A Lecture Companion

″The Mind of the Body: A Window into Embodiment and our Future″

Michael Levin

Recorded on August 3, 2024

About this document

This document is a companion to the recorded lecture *"The Mind of the Body: A Window into Embodiment and our Future"*, recorded on August 3, 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

"Metaphysics and the Matter with Things: Thinking with Iain McGilchrist was a collaborative conference put on by the Center for Process Studies (CPS) and the California Institute of Integral Studies (CIIS) in March of 2024. This three-day conference brought leading process thinkers across various disciplines, including physics, neuroscience, psychology, philosophy, and theology into critical dialogue with McGilchrist's work in a collegial effort to assess, question, extend, and apply it. For more information on the conference and to purchase recordings, please visit https://ctr4process.org/mcgilchrist-conference/"

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.

Slide 1 of 2 · Watch at 0:00

Hi, everyone. While they're getting the screen sharing set up, thank you so much for making the opportunity for me to speak to you remotely. I wish I were there in person. I have some things to talk to you about.

Slide 2 of 2 · Watch at 0:20

So what I'd love to do today is to talk about our work. And in 20 minutes, I can only give you the highlights of some things that I think are relevant.

What I'm interested in is embodied minds, and in particular, embodied minds beyond the ones that are really easy for us to recognize. As a model system, we use the collective intelligence of cells as they navigate the space of anatomical possibilities.

This gives us the ability to ask questions in a very serious, empirical way. What does the mind of the body think about? How do we communicate with it? How can we reset some of its priors and its beliefs?

This gives rise to a very practical research program in regenerative medicine: birth defects, injury, cancer.

There are a few main points which are related to Iain's work. I'm delighted to be able to talk about these things in this context. You'll see these crop up during the talk. But fundamentally, what I do is diverse intelligence research.

One of the things I'd like to introduce to start off is this old painting, this is Adam naming the animals in the Garden of Eden. One of the interesting things about this old story is that it was Adam that had to name the animals. It wasn't God or the angels, it was Adam that had to do it. We wonder why was it on him to name the animals?

In many traditions, discovering the name of something or naming something means that you've understood its true nature, you've discovered its inner nature. This is exactly what we're doing because it's going to become very important to understand the true nature of all kinds of beings beyond what is shown here and what's conventionally thought about.

My framework is, the goal of it is to recognize and learn to ethically relate to all kinds of diverse intelligences. That includes, of course, us and various animals that are easy

for us to recognize, but also all sorts of weird creatures, colonial organisms and swarms, synthetic biology, engineered new life forms, chimeric life forms, things like cyborgs and hybrids and various other beings that are going to be with us in the next few decades, starting off with humans with various augmented technology and radiating out from there. Eventually, maybe even truly alien life.

I think that it might be shocking to see some of these things in the same list as some of these things. We really have to take seriously the humility of not assuming that we understand what kinds of embodiments various aspects of the ineffable can take. My point here is not to neglect the things that make us special, but actually to really understand what other embodiments they can take.

Of course, I'm not the first person to think about these things. Rosenblueth, Wiener, and Bigelow tried for this scale so that you can put on the same continuum all kinds of different degrees of agency. What I'm interested in is a framework that moves experimental work forward to new capabilities, and also enables better ethical frameworks that are going to avoid some of the lapses that I'm sure will result from naive and simplistic ways of trying to make a distinction between all of these different things.

This idea of making distinctions between agents in terms of where they come from—where they evolved in a random search process of evolution, where they were engineered by human engineers, or somewhere in between or a combination. Let's look at where we actually come from.

This is the journey that each of us takes. We were all, quote unquote, just physics at some point. I hate that phrase, but people use it a lot.

We all begin life as an unfertilized oocyte, a blob of chemistry. And slowly, gradually.

Developmental biology gives you no specific point at which something unique happens to turn this process into this or one of these other amazing creatures. We all have made this journey. If you look inside any one of these things, for example, Descartes was really interested in the pineal gland because it's unitary in the brain. He felt that was appropriate for the common unitary feeling of the integrated feeling of the human experience. But if you look in the pineal gland, you see a bunch of cells. If you look inside each of those, you see all this stuff. There's an amazing depth of mechanism going all the way down to molecular networks. This happens to be a free-living organism, but this is the kind of thing we are made of. We are made of an agential material. We are all a kind of collective intelligence because the substrate that we are made of is not a passive material. It actually has agendas; it has the ability to form memories, it can learn, it has preferences and so on. This is our architecture.

