Hessameddin Akhlaghpour outlines how RNA may implement universal computation

Could the brain's computational abilities extend beyond neural networks to molecular mechanisms? Akhlaghpour describes how natural universal computation may have evolved via RNA mechanisms.

26 November 2024 | by PAUL MIDDLEBROOKS

This transcript has been lightly edited for clarity; it may contain errors due to the transcription process.

Hessam Akhlaghpour
One of the most important insights of the 20th century, in my opinion, was the finding that with a very simple set of rules, you can achieve what's called universal computation. It's common wisdom that our models of computation achieve universality, but it's wrong [chuckles] and I'll explain why. When you take a step back, you see that there are these molecules within cells that resemble strings of symbols, and they also fold up into these tree-like structures that would be very useful for doing computational stuff.

Paul Middlebrooks
This is *Brain Inspired*, powered by *The Transmitter*.

[pause]

Paul Middlebrooks
Hey, it's Paul. My guest today is Hessam Akhlaghpour. Hessam is a postdoctoral researcher at Rockefeller University in the Maimon lab. His experimental work is in fly neuroscience mostly studying spatial memories in fruit flies. However, we are going to be talking about a different (although somewhat related) side of his postdoctoral research. This aspect of his work involves theoretical explorations of molecular computation, which are deeply inspired by Randy Gallistel and Adam King's book Memory and the Computational Brain. Randy has been on the podcast before to discuss his ideas that memory needs to be stored in something more stable than the synapses between neurons, and how that something could be genetic material like RNA. When Hessam read this book, as you'll hear him describe, he was re-inspired to think of the brain the way he used to think of it before experimental neuroscience challenged his views. It re-inspired him to think of the brain as a computational system. But it also led to what we discuss today, the idea that RNA has the capacity for universal computation, and Hessam's development of how that might happen. So we discuss that background and story, why universal computation has been discovered in organisms yet since surely evolution has stumbled upon it, and how RNA might and combinatory logic could implement universal computation in nature.

Show notes are at braininspired.co/podcast/199. If you enjoy this episode, you might also like episodes with Randy Gallistel and David Glanzman, episodes 126 and 172, respectively, which I also link to in the show notes.

Thank you to all past, present, and future Patreon supporters, one of whom actually just created a Brain Inspired specific search engine which was shared in the Discord. Thanks for that, Bryan. I hope it proves to be a useful resource for our little community here.
Ok, here we go with Hessam.

[pause]

Paul Middlebrooks
Last time, well, I guess we were off the boat, so I was at this workshop in Norway that you were at. That's where we met, and you were talking Combinatory Logic and RNA then, and that's where we're going to talk about now. It was fun on the boat with you getting to know you a little bit, and good to see you again.

Hessam Akhlaghpour
Yes, good to see you too. Yes, I'm super excited about this opportunity to talk to you. I told you that I was a long-- I was an old fan of this show. I started listening to it very early in my podcast, and to imagine that I'd be speaking on it is a very exciting thing.

Paul Middlebrooks
Well, I would be remiss to say, you actually had mentioned the *Brain Science Podcast* by Dr. Ginger Campbell to me and how that was an early influence. She was like an early-- I loved her podcast too, and that was part of the inspiration eventually when I started *Brain Inspired*. Shout out to Ginger.

Hessam Akhlaghpour
Yes. I love that podcast. I wish it was still going on, but yes, sometimes I just catch myself going back to listening to very old episodes.

Paul Middlebrooks
Oh, yes, because she does a really good job. She's a really good host. I'll just leave it at that.

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
Unfortunately she doesn't make it anymore, but I remember going on runs in Nashville, Tennessee. You know how you have that memory of where you were when you heard something or when you were reading something, and maybe we'll talk about that with that Gallistel book that we'll mention in a few minutes. I remember specific places in Nashville listening to her podcast and just enjoying it a lot.

Anyway, good to have you here. What we're going to talk about is what you've been on lately-- Not lately, for the past few years-- which is the RNA and universal computation. That's not how you came here. I know you've worked with Drosophila, you've done a lot of experimental neuroscience work up to this point. What do you do in your-- what's the right way to say this? In real life, in your day job?

Hessam Akhlaghpour
My day job is basically doing experiments on flies. I'm in, Gaby Maimon's lab here at Rockefeller. Basically, I'm doing fly neuroscience, doing behavioral experiments, using your genetics-

Paul Middlebrooks
You're a post-doc.

