A Lecture Companion

"Aging, goal-directedness, and bioelectricity" by Michael Levin

Michael Levin

Recorded on June 26, 2026

About this document

This document is a companion to the recorded lecture *"Aging, goal-directedness, and bioelectricity" by Michael Levin*, recorded on June 26, 2026.

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

Lecture description

This is ~23 minute talk overviewing our recent results and the way we think about aging and longevity. It was for a conference with quite short speaking slots so I had to talk fast... Longer talk coming with more data but this outlines some ideas and overviews our published recent work on aging.

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 13 · Watch at 0:00

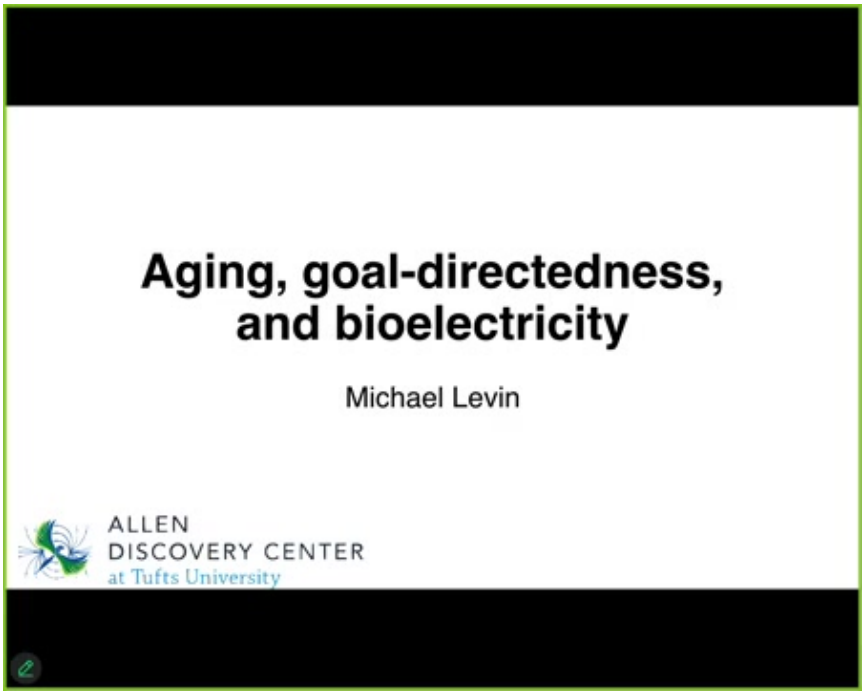

I would like to talk to you today about some of our work on aging at the intersection of the concepts of goal directedness in biological systems and bioelectricity.

So just to give you an overview of how we think about aging. There are two very popular ways to approach the fundamental cause of aging. There are your damage theories, right? So accumulation of thermodynamic noise and things like this. And that says that aging is fundamentally caused by a problem of physics. Then there are the programmatic theories that say that aging is mostly caused by a problem of biology, evolution, specifically. There are specific programs that want you dead. And we are going to introduce a new potential source of aging, which has its root neither in physics nor in biology, but actually in cognition. Or if you prefer, you can think about cybernetics.

So today I'm going to talk about data in three brief areas. First, this notion of atavistic dissociation and what happens to cells and tissues during aging. Upstream of that, I will talk about the degradation of instructive bioelectric patterns. And upstream of that, I will float a new idea for a potential ultimate cause of aging.

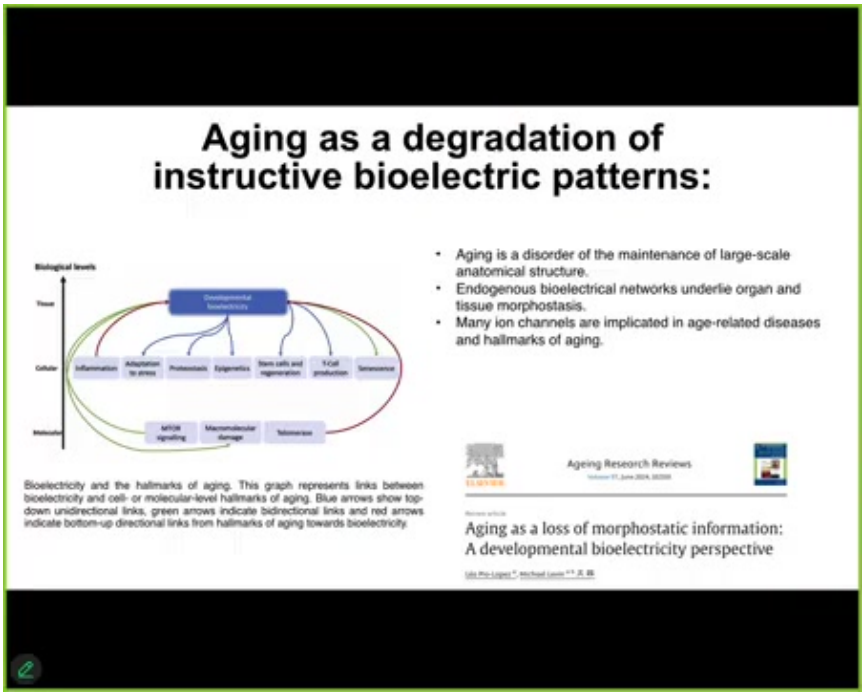

So some of this can be seen in this review. And basically, what I would like to start by focusing on is this notion of aging as a disorder of the maintenance of large-scale anatomical structure. And one of the things that helps maintain that large-scale anatomical structure over time is developmental bioelectricity. I'll talk about that momentarily.

