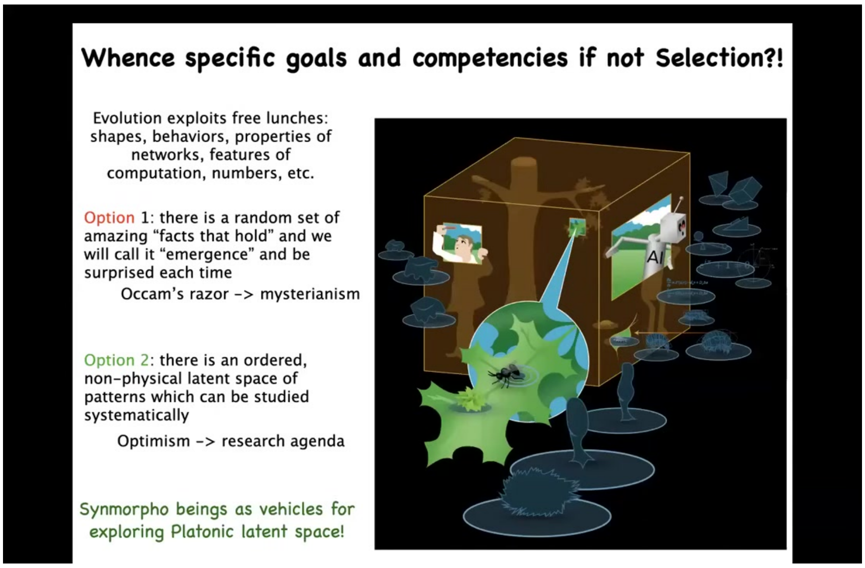

We have an option. We are starting to see that genetics and history are not the only source of these amazing patterns. Evolution exploits all this. Shapes, behaviors, and properties of networks, properties of numbers of computational systems.

There are two options here. One is that we can call them emergent, and we can assume that there's a random set of these emergent things that we will come across from time to time, and we will write them down and catalog them. This has the benefit of Occam's razor in that it's simple enough. But it's also a very mysterian approach because it suggests that we're going to commit to being surprised by these things when they happen.

Option 2 is more unusual, certainly for molecular and developmental biologists. Although the mathematicians have been on top of this for a really long time, a different assumption is that there's an ordered non-physical latent space of patterns, which can be studied systematically.

Not just a grab bag of random, amazing facts that show up, the distribution of primes and the truths of number theory and all these kinds of things. They don't just show up from time to time. There's actually a space that we can explore.

Then what we can do is view the kind of creatures that I was showing you, the anthrobots and the xenobots and AIs and so on, as vehicles or as a way to explore the different regions of that latent space and see what's actually in it.

This is what I think is going on. I think we're actually exploring a pre-existing space.