Hessam Akhlaghpour
-imaging. Yes, I'm a post-doc. Yes.

Paul Middlebrooks
All right. I just want to bring that up because what we're going to talk about is something that you and I also shared sort of a-- Well, I want you to tell your story of how you came to this. How you came to what you're studying now, just as a background, because I had the same-- I wonder how many-- what percentage of graduate students have this-- what would you call it?

Hessam Akhlaghpour
Disillusionment.

Paul Middlebrooks
Well, a lot have that, but very specific kind of disillusionment in that like, "Oh, is this all wrong?" That is a pretty major disillusionment, but, not necessarily is this all wrong, but a conniption about what you're doing and stuff. Tell the listeners.

Hessam Akhlaghpour
Yes, sure. I did my undergrad in computer engineering. I was really into computer science, algorithms, data structures. I felt that I was very proficient at that stuff. For grad school, I decided that I want to go into neuroscience because this is the most exciting field right now. The brain, poses a very challenging problem to scientists. It seems like I can use all of the skills that I learned at computer science to try to understand this very complex system that's mainly known for being a computational organ.

I came in naively thinking that, okay, all of this stuff that I learned about designing algorithms, data structures, figuring out what algorithmic complexity this algorithm runs at, what's the memory complexity, all this kind stuff, I thought that would be relevant to the study of the brain.

Paul Middlebrooks
Just relevant or you thought, "Oh, I'm going to find all these algorithms in the brain. I'm going to find the computational complexity and it'll map on to processes and stuff. "Was it more direct or was it just relevant, you thought?

Hessam Akhlaghpour
I don't remember exactly what I felt, but I feel like I felt that I had the right skill set, but very quickly I was humble to learn that actually not any of this is useful. It might be useful when you're like data analysis or coding up some behavioral experiment. To understand the brain, it's not-- classical computer science isn't very relevant. I guess the talking point that everyone would use is that brains are not designed, they're evolved, they're messy, they don't conform to engineering standards of design.

Paul Middlebrooks
We're going to come back to that very point when we talk about-- Sorry to interrupt, but we're going to come back to that, of course.

Hessam Akhlaghpour
Yes, definitely. Yes. Basically, the first few years of my grad school, I learned that the brain is not a computer. It's a messy, wet organ. You're going to have to understand it the way it is, and not try to impose your own idea of how computation should be on that organ. Towards the end of grad school, I felt that sense of disillusionment that you were talking about, about the whole field, like, what are we doing? It didn't feel like we're making any progress.

I'm saying this in a very, I guess, not in very generous way. Let me just say in the most extreme way that we're just collecting data. We're just collecting more facts about the brain and not really making any insight.

Paul Middlebrooks
That's the old criticism of biology when it was called it was called-- was it Ernst Mayr who called it stamp collecting?

Hessam Akhlaghpour
Yes, stamp collecting. I agree that this is not a very generous way to frame it. I'm just expressing the feeling that I had at the time. Because people are doing amazing work.

Paul Middlebrooks
Of course.

Hessam Akhlaghpour
Not everyone's stamp collecting.

Paul Middlebrooks
Good save. Good save.

Hessam Akhlaghpour
[laughs] I really beat it. It's not as a save, but-

Paul Middlebrooks
It's true.

Hessam Akhlaghpour
-just like the general direction of the field seemed aimless to me, and then that was-- Go ahead.

Paul Middlebrooks
Did your own work feel that way also? Because often people think that but then except my work. What I'm doing is right on course to solve the thing that I need to solve, right?

Paul Middlebrooks
Yes. I like the work that I was doing. The issue was I had high hopes for a certain direction. What I was trying to do-- I worked on rodent neuroscience in grad school with Ilana Witten at Princeton. I started off thinking that I'm going to solve basically some very fundamental thing about how working memory functions.

I was thinking in terms of, "Okay, I'm going to optogenetically turn off all these cells, and then and then get the rat to forget the short-term memory that was in its head." That was a very ambitious goal that I had in mind, and I didn't get to that. I got to something that was valuable, understanding how the striatum is involved in working memory. I made a small contribution to that field, but it wasn't that related to my own stuff. That disillusionment that I felt was more just about the whole field.

Paul Middlebrooks
Are you familiar with the quote-- I'm going to misattribute every single quote I try to quote here, but I think it's Mike Tyson, or at least it's usually attributed to him, that everyone has a plan until you get punched in the face. That's like how experimental neuroscience works, for a lot of experimental science, I guess.