Slide 3 of 13 · Watch at 1:38

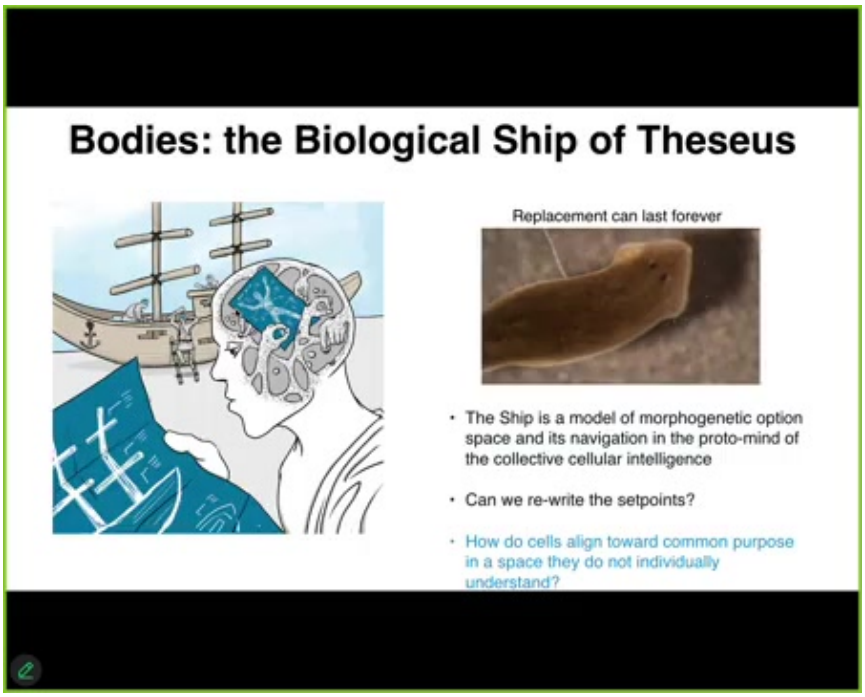

But the bottom line is that, as we all know, our bodies are a kind of ship of Theseus, right? So materials come and go, cells come and go. But what you would like to do for longevity is to have a large-scale coherent structure, the body, including the cognitive individual that inhabits it. You would like to have some continuity of that, even though the physical hardware is swapped out over time. And I think one of the more instructive ways to think about this whole ship of Theseus thing is that if we ask, where is the actual ship? I would propose that this is not the ship. The ship isn't the actual physical structure. The ship is the plan in the mind of the repair machinery, which in this case are human workers, but I'm about to tell you the same story in the context of cells within the body, that basically what's actually the ship is the information structures that allow the repair machinery to aim towards a specific thing. So that the ship of Theseus, as the wood, the planks are swapped out over time, it doesn't become a shed or a bicycle or anything like that, it remains the shed. In order for that to happen, the repair machinery has to have a very clear picture of what it is that they're actually working on. And we do have some species that I'll briefly mention, such as these planaria that apparently can do this trick, as far as we can tell, forever. The asexual version of these flatworms seem to be able to continue without any obvious limit. So what we want to understand is how exactly large-scale bodies are maintained? Because that would, how are they constructed and then maintained? Because that would let us intervene in disorders of this process, aka aging and cancer and some other things. Of course, as bioengineers, we're also interested to know whether we can rewrite the set points. Can you make the cells build something different? Again, that's

important because you may not just want to remain as status quo, you may want to introduce improvements, enhancements. So that's, I think, clearly a part of the longevity agenda. And fundamentally, we need to understand how is it that cells align towards common purpose in a space that they don't individually understand. So there's no individual cell that knows what an eye is or how big and how many of them you're supposed to have, but the collective absolutely does. And the reason you know it does is that in species such as these, in highly regenerative species, you see this example of anatomical homeostasis. So if you amputate, for example, a limb anywhere along the axis here, the cells jump into action, they rebuild very quickly, they replicate, they do all the things that are needed, and then they stop. And that's the most amazing part of this whole process. They know when to stop. When do they stop? Well, they stop when a correct structure has been achieved to some small tolerance of error. But basically, when they are deviated from the correct set point, they will take action to reduce the delta. They're an error minimization scheme. So, and this isn't just about damage. There are some remarkable examples like this where, and this is old work from the '50s, it's not our work. When you take a tail from one of these guys and graft it to the mid flank, what happens over time is it gets in place, it gets remodeled into a limb, and in particular the cells at the tip of the tail, which locally there's nothing wrong, there's no injury, there's no damage, but they become fingers. Why do they become fingers? Because the error is propagated from a very large-scale, very abstract information structure, which stores, roughly speaking, the information about what a complete body should look like. And a correct body of this amphibian does not have a tail in the middle here, it has a limb. And so that delta then propagates down into the cells in the molecular cascades below to execute the changes needed to reduce that error. Again, error minimization scheme. In this case, it's not just damage at the site of the missing tissue. It's the fact that local order obeys a global plan.