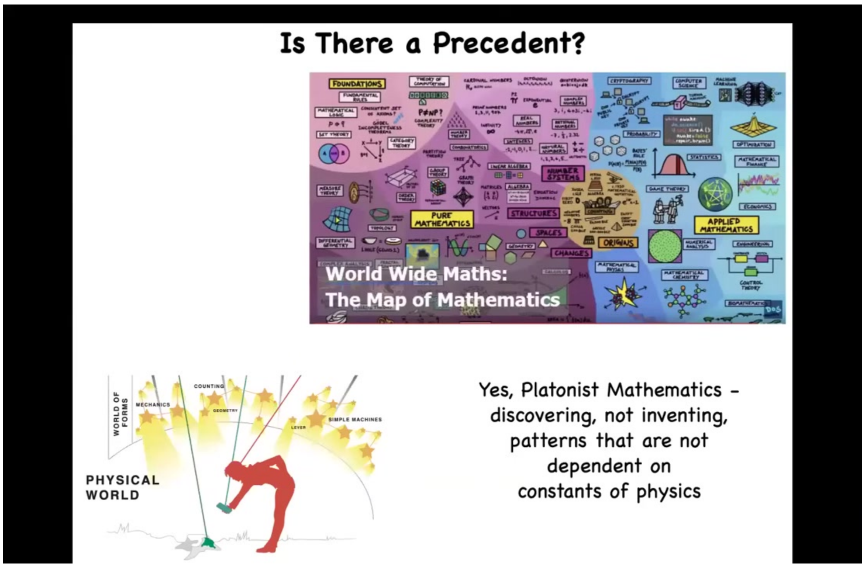

There is a precedent for this, which is many mathematicians have this Platonic view where they don't think they're inventing mathematics, they think they're discovering it. And by discovering it, they make this map of mathematics. It's not just random collection of truths. There's regions, there's the topology and the number theory and these other things that are related to each other. It has a metric. They can smoothly explore the space. And all of these patterns are not dependent on physics in the sense that they would all be the same, even if the parameters of the Big Bang were all different. So this is already an existing view. Plato and Pythagoras already expounded this idea that what the physical world is doing is implementing these particular patterns. What I would say is that what evolution is doing is exploring the space of pointers into this space. That physical objects, be they simple machines or biologicals or hybrids, are pointers into that Platonic space of forms and evolution exploits this readily.

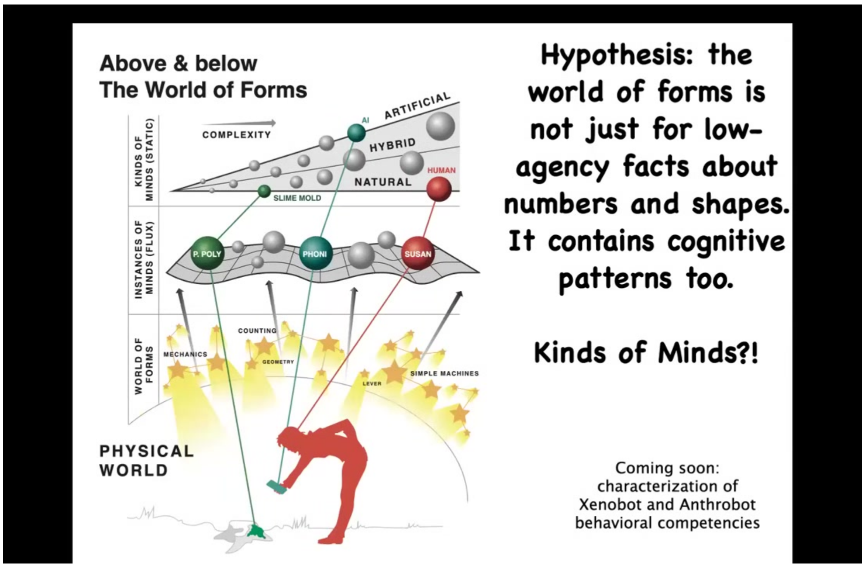

But we can go one step further and say, what if the platonic space isn't filled with only these low-agency, boring things such as triangles and prime numbers. What if behavioral patterns, meaning kinds of minds in Dennett's terminology, are also in this space? That is, you have these very low-agency things that we see in physics, but also maybe patterns of behavior and morphology, which is where the information from these novel systems comes from, where they don't have an evolutionary past.

We're doing a lot of work now on characterization of problem solving and behavior in Xenobots and Anthrobots. Stay tuned for that.

My bizarre conjecture is that I think we are already committed in many of the sciences to the idea that there is a space of forms that can contribute to the control in the physical world. We don't want to simply be surprised by them. We want to explore them. Perhaps, as Whitehead said, what you're seeing is the ingression of these patterns, both simple patterns and kinds of minds, into the physical world when we make appropriate portals. That would be embryos and biobots and all of these things that I showed you.

The last thing that I'll mention is that it doesn't take complex machines or large language models or embryos or cells before you start seeing surprising ingressions. I don't have time now to go through all this, but I will just point out this paper where we took a classic sorting algorithm, bubble sort, a few lines of code, perfectly deterministic, totally transparent, no new mechanisms to be found. We study them using the tools of behavior, and we confront them with problems they haven't seen before. We find out that even something as simple as bubble sort has interesting properties, delayed gratification and the ability to take on various side quests during its sorting task that we did not put in the algorithm. Not just complexity, not just unpredictability, but actually problem-solving behavior and novel goals in a very simple, minimal system.