Hessam Akhlaghpour
Yes, experimental biology, yes, exactly, yes. I hadn't heard of that, but it sounds like-- yes.

Paul Middlebrooks
All right, so you were doing good science, and Ilana's lab's very good, and you continue to do good experimental neuroscience research, but eventually, you felt that disillusionment with the field as a whole.

Hessam Akhlaghpour
Yes, until I suddenly got my hands on Randy Gallistel's book, *Memory and the Computational Brain*. The reason why it resonated with me was that it allowed me to unlearn what I had learned about computer science being irrelevant to the brain.

Paul Middlebrooks
Oh, I see.

Hessam Akhlaghpour
I'm talking about classical computer science like algorithms and data structures, because the first nine chapters of the book are basically just-- Honestly, I skipped the main text of those nine chapters, because it was like stuff that I already knew from, studying computer science. I just skimmed through them and read the summaries at the end, but then the rest of the book was actually-- I was a grad student and I was surrounded by very smart people, people who were very knowledgeable that had a certain perspective about neuroscience.

Then here comes along Randy Gallistel, this professor in psychology with a very good reputation of being, a serious scientist, as saying that, actually, you know what? It's okay to ignore this common wisdom that everyone is saying and treat the brain as a computer and use principles of computer science and theory of computation in your study of the brain. That got me super excited because that's the reason I came into neuroscience.

Paul Middlebrooks
Wait, so I don't know if we said the name of the book. It's *Memory and the Computational Brain*. I'll link to it in the show notes, of course, but, so I've had Randy on, and I've had, in a similar vein, Dan Glanzman on, and now you will be the third person to be talking about this. It could be RNA or something sub-cellular, something molecular.

In that book he talks about path navigation, how ants keep track of where they are and how some of the stories that we, neuroscientists, we-- I don't think about ants too much, but the field of neuroscience has a story about how it works. He goes into arguments why it wouldn't work this way. Same with bees. He goes through lots of examples carefully saying, "Well, this would not work." Then of course some of the learning studies that he and others have been involved in.

That opens you up into feeling that it was okay to treat the brain like a computer again. Did that make you feel like the brain is a computer again?

Hessam Akhlaghpour
I particularly remember in the last chapters of the book, when they started speculating on where the solution might be, and that's where they brought up the authors of the book, Gallistel and King, they brought up the idea that it could be stored in molecules the same way that we have genetic information stored in DNA. Maybe that's how cognitive memories are stored or it could be something else.

It could be the same way that you could have, specific changes to molecules like, I don't know, phosphorylation of some molecule and the way that those phosphorylation rates are distributed across cells or something like that. There could be various ways that you could imagine memories being encoded. It allowed me to let go of this synaptic hypothesis. That's the dogma of the field.

Paul Middlebrooks
Right. I just want to spell that out really quickly. Throughout the book, Gallistel and King build the argument, and Gallistel in his other works, that there are these problems, for behaviors and memory and learning that we don't have solutions for in the spiking patterns of neurons, which has been the hope and the assumption of neuroscience. Everything is spiking and everything is how the neurons communicate with action potentials and the patterns.

However, he goes to great pains to show in multiple cases that there is not a good story and that there doesn't seem to be possible, a good story, on principles. Just correct me if I'm wrong as I'm sort of spouting a self-memory.

Hessam Akhlaghpour
That makes sense. The arguments were-- a lot of them were conceptual in that book. For example, there was an emphasis on the need for a read-write memory.

Paul Middlebrooks
Right. Right.

Hessam Akhlaghpour
Synapses aren't really a read-write memory. You can't go in and write a specific value into a synapse, or read a specific value that's stored in a synapse. Many other conceptual arguments that Gallistel has made in his other writings. How do you encode a number? What's the code? A lot of people brush those questions aside, and I understand their arguments, but I just don't agree with them anymore.

Paul Middlebrooks
Anymore. Okay, so you had this experience reading the book. You couldn't put it down. That book specifically, and most of Randy's work is on memory and learning, and how those could be implemented at the subcellular level with hypothesized subcellular substrates like DNA, RNA,

proteins, et cetera, phosphorylation, methylation, other various possible means of doing things. Then you took a different course on it because-- is this a good time to talk about your interest in universal computation and Turing equivalence, et cetera?