This is really critical towards making and maintaining bodies against all the weird things that can happen to them, not just aging, but damage and these kinds of things that scientists do in their experiments. So what's important here is that in these kinds of examples, we're looking at anatomical homeostasis towards a stored set point. It is, in fact, not unlike what we've often been told when we start to study biology. When you say, but how does it know how to do that? And people say, nothing here knows anything. It's chemistry. Chemistry doesn't know anything. Chemistry just rolls forward and eventually, in this kind of open-loop process, emergence is the way to understand these complex structures that happen. I think that's a very limiting perspective. There may be situations where that's good enough, but I don't think it's good enough for most regenerative medicine, and I don't think it's good enough for longevity. So we have to understand that these systems actually do have a stored set point that they're working towards, because that's a key to an important set of interventions. So how is that possible? How could a bunch of cells store a memory of what it is that they should build?

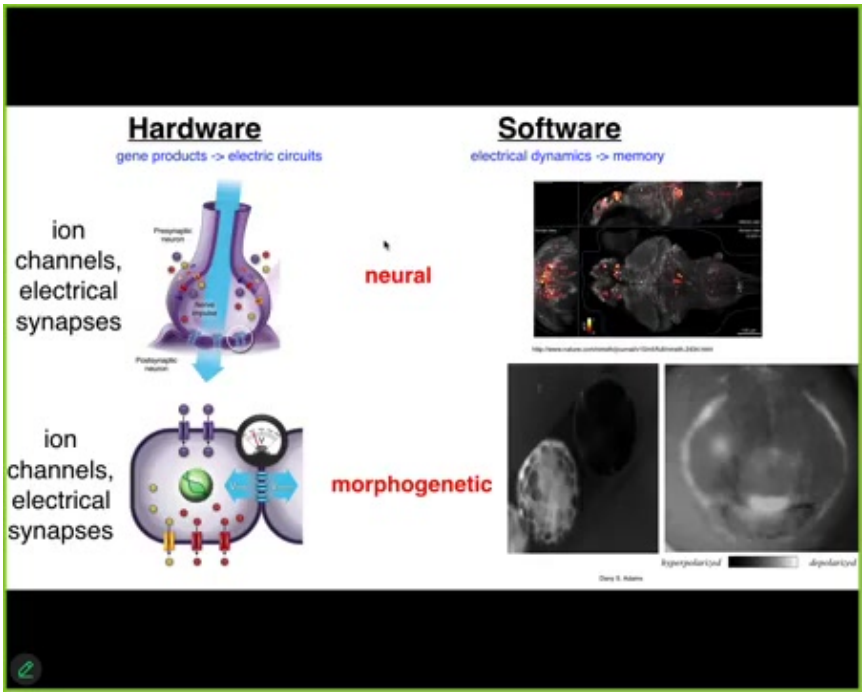

We already know from neuroscience one example of a collection of cells that remembers what it should do, and that's the brain. So in the brain, we have this group made this amazing video of live zebrafish and live zebrafish brain, and the commitment of neuroscience is that if you were to understand and decode this electrophysiological activity, you would be able to know the memories, the goals, the preferences, and so on of the animal and understand what it was capable of doing and why it does any particular thing. So it turns out that this scheme of storing goals for directed navigation of problem spaces is extremely ancient. It's not about neurons, it's not about brains. Basically, it was discovered by evolution around the time of bacterial biofilms. And since then, it's been used by every cell in your body. So every cell in your body has ion channels. Most cells have electrical synapses to their neighbors. And what they do is they end up making electrical networks. And what do these electrical networks think about? Well, the electrical networks of the brain think about moving you through three-dimensional space. The electrical networks of your body think about the shape that you should be and the journey that you took from being a single cell, a fertilized oocyte, to a mature target morphology. And this right here, we call this the electric face. It's an example. We created the first tools to read and write, the first molecular tools to read and write this kind of pattern memory in non-neural tissue. So when you look at a neural frog embryo and you want to know what the future face would look like, this is it. You can read out this pattern. You can see here's where the mouth is going to be, here's where the eye is going to be, here are the placodes. This is

a subtle scaffolding, informational, an informational structure that tells you what the future gene expression and what the future anatomy is going to be. Okay, and so...