Now, when we start life, this is an embryonic blastomere, blastoderm, at least a small view of it. Let's say we have 50,000 cells and we look at those 50,000 cells and we say, okay, there's an embryo, one embryo. What are we counting? What are we really counting when we look at 50,000 cells and we say that that's an embryo? Well, what

we're really counting is, first of all, alignment, functional alignment towards a goal. What we're saying is that all of these cells are going to work together to make a very particular journey in anatomical space, starting with one cell and ending with a normal, let's say, human in this case, target morphology. Fundamentally, what we're counting here is the commitment to the same story. All of these cells are committed to the same set of preferences and goals in anatomical space, and that's why they're all going to work towards making this.

This is not simply a mechanical process of complexity emerging from simple rules, because we know there are many examples that show that if you try to deviate them from this path, they are very good at finding new ways to get there.

One of the key things is that there are lots of similarities here between the construction of the body and the construction of the mind. If you didn't know what a human was and somebody showed you the material of a brain and asked, "How many cells can actually fit in here? What's the density of cells per medium?" we actually don't know. We don't have a way of knowing that. It isn't genetically determined in the case of development either because if you take this blastoderm, let's say a chicken or a duck egg, and you use a little needle to make some scratches here in this blastoderm, for the few hours before these things heal up, each of these islands, not being able to feel the presence of the others, is going to self-organize into an embryo.

Each one of them is going to be its own self. You can make twins, triplets, probably up to half a dozen or so. There are a couple of interesting things you see. First of all, the number of individuals that come from this blastoderm, and this happens in humans too, whether conjoined or not, is not known. It is not fixed in advance by genetics. Most of the time it's one, but it could be anywhere from zero to half a dozen or more. The next thing that happens is that each of these embryos commits to a particular path in anatomical space. The cells here in between have to decide, am I part of this embryo or am I part of that embryo? There's a scale-up going on here in which individual cells will bind into a collective, which will go off and do specific things. The cells at the border have a difficult time.

Some of them have problems deciding, am I the left side of embryo 2, or am I the right side of embryo 1? For this reason, it's been long known that conjoint human twins will exhibit laterality defects, so left-right inversions of the heart, the gut, in one of the twins, because some of these cells have a hard time figuring out where they are. This is the process of coming into the world as a self, the first thing you have to do is figure out where your borders are. Where do you end and the outside world begin? This has many implications in neuroscience as well, issues in individuation for dissociative disorders. One of the first problems that has to be solved is the problem of understanding left versus right.

Now even individual cells: here's a chiral paramecium, here's a slime mold, this whole thing is one cell and they have a preferential rightward bias when they hit a T maze, they have a built-in asymmetry. Very interesting because macroscopically the universe

does not distinguish left from right. At the quantum level it does, but macroscopically it does not.

All of the cells in the embryo have to do a very interesting thing. They all know which direction is, let's say, rightward. All the cells have that same information that way is rightward. But they have to convert this into knowing where relative to the midline they are. This is a chick embryo. Here's the midline. The cells here on the left side have turned on a particular gene called nodal because they already know that they are on the left side of the midline relative to the other cells in the midline of this body.

The way that happens is that from the very earliest moments of fertilization, in fact at the first cleavage, there is particular bioelectrical machinery: proteins that create an electrical gradient migrate to one side; they all go rightward. This sets up a voltage gradient, which we can measure, and that voltage gradient yanks along something called serotonin, which is a neurotransmitter.

One reason this might be interesting in this context, and related to some of Iain's work, is that right from the earliest moments of development you see questions of laterality, questions of asymmetry, and the use of a neurotransmitter to solve these problems long before neurons or the brain shows up. In this case, bioelectricity, by virtue of its control of neurotransmitter signaling, is involved in the collective decision-making that's required to set up an embryo as distinct from individual cells.