I've talked to many people who will say, "Well, I build language models, there's nothing emergent there. I know I built it from scratch, linear algebra, and that's all that's in there." My point is, if we don't even know what bubble sort is doing, we have to have a lot of humility about what's going on in something as complex as a software AI. My claim is not that it's going to be cognitive and it's not going to have its intelligence because of the algorithm. My point is that for the same reason that biochemistry doesn't tell the whole story of the human mind, the algorithm and the mechanisms at the bottom of these computational systems do not tell the whole story of what they're capable of. These ingressions are surprising. We're very bad at being able to predict them. We have a lot of surprises left in store for us.

Slide 51 of 54 · Watch at 1:09:50

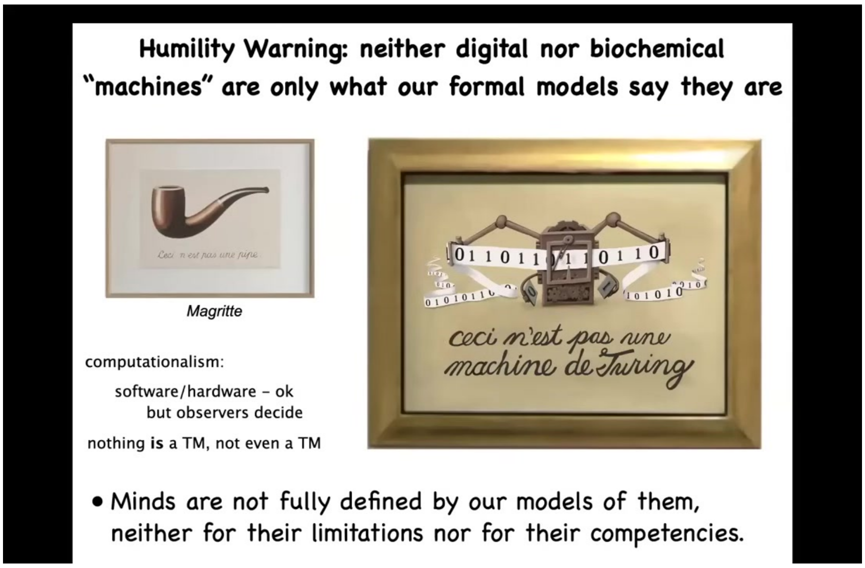

And so for this reason, I want to make this final point: Magritte drew a pipe and wrote in French, "This is not a pipe." This is not a pipe; this is a picture of a pipe. We need to remind ourselves that our formalisms have limitations. This includes computationalism: whether our computational models do or do not describe biology, I don't even think they describe physical machines. We have to be clear that our formalisms have limitations and various properties, but those are not the things themselves. Those are just the formalisms. And even something like this, which looks like a Turing machine, is in fact not the formal model of a Turing machine. It's something quite different. We're going to have lots of surprises. So again, I think neither living things nor "machines" are what our formal models say they are. We need to understand better this input of these ingressions.

Slide 52 of 54 · Watch at [1:10:52]