Slide 5 of 13 · Watch at 8:40

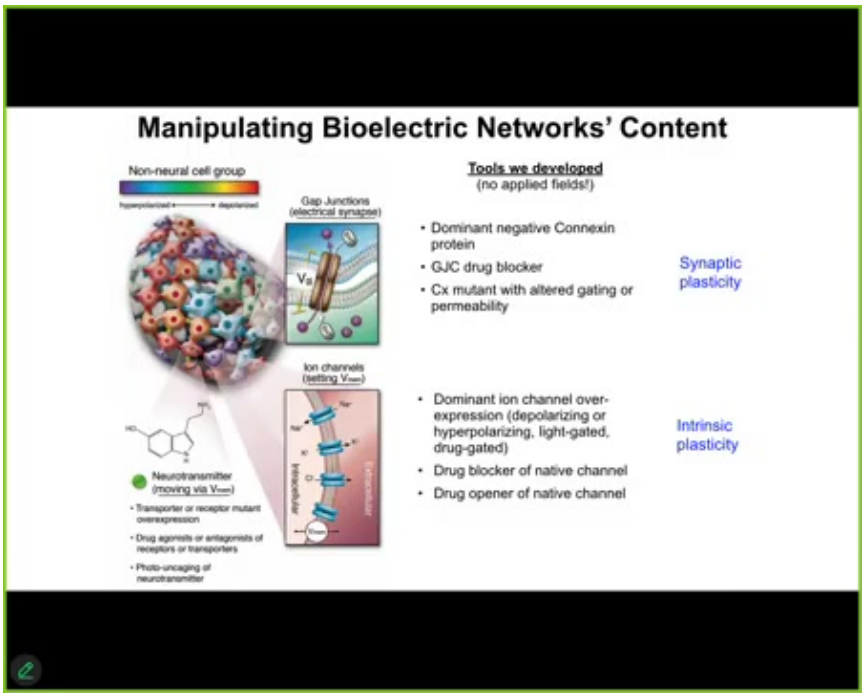

What we're able to do then is to manipulate, we've developed tools to manipulate this. We don't apply any electric fields or anything like that. We open and close ion channels and gap junctions using optogenetics, using ion channel drugs, and so on, in accordance with computational models that we develop. And so what happens, what can you do when you rewrite these bioelectric pattern memories?

Slide 6 of 13 · Watch at 9:03

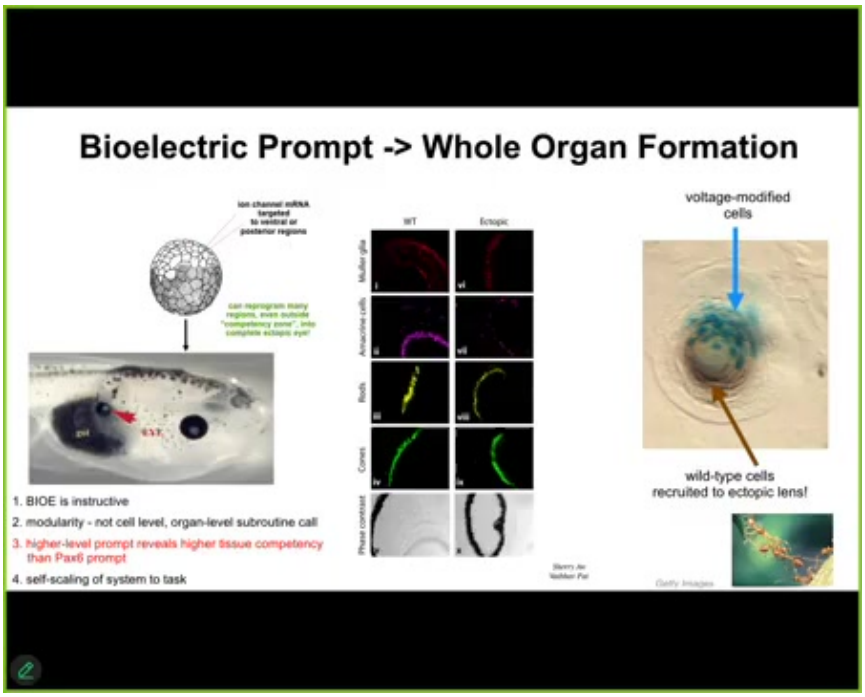

One thing you can do is induce whole organ formation. So for example, by injecting RNA for a particular ion channel to establish a voltage state that looks like the eye spot in that electric face that I showed you a moment ago, you can tell a bunch of gut cells to form an eye. These eyes have all the right internal components, lens, retina, optic nerve, and so on. You don't even have to get all the cells. If you just get a few cells, they will recruit the neighboring material to help them build the structure. It's a prompt. We don't tell the material how to build an eye. We have no idea how to micromanage the tens of thousands of cells that need to be involved and genes that need to be involved. The bioelectric layer offers a large-scale, high-level set of subroutine prompts that can let you communicate with the material and ask it to do very complex things that we don't know how to do.