Hessam Akhlaghpour
Yes, this would be a great segue into that. Most of Randy's arguments come from an angle of understanding memory. The concept of memory isn't-- not everyone agrees on what memory is. There could be semantic debates that just pop up on the side when you're discussing what is the physical substrate behind memory. There's another angle which you could take which is just as rigorous, if not even more rigorous, which is computation.

You can ask, what is the computational scope of a system from the lens of theory of computation? When you ask that question, that also leads you down towards molecules and RNA. A lot of the paradigms that we have, the models for computation in neuroscience, fall short of what's called universal computation. Actually, maybe it's better for me to just go straight into what I'm talking about. What is universal computation?

Paul Middlebrooks
Yes, and why we care.

Hessam Akhlaghpour
Yes, exactly. Yes. Okay, so in the theory of computation there are various levels of computation power that a system can have. One system might be able to compute a certain set of functions. Another system might be able to compute more functions, just the same set of functions as another system, but even more. For example, finite state machines, they can compute things like, what's the remainder of this number when you divide it by 7? There's a single finite state machine that does that, and it will do that for any number. It doesn't matter how many digits you give it, it's always going to be correct.

Then there's some problems that finite state machines can't solve. I don't know, what's one example? I don't know, "Is this string of parentheses balanced?" That's a problem that there's no finite state machine that can solve that problem for any given string. However, there are other systems of computation which could solve that.

Basically, the point is that you can have different computation systems that are able to solve different sets of functions. Now, one of the most important insights of the 20th century, in my opinion, was the finding that with a very simple set of rules, you can achieve what's called universal computation. You can build a system that's capable of solving any solvable function, any computable function.

When I say capable of solving, it requires a description of the algorithm. It's not like, "Okay, I have a universal computer and I can solve everything." No, you need to find the algorithm that solves certain problems. When I say capable of solving, it means there is a description of a program for every computable function.

Paul Middlebrooks
It has the capacity for that description.

Hessam Akhlaghpour
Yes, that's right. Yes. A really cool thing is that they had these competing models of computation back in the 1930s, general recursion theory, and lambda calculus, and then Turing's automatic machine, which we now call Turing machines, came along. Within a few years, they realized actually all of these systems, which were intended to be models that capture what it means to compute, they all are equivalent, meaning that you can simulate any one of these systems with another of these systems.

That led to the idea that maybe we've arrived at something very profound. Maybe we've found a limit to what's computable. computable because there are functions that you can describe that are not computable, but-

Paul Middlebrooks
Like life.

Hessam Akhlaghpour
Well--

Paul Middlebrooks
We'll get to it. We'll get to it. [laughs]

Hessam Akhlaghpour
One example is Chaitin or Chaitin's constant. Actually, I'm embarrassed to say, I don't know how to pronounce his name, but let's say Chaitin's constant. It's a very well-defined number, but you can't compute it. Let's not get into the things that are not computable. The point is that you can really easily reach that level of computation power, where for every computable function, you will have a description.

That description can be the description of how a trade machine's operations work. It could be a description of how a lambda calculus function works. The point is, for every computable function, you're going to have a finite length string that determines how the system operates through time. That will lead to solving a certain function. Now, this is kind of the theory of computation lens.

Now, you can ask, "We have these models of computation and neuroscience like we have neural networks. How is a neural network, a computation system?" Well, for every function you can have a description of a network that may be able to solve that function. The description of the network would be the set of neurons, the weights between these neurons, and the activation functions that each neuron has. In a string, I could describe a network, and this network would be solving a function.

Then you can ask, "Okay, well, what are the set of functions that neural networks can solve?" It's common wisdom that our models of computation achieve universality, but it's wrong. [laughs] I'll explain why.

Paul Middlebrooks
Okay. Wait, our neural models of computation?

Hessam Akhlaghpour
Well, our models of biological computation.

Paul Middlebrooks
Okay. Yes.

Hessam Akhlaghpour
Back in the 1990s there was a series of papers that showed that you could simulate a Turing machine with neural networks. The problem was that the kinds of neural networks and dynamical systems that were shown to be able to simulate Turing machines, they are irrelevant to biology, because they lack structural stability. They're even irrelevant to engineering, you couldn't even engineer these systems.

Paul Middlebrooks
Can you go-- I think you're about to go into more detail on why that's the case.

Hessam Akhlaghpour
Yes. The crux of the matter is structural stability. When you're describing a dynamical system, the system includes a number of parameters. Then you can ask, "What happens if I change these parameters by some infinitesimal small amount? Will it still resemble the same dynamics of the original dynamical system?"