Now, these collectives have all kinds of interesting problem-solving properties. First of all, upon injury: this is a salamander—if they lose their arm, they will regenerate, and they will make exactly what's needed, and then they stop, no matter where they were amputated. That's the most amazing part of regeneration: they know when to stop. The same thing happens in these tadpoles that have to rearrange their face to be a frog. If you create a so-called Picasso tadpole where everything is in the wrong place—the eyes on the back of the head, the mouth is off to the side—they will actually rearrange in novel ways and give you a very normal-looking frog.

What the genetics does not give you in any of these cases is a hardwired machine that always does the same thing under all circumstances. There may be a few species that do that, but by and large that's not what you get. What you get is a highly flexible error-minimization scheme. You get something that's able to correct despite novel circumstances, which is William James's definition of intelligence, the ability to hit the same goals by different means.

Highlight the incredible problem that these organisms solve. This is a cross-section through the kidney lumen of a newt, and you can see that there's about eight to ten cells that normally are used to make up this structure. There's a trick you can use in the early egg where you allow the chromosomes to multiply, but the cell itself does not divide. They end up with lots more DNA.

In that case, the cell itself will get bigger. If that happens, what you see is that the newt is exactly the same size. It just uses fewer cells to do the exact same job.

Now, what's really remarkable is that if you make the cells truly gigantic, then something new happens. Instead of using cell-to-cell communication here, it uses a different molecular mechanism, cytoskeletal bending, and each cell will bend around itself, leaving an empty spot inside. Coming into the world, you don't know how many cells you're going to have. You don't know what the size of your cells are going to be. You don't know how many copies of your genome you're going to have. You can't rely on any of this stuff. All of these things that supposedly were honed by evolution, you can't assume any of that. You have to be able to, in this kind of creative problem solving, deploy new molecular mechanisms that are at your disposal to get the job done, to complete that journey in Morphospace, despite the fact that it's not just external injury, it's not just that the environment's changed, your own parts have changed. You can't even rely on the consistency of your own parts. This ends up being actually a very interesting property that I think is the ladder for intelligence. It's kind of a ratchet for intelligence in biology.

So how does all this work? Years ago, we started to take inspiration from neuroscience, which is used to these multi-scale problems and understands the scaling of information from very small components all the way up through high-level cognitive functions. We know that there is this bioelectrical machinery that sets voltage states that propagate through networks, and you can try to decode those electrical states here. This is a living zebrafish brain. You can see all the physiology that's going on here as this fish is thinking about whatever it is that fish think about. What we realized is that this is actually not unique to neurons or brains. Every cell in your body does this. Evolution discovered this many, many eons ago. This is extremely early, about the time of bacterial biofilms. It discovered that electricity was a great way to scale minds from tiny little microbial components all the way up through larger kinds of things. Every cell in your body has this electrical signaling. Most of them become coupled into networks with electrical synapses. We can ask the same question. We know what these networks think about. They think about moving you through three-dimensional space, for example. What do these networks think about? They think about moving you through anatomical space.

You can see what this looks like. This is a voltage-sensitive fluorescent dye technique that we developed. This is a frog embryo that's about to put its face together. Long before the genes come on to regionalize the face, you can see here, you're now reading the mind of this collective intelligence. You can see this is what it's going to build. Here's the electric face. I'm showing this one to you because many of the patterns are very hard to decode. This one's easy. It looks like a face. Here's where the animal's right eye is going to be. Here are the placodes. Here's the mouth. This is what's going to guide the gene expression and the subsequent anatomy. There are pathological patterns. If a human oncogene is injected that will eventually make a tumor, you can already detect at an early stage these are the cells that are going to disconnect from the environment, go off like amoebas, treat the rest of the body as just external environment. We'll talk about that in a minute. These bioelectrical patterns are literally

the memory in the excitable medium of this collective intelligence that control where they go in anatomical space.

We've learned partially to communicate with them and convince them to do other things. Here, by injecting mRNA and coding a potassium channel, this is something that's going to change the voltage of the cell. There are no fields, no waves, no magnets, no electromagnetic radiation. That's not how we do it.