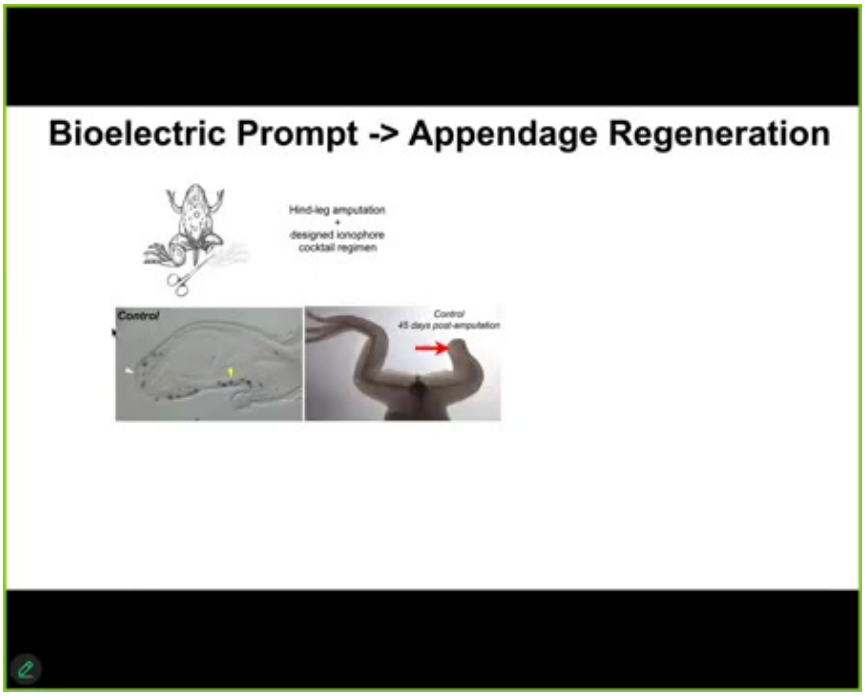

So organ formation here, for example, triggering appendage regeneration. Frogs, adult frogs, unlike salamanders and axolotls, do not regenerate their limbs. We can, by interacting with the wound for just 24 hours, kickstart the blastema with all of its appropriate markers. By 45 days, you've got some toes and a toenail. The leg eventually is quite serviceable. It's touch sensitive, it's motile.

The amazing thing is that you interact with it for 24 hours. You then get 18 months of growth. We don't touch it during that time at all. This is not scaffolds. This is not stem cell manipulation. This is not micromanagement. This is convincing the tissue on the first day that they should travel down the leg-building route in morphospace, not the scarring route. And so this is the power of these kind of high-level control points where the tissue actually stores a memory of what all the future activity is going to be aimed towards.

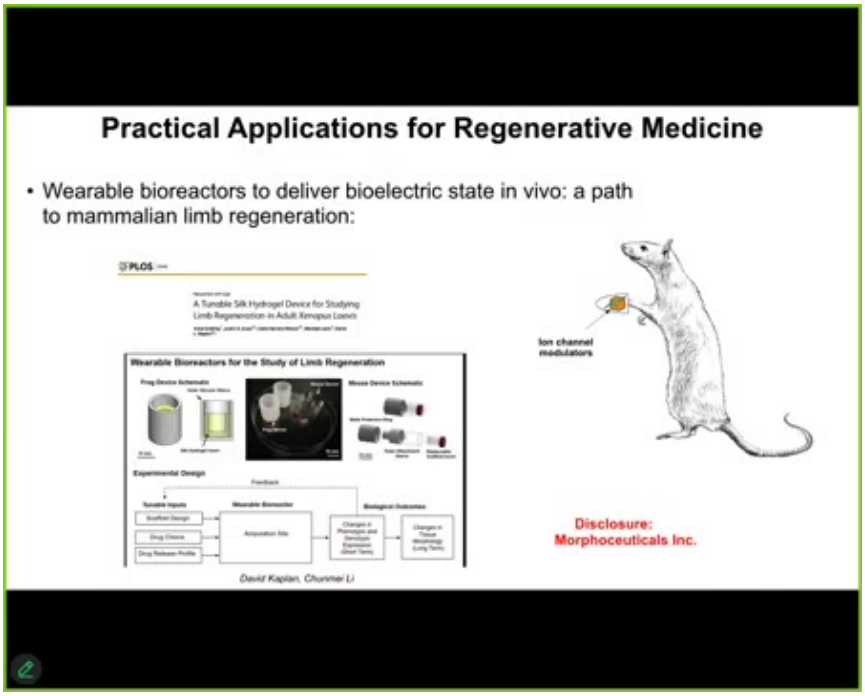

And so, of course, now we're trying this in mammals. And so with David Kaplan's group, we're making these bioreactors to apply the interventions. And it'll be a very interesting question to find out what the age of the new tissue is and things like that. So that'll be interesting.