In other words, is there, in technical terms for those who are interested, is there a homeomorphic neighborhood of dynamical systems to this system that you're describing? If there isn't one, then you're describing a singular point in parameter space.

Paul Middlebrooks
Fragile, could be fragile.

Hessam Akhlaghpour
Yes, exactly. The smallest error in your parameters when you're trying to implement this system would result to something that's vastly different. This comes from-- this is not my argument. This is Chris Moore's argument, which was actually the first person to show that dynamical systems can be used to simulate Turing machines.

He argued that structural stability is a reasonable criterion for systems that either an engineer can build, or you would be able to find to occur in nature. He also conjectured that no universal finite dimensional dynamical system would be structurally stable.

Paul Middlebrooks
Okay, so just to pause here, make sure I'm getting this. How do you square this with the idea of degeneracy in circuits in the brain, for example? You can use this exact same circuit to produce different rhythms or you can use different parameters in the same circuit to produce the same rhythms in this case. That seems unstable. Its robustness--

Hessam Akhlaghpour
Well, I'm not sure if that would be structurally unstable, because the thing is, almost all of the models that people use, not just for not just in like, studies of biological neural networks, but even in AI, almost all the models that people use, they are robust to a small enough amount of error in their parameters. Otherwise, they would just be irrelevant to deep learning. In deep learning, you're searching for a network that, for example, fits a certain function, and you're just moving in parameter space. If the solution is a single point with no clues nearby, then it's hopeless. You can't find that solution.

Paul Middlebrooks
Right. Yes. Okay. I thought you were saying the opposite, so I misunderstood. I thought you were saying that biological systems are inherently stable. That is what you're saying. Biological systems--

Hessam Akhlaghpour
Yes. I would say our models of biological computation that we actually use, that we actually think might be relevant, they're all structurally stable.

Paul Middlebrooks
Fair enough. The model is stable.

Hessam Akhlaghpour
Yes. Yes. We're always talking about, "I think this model is how the brain computes." Those models, they don't have this weird feature of structural instability. It's as if we're paying lip service to universal computation. If we just say, okay, look, RNNs are universally powerful. Then we never even talk about that neural network that is universally powerful. There's a subset of RNNs that we actually study, and there's another set of RNNs, which are universal. As far as we know, those don't overlap.

Paul Middlebrooks
For it to be universal, it has to compute that single point. Is that correct? Because you have to be exact when you're computing. I'm sorry I'm so naive here, by the way.

Hessam Akhlaghpour
Well, no, no, no. That's a good question. The way I would say it is people have come up with a way to describe a single neural network for every computable function. However, each one of those networks is a single point in space.

Paul Middlebrooks
Oh, okay. I see. Just to be clear, are you talking about the universal approximator theorem?

Hessam Akhlaghpour
No, no, no, no. That's very different. I'm talking about Siegelman and Sontag. They had a neural network system that basically uses the-- I think it was conceived as the membrane potential of a single neuron as a unary stack. You can imagine there's a string of digits after the decimal point that represent the membrane potential of a single neuron. If you treat that as a unary stack, you can compute with it. If you have three stacks, you can compute with it.

There are other ways to do it, like with dynamical system models that aren't necessarily like neural networks, but they essentially treat digits after the decimal point of a number as a string of symbols. Strings of symbols are really important, actually. That's one of the arguments for RNA. They really allow you to express computational power in a computational system. You can achieve it with like if you treat a number as a string of symbols. I really don't think that's how the brain works.

Paul Middlebrooks
Okay, so let me just rephrase this then and see if I get it right this time in my own words. There is a space of possible neural networks, parameter sets, architectures, activation functions, et cetera. Of that entire space, there are discrete points of the combinations of all of those different things that lead to universal computational abilities. Every other point that's not those discrete points are not universally computable.

Hessam Akhlaghpour
Yes. Basically, the point is, if you're talking about the subset of RNNs that are relevant to biology, we don't know if that covers all computable functions.

Paul Middlebrooks
Okay. We don't know if it matters either, right? It matters to you.

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
[laughs] Well, this is a big deal, right?

Hessam Akhlaghpour
Yes. No, that's a fair point. The thing is, I find it really hard to accept that biology would not have stumbled upon universal computation. Because it's such an easy thing to accidentally stumble upon. When you're working in abstract systems, there are several examples of this. For example, linear cellular automatons. Wolfram's rule 110 accidentally stumbled upon universal computation. It wasn't intended to be a powerful computation system, but it was discovered to be.