We use the native interface that the cells are using to hack each other during embryogenesis. We're taking advantage of it. What we can do is we can instill in these cells a particular voltage state that says to them, make an eye. If you do that, in cells that normally will give rise to the gut, here, so this is a tadpole, here's the native eye, here's the gut. Those cells will happily build an eye and that eye will have all the lens, retina, optic nerve, all the same stuff.

We learned a few things here. We learned that the bioelectric memories are instructive. They tell the organs where to go. They're very modular. They're top-down control. We didn't tell these cells how to build an eye. They already know how to build an eye. All we said, we just gave them enough information to say, go ahead and do that. We didn't give all the information on all the strategies. We're not talking to the individual stem cells. We're not micromanaging the gene expression, nothing like that.

And we discovered a counter to the developmental biology textbook, which will tell you that only these cells have the competency to become an eye, that if you prompt them in the correct way, pretty much all cells in the body have the competency to do that. That reminds us, when you're judging the competency of some system, you're really taking an IQ test yourself because you may just not have found the right way to interface with the capabilities of that system.

There are some very interesting dynamics that happen when you do this. The cells that we've injected are saying to their neighbors, let's build an eye. In particular, if there's only a few of them, they will sometimes convince all these other cells around to participate in the process. The blue ones are the ones we injected, but there's not enough of them. They know that. And so they activate all these others to work with them. The cells around them that are trying to suppress cancer are saying, you're crazy, don't be an eye, you should be skin. We're all going to be skin or gut. They have this back and forth conversation that we can now track that determines what happens.

We can rewrite this memory in creatures like planaria. These are flatworms. You can cut these flatworms into pieces. Every piece reliably knows what to make. There's this bioelectrical pattern memory that says one head, one tail. If we edit that memory and convince them to do this kind of pattern, if you cut this animal now, you will get a two-headed flatworm. This is not Photoshop or AI, these are real animals.

Very importantly, this is not a bioelectrical pattern of this two-headed animal. This is a bioelectrical pattern of this perfectly normal-looking one-headed animal whose anatomy is normal, whose molecular biology is normal. This is a counterfactual

memory. This is the early evolutionary version of our mental time travel because what this animal is doing is the body is storing a memory of what you will do if you get injured in the future. This is not what you're doing now. This is what you think a correct planarian looks like. The body can store at least one of two different representations of what a correct animal looks like, which is what you will use if you injure it.

I call this memory because once you've reset it, it stays. If you take these two-headed animals and continue to recut them, this is how they normally reproduce as they tear themselves in half and regenerate and go on. If you recut them, remove the primary head, remove this ectopic secondary head, the middle fragment is now forever convinced that normal planaria should have two heads. And we can actually reset it back to one. We have the ability to actually incept new memories into that anatomical collective intelligence and tell it where it should go if it gets in trouble. Here you can see the video. These are what these two-headed animals are like.

The last bit of this I want to tell you is a story about cancer. What evolution has done, remarkably, is use bioelectricity as a kind of cognitive glue to scale up the goals of these biological systems. Individual cells have little tiny goals. All they care about is the local physiological states here. But together, what evolution has allowed is for them to join into electrical networks that maintain very large-scale goals like building a limb. If you deviate them from that goal by injury, they will continue and rebuild and then they stop. That process has a failure mode and that failure mode is cancer.

What can happen is that if there's an oncogene, or many other things that can cause cancer, it will cause the cells to disconnect electrically from their neighbors. At that point, the rest of the body just becomes environment and they do what all single-celled organisms do. They reproduce as much as they can. They go where life is good; that's metastasis. These cancer cells are not more selfish than normal cells; their selves are smaller. The boundary between self and world, whereas before it was huge, probably the size of an organ or maybe the whole animal, is now literally shrunk to the size of a single cell, and it's still every man for himself, but now that border is tiny. Their model of where the self ends and the outside world begins has shrunk.

We can use that kind of thinking to design a therapeutic. What if we artificially force these cells to continue being part of that electrical network? We won't kill them. We won't turn off the oncogene. We will simply force them into electrical communication with their neighbors. If you do that, this is the same animal. Here's the oncoprotein blazingly expressed. There's no tumor. That's because we've injected an ion channel that controls the bioelectric potential. It doesn't fix the genomics, but that's OK. You don't need to.