Slide 10 of 13 · Watch at 11:10

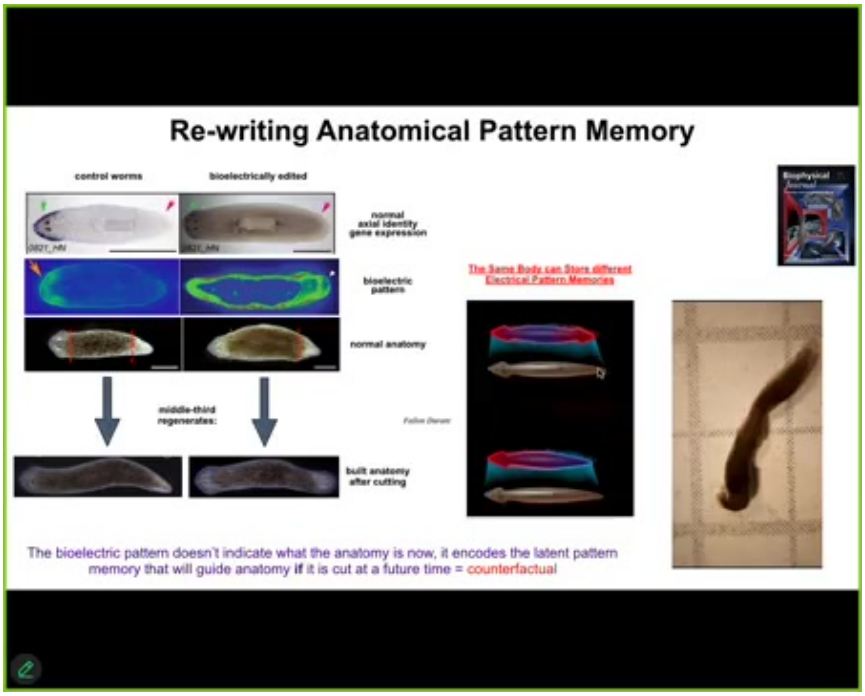

We can also rewrite anatomy on a very drastic scale. So these are, again, these flatworms. You can cut them into pieces. Each piece makes a new worm. How do they know how many heads a piece is supposed to have? You might think it's genetics. It appears not to be because we don't touch the genetics. What we do is use ion channel drugs. We rewrite the bioelectric state. So instead of one head pattern, we make a two-head pattern. And sure enough, you get these two-headed worms. And here you are. You can see them doing their thing. So the body of a planarian can store at least one of two different targets for what a correct planarian should look like. And the memory is solid. If you keep cutting these two-headed worms, they'll continue regenerating two-headed. They don't forget. So what I'm telling you is that the bioelectric layer stores set points that guide cell activity. We can now visualize them. We can reset them, at least in some cases. And what that means is that we can, without having to micromanage all of the molecular and cellular events, rewrite a new set point, and the tissue will do what it does best, which is build to a specific set point. So we develop these things in systems that are very good at that, right? So it's a practice. So these planaria are amazing regenerators. Xenopus is sort of not as good, but the tadpoles can.

Slide 11 of 13 · Watch at 12:30

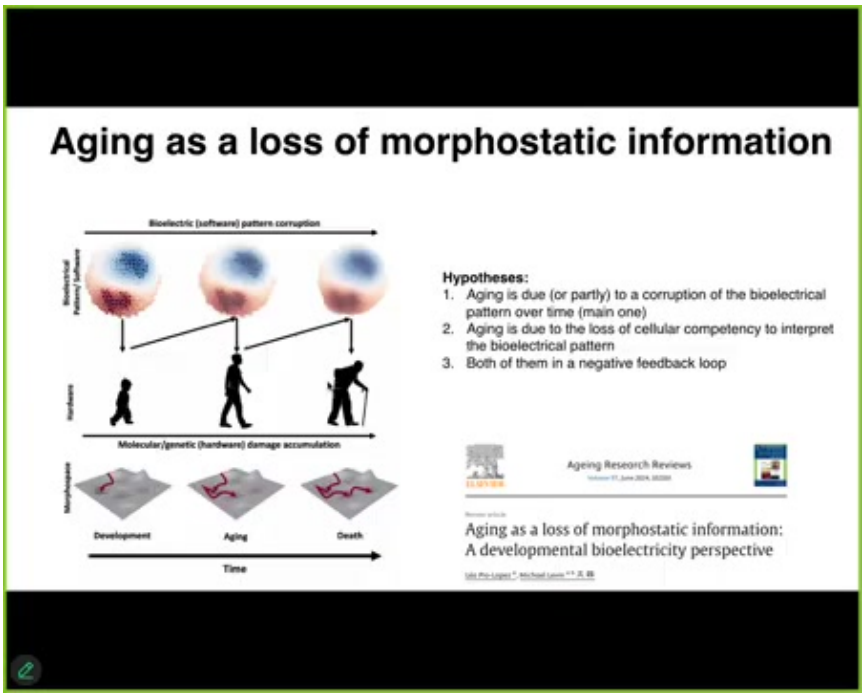

So here's where we finally get to aging. So I've shown you examples of healing and regeneration and things like that. But one of the hypotheses that we had is that perhaps what's happening with aging is that initially crisp bioelectric patterns are becoming fuzzy. They're degrading over time. And maybe what's happening is the reason the ship of Theseus isn't doing as well at this end is because the pattern, the information pattern that reminds all of the cells what it is that we're trying to make, is actually getting fuzzier and less precise over time. So this requires us to do a few things. First, to characterize the changes in bioelectric state over the lifespan, first of all, and then to see what can we do about this. Can you take these fuzzy images and get them nice and crisp? So I don't have time to go into details, but we've looked at senescence in vitro, and it turns out that, first of all, there are some very consistent and interesting changes in the bioelectrics of human cells in vitro as they undergo senescence. And by the way, also interestingly, these bioelectric patterns that I speak of, so the spatial distribution of voltage gradients, are also present in vitro, even in a Petri dish, when you have individual cells and it's not doing any kind of significant morphogenesis, even there you have large-scale patterns. And Hamid Siddiqui in my group characterized this extensively.