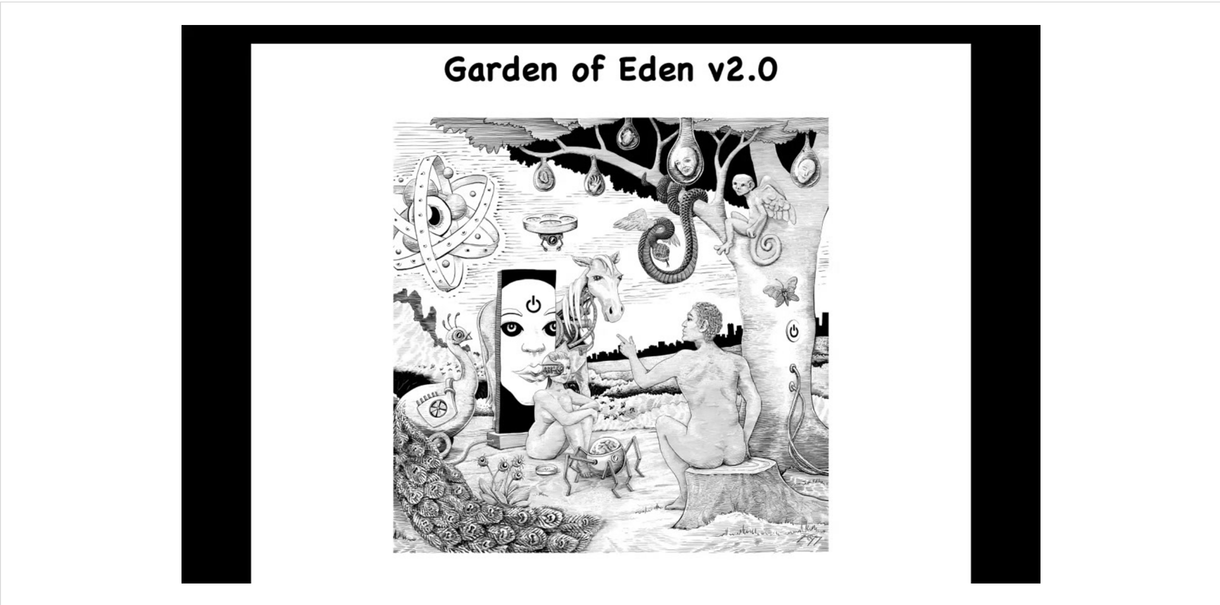

For that reason, I think the new version of Adam and the Garden of Eden is going to look something like this. We as humans are going to have to learn to discover the true nature of a wide range of beings that are alien to us now. We have to rise beyond our own evolutionary firmware to be able to recognize intelligence and ethically relate to all these beings.

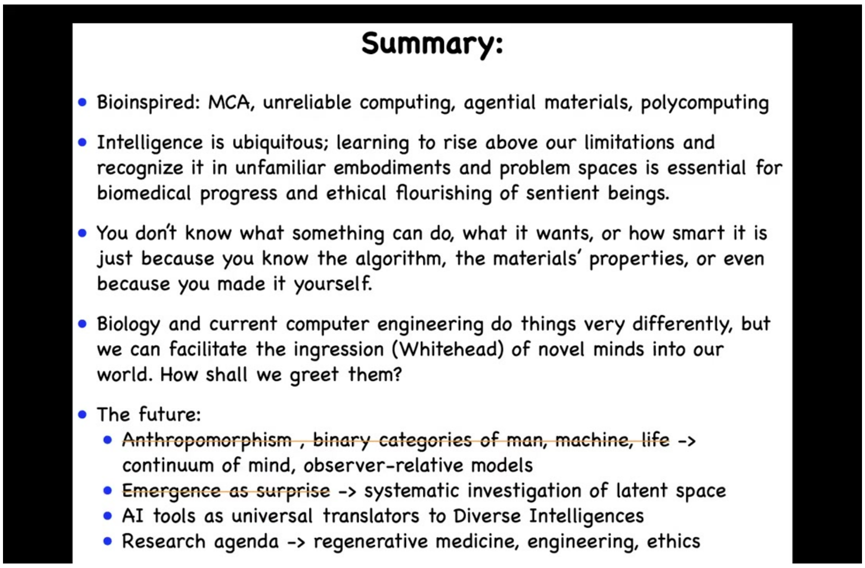

To really do bio-inspired engineering, we have to implement multi-scale competency, unreliable computing, and be able to work with an agential material where everything is interpreting everything else in different ways. Intelligence is ubiquitous. We do not know what something can do, what it wants, or how smart it is, just because we know the materials, the algorithm, or because we made it ourselves. Biology and current computer engineering do things very differently, but they don't have to.

We can certainly take advantage of this knowledge to build new structures, but we must remember that, just like when we make embryos, we have a responsibility because we are ingressing new minds into the physical world, maybe very alien ones that we don't know.

The future is all going to be about the continuum of mind. A systematic investigation of this latent space and perhaps AI tools as universal translators to this diverse intelligence. There's a huge research agenda here across regenerative medicine, which I didn't even talk about, engineering, and ethics.

There's lots of papers. Here are a few, if anybody's interested in diving more deeply.

I want to thank all the people who did the work. Here are all the postdocs and graduate students who did the experiments that I showed you today. Lots of technical support. We have some amazing collaborators who work with us. Here are our funders whom I need to thank.

Disclosures: there are three companies that fund some of this work. Here they are.

Most of all, the various model systems that we work with are the ones who teach us the most. Thank you very much. I will stop there.

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.