I think that in the future, biomedicine is going to look a lot more like a kind of somatic psychiatry than it is like chemistry. This is what we do now, all of these bottom-up interventions, because we treat cells and tissues as simple mechanical machines. They're machines, but they're not simple mechanical machines. They're agential beings

with problem-solving capacities. We can interface with them using all kinds of interesting novel technologies that are now forthcoming.

What I've shown you in terms of regeneration and so on is the ability to restore and rebuild to the correct target morphology. What about the plasticity? How far away can we go from that? Normal development really puts us into a false sense of understanding because acorns always make oak trees and we think that we know what the acorn genome can do.

But if prompted by a non-human bioengineer, this wasp, these exact same cells with no genomic change could build something completely different. They can build these amazing galls. If not for this little guy, we would have had no idea that these cells are actually capable of doing that.

What we do in our lab is we'll borrow some skin cells from an early frog embryo. We put them in a little depression. They don't do any of the things you might assume they would do. They actually come together to be something we call a xenobot. These are autonomous novel kinds of constructs. They are a proto-organism. They swim along. They can go in circles. They can patrol back and forth. They have this group behavior.

Here's one navigating a maze. It swims along here. It's going to take this corner without having to bump into the other side first. It takes the corner. Then, at this point, it spontaneously turns around and goes back where it came from.

So we're in the process of studying their ability to learn their various behaviors. One thing they do is if you provide them with loose skin cells, they will run around and collect those skin cells into a little ball. Because they're dealing with an agential material, these are not passive pebbles, these are cells; those balls mature into the next generation of xenobots. They do exactly the same thing. So you get multiple generations. This is kinematic self-replication.

You can ask, what did the frog genome actually learn during its many years of evolution? Well, it certainly learned by default to do this. This is the normal developmental sequence. These are the behaviors. But it also can do this. This is a xenobot developmental sequence. I have no idea what it's trying to become. There's never been any xenobots before. There's never been any evolutionary selection to be a good xenobot. This is something completely novel that these cells can do to repurpose their hardware in a new way.

I want to point out that this is not some weird fluke of embryonic cells or amphibian cells. Look at this little proto-organism. If I showed you this, you might think that this came from the bottom of a pond somewhere. If you sequence this, what you will see is 100% Homo sapiens. These are human cells. These are adult patient tracheal epithelial cells donated by human patients during trachea biopsies. In a particular novel environment, they will self-assemble into this little thing that we've called an anthrobot.

Now, they have all kinds of cool properties. One of the coolest is that if you have some

neurons here growing in a dish and you make a scratch wound, they will navigate down that little scratch wound. When they settle into a group, this is a superbot cluster—there are probably 10 of them here. What they will do is sit there and knit the sides of the neural wound together. Four days later, you pick them up and you see that what they were doing here is healing this neuronal scratch wound culture. That's just the first thing we found, and who knows what else they're capable of.

Just imagine: who knew that your tracheal cells that sit there for decades quietly in your airway actually have the ability to make an autonomous multicellular little organism. It's never existed before. It's not how humans develop. It's not anything in our lineage that actually has the ability to run around your body and make these repairs.

What I think is that by making these kinds of things—I didn't even tell you the story where we can make planaria take on the head shapes of other species. The same hardware with no genetic modification visits attractors in that anatomical state space that belong to other species. I think all of these synthetic beings and everything else that are made here are vehicles for exploring a platonic space of affordances that go way beyond the one little pinhole that we normally see in normal development and normal regeneration. By using various techniques, we can actually start to map out the space and ask, if evolutionary selection is not what created these things, where are these patterns coming from? I think we now have a way of starting to understand this.

The very last thing I will say is that I'm very conscious of the fact that the talk that I gave was exclusively a left-brain talk. In the week leading up to this, I attempted a right-brain contribution. You can see here at this URL or here; I wrote a little bit of what this topic of diverse intelligence might look like if addressed from a more right-brain perspective.

I'll stop here and thank the postdocs and the grad students who do all the work. Thank our funders who support us. Some disclosures: a few biotech companies are supporting our work. And of course the animals, because they do all the hard work. Thank you so much.

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.