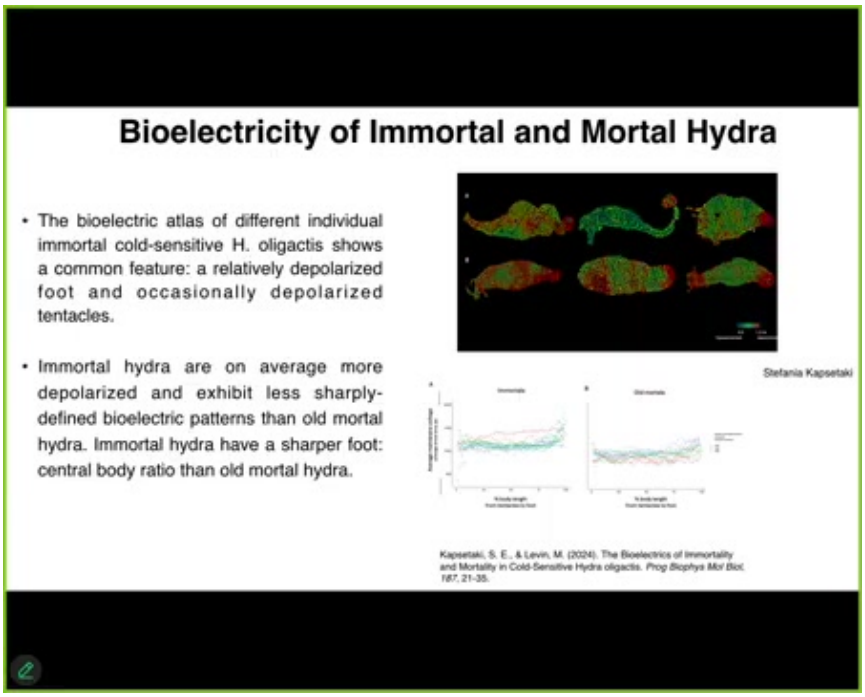

We then looked at hydra. So it turns out that if you look at mortal and immortal hydra, there are hydra that have been induced to age, they have quite different bioelectric states. And so this is all ongoing, of course. We're very interested in human tissues, in vivo, and so on. And so stay tuned for that.

But here's the best part. If you're wondering, OK, so the bioelectric pattern degrades, but it's going to be really, really difficult to go back in and correct that bioelectrical pattern in vivo. How are you going to get it to the right spot? So what we were able to show is that when you induce, so this is a model of repair of birth defects, but I think exactly the same strategy is promising for longevity work, and we're doing that now.

So here's a normal tadpole. Here's the brain. It's got the forebrain, midbrain, and hindbrain. Here are the eyes. This tadpole has been injected with a dominant mutation of the notch gene. So this is a brutal genetic intervention. Basically, there is no forebrain, midbrain, and hindbrain, or a bubble, because not such an important gene for neurogenesis and brain patterning. These animals have no behavior to speak of. They're profoundly deformed.

And then what we asked, what we did was we looked at the bioelectrical pattern here, and we realized that basically it was much flattened. It was basically the difference between the outer edges of the neural tube and the voltage of the inner edges of the neural tube that tells the embryo how big should your brain be, what should be the structure of the brain. So there's a bioelectrical gradient that says this is what a normal

brain should look like. In these systems, once we hit them with various teratogens, including notch mutations, that difference evens out.

And what is even more challenging is that you can't simply raise the voltage. That's not enough, because this is as wrong as this. What you want is this actual pattern, this complex pattern where the outer edges are depolarized, the inner is hyperpolarized. And so we had a computational model that was built by our collaborator, Alexis Pytak. We asked the model, what channel could we, what ion channel could we open or close to get to sharpen up this distinction, which is very, very fuzzy at the beginning. And it made a suggestion of one specific channel, HCN2, and we were able to show that if you either overexpress or chemically crank up HCN2 activity, you get a normal brain. Okay, and so normal gene expression, normal morphology, normal learning rates. These guys, even though they still bear the mutation, they're completely repaired.

And so what you see from here is two things. That, first of all, at least in some cases, you can fix even very difficult hardware defects such as mutations, you can fix the phenotype at the software level with this bioelectrical manipulating the bioelectrical information, that's A. And B, what we literally did here was to, with a very simple intervention that did not have to be patterned either in space or time, we didn't have to go tattooing the cells with individual channels, we're able to sharpen a fuzzy pattern that's causing a whole body level defect, able to sharpen it and get it nice and crisp so that a normal embryo results. So this is the kind of system, ion channel drugs, aka electroceuticals, which there are an incredible amount of guided by a unique computational platform that tells us which channel drugs you need at any given moment in time.

Okay, so that's the story of our bioelectrical approach. Now, upstream of this...

Slide 13 of 13 · Watch at 17:36

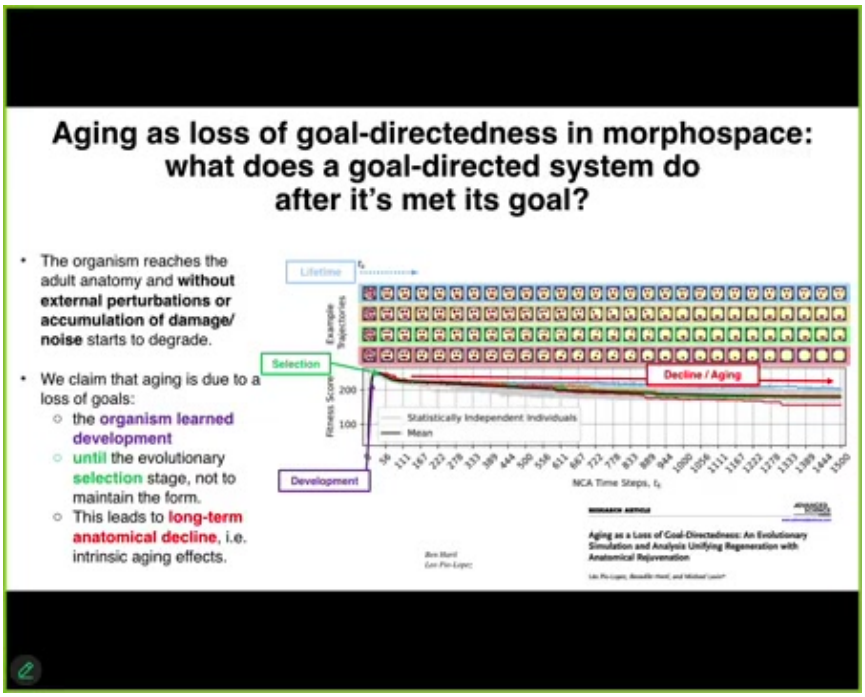

Let's talk about this for a second. Aging, as a loss of goal-directedness in morphospace. Why would bioelectrical patterns get degraded? Now, there are, of course, thermodynamic noise theories and so on, and so that could be. But there's another thing that I think is very interesting, which is the following. You could ask the question, what does a goal-directed system do after it's met its goal? And what I mean by goal-directed is a homeostatic cybernetic system such as the networks of cells that build your body, that repair it to a specific pattern, after it has met its goal during embryonic development and achieved the pattern. What will it do after that? And you can start to sort of think about ways that this could go wrong that have nothing to do with hardware defects by thinking about psychologically, if in the sort of the Judeo-Christian version of heaven, if you got to heaven and there was no damage, there was no entropy, there was no infection and nothing like that. You know, we could all keep ourselves busy for, I don't know, 10,000 years, but after a billion years, would we still be sane? Does that seem reasonable that a coherent psychological being can remain intact for very long periods of time after all of its goals have been met? And so we made a simulation that basically just studies cells that are rewarded for building a correct body, in this case of the electric face that we showed you. And what we found as an emergent feature of this model, we did not bake in any programmed aging, we did not put in any noise. As an emergent feature of this model, it starts to degrade. This is basically a psychological model of aging that has nothing to do with the physics or evolution, because neither of those forces are present here, it's a completely different cause that I think is very interesting. And then what we found in that same model is

that if we force regeneration, it gets rejuvenated again, which is maybe why the planaria are able to do this for infinite time scales, because basically every two weeks they give themselves a new challenge, they rip themselves in half, and this is enough to keep the cognitive loop engaged. And so the last thing I want to say is basically this, and we want to talk about cancer as an example of this. One of the things that happens during embryonic development and evolution is that individual cells with tiny little goals, they have tiny little cognitive light cones, they only care about very small scalar metrics at their scale, they connect into electrical networks that have grandiose goals. They're building limbs, they're building other organs. And so when cells disconnect from that network, this is like all of our stories on aging, but the whole point here is that what prevents things like cancer and aging and so on is the continued connection to a large-scale network, both electrical and chemical and probably biomechanical, that stores the target of, the large-scale target of what it is that they're doing. And so we asked, so the final part of the story is we asked what happens to that concordance? any large, any high level of organization has to align its parts towards goals that the parts don't know anything about. And so we analyzed cells, existing data on transcriptomics to find out what actually happens to the concordance of cells. And what we found, we've called atavistic dissociation, because what we found, and I don't have time to take you through the data now, but the bottom line is this: that, basically, when you're young, all of your cells, as far as their transcriptome goes, are in exact agreement on what the evolutionary age is of the body. In other words, not the physical age or the epigenetic age, the evolutionary age. So we're looking at phylostratigraphy here of the transcriptome. And what happens as you age is that some tissues, and they don't all do this to the same extent, but many tissues in the body become divergent. The cells are no longer in agreement on where you are on the tree of life. So the transcriptomes start to float off to other parts of the tree of life.

So basically, just to put everything together, this is what we have: We think that upstream of all these things may be a new cause, which is fundamentally cognitive or cybernetic at the level of the tissue, not at the level of the human mind, but at the level of the decision-making of the individuals. Downstream of this, we're characterizing a set of mechanisms that have to do with a blurring of the bioelectrical information. And downstream of that, there are many consequences. But one is this amazing dissociation in terms of the evolutionary age that your different tissues think they are in terms of their transcriptome.

And the very final thing I'll say is that we also spend a lot of time thinking about the longevity of the species. So there's this paradox of change. Like, if you don't change during evolution, if you don't change, you will eventually die out. If you do change, are you still the same? And what does it mean for our species to have longevity? What does the Ship of Theseus look like for us as a species? We know we've changed a lot on an evolutionary and developmental scale, but now that we can make biological changes and technological changes and so on, what does that look like? What does longevity look like for us? And are we going to repair this particular body and try to

continue that? Or is it going to be more like this kind of metamorphosis situation? I think it's the latter.

So I'll stop here just to thank the people who did most of the work that I showed you today, our amazing collaborators and our funders. And then three or four disclosures: Here are some companies that fund some of this work.

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.