Wang tiles, Conway's Game of Life. There's a lot of examples of people stumbling upon these very complex, unpredictable systems, and later discovering that they're universal. It feels, I don't know, hard for me to believe that biology can evolve something as complex as the eye that conforms to the principles of optics, that uses a lens and an aperture. Somehow it doesn't care about the principles of computation, and it can't achieve something that's so much easier to build than an eye.

Just from my intuition, it feels like, and this is an intuitive-based argument, and it might not be convincing to everyone, but I just feel like a universal computation system would have enormous selective advantages for organisms that are striving to survive and reproduce and solve complex problems. That's why I think it's a meaningful and important question to ask, where is life's universal computer? Where can we find it?

Paul Middlebrooks
You had mentioned to me that you think that this is-- We've been talking basically about neuroscience and the models in neuroscience, but you think this is relevant to artificial intelligence as well.

Hessam Akhlaghpour
Yes, basically, one of the things that I've noticed right now in the interaction between AI and neuroscience, which is actually--

Paul Middlebrooks
There's no interaction. There's no interaction.

Hessam Akhlaghpour
You've talked about this before on previous episodes, I know. No, so there's this one interaction that I can confidently say exists. I can't tell you the number of times that I've spoken to someone in machine learning or just non-neuroscience AI. I've explained how there's a problem that our current models of machine learning are learning functions in a space that's not Turing equivalent. I can get into that in a moment, because that's also something that would seem contrary to common wisdom.

I have a very similar critique to current approaches in machine learning. My argument is that, "Hey, we're not taking universal computation seriously." Then the response that I get, I can't tell you how many times that the response that I got was, "Okay, well, if the point is to create an intelligent system, aren't we intelligent and aren't we neural networks?" At the end of the day, if your argument is against neural networks, how are we intelligent?

In a sense, those people that are working in AI and working on these neural net models, they're relying on the confidence of neuroscientists that this is it. It's a neural network system that's doing this computation.

Paul Middlebrooks
Oh, no, they're not. No, they're not. No, they're not. They're not relying on neuroscientists. They're just building their models.

Hessam Akhlaghpour
They're building their models, but there is an assumption that there's no barrier to the computational ability of neural networks if the target is an intelligent system. because if you believe that our brains--

Paul Middlebrooks
Let me just [unintelligible 00:36:03] because that has nothing-- I believe that from the common AI engineer's perspective, that has nothing to do with neuroscience. You disagree?

Hessam Akhlaghpour
I don't know. Well, at least in the discussions that I've had with people, I find them referring to the fact that we are intelligent and we are neural networks.

Paul Middlebrooks
Oh, okay. Fine.

Hessam Akhlaghpour
There's got to be a neural network solution to intelligence.

Paul Middlebrooks
That's true. That is the common assumption among neuroscientists as well.

Hessam Akhlaghpour
I think the reason that--

Paul Middlebrooks
It's true also. We have neural networks, not we are neural networks, but we have a few other things as well.

Hessam Akhlaghpour
What do you mean that it's true?

Paul Middlebrooks
We have a brain. It's true. We have neurons.

Hessam Akhlaghpour
Yes, yes. The question is is-

Paul Middlebrooks
Is that enough for universal computation?

Hessam Akhlaghpour
-thought implemented through a neural network model versus some other model that might be at the molecular level? What I was trying to get at was this mutual interdependence of neuroscience and AI. How AI researchers are relying on the confidence of neuroscientists that computation is happening through neural networks, and the other way around I feel. Neuroscientists see that the most advanced cutting-edge models for AI look really like neural networks.

Maybe it's not exactly biologically plausible yet, but there's going to be some mapping at some point. That's how both fields are relying on each other's confidences that neural networks by themselves can solve intelligence.

Paul Middlebrooks
You think universal computation is required to solve intelligence, whatever the hell that means? Something you said earlier was we are human, we are intelligent, therefore we think we can solve intelligence. I wanted to jump in and say, "Yes, we define what intelligence is," so it's not like intelligence is out there and we have some, and we know what it is, we actually define it. That's a semantic issue.

Hessam Akhlaghpour
Then you could come up with a new definition that's not rounded on us. I don't really want to get into the semantic argument of what intelligence is.

Paul Middlebrooks
My point is universal computation doesn't care about the needs of an organism, for example. Every definition of intelligence of the million that are out there, there's something about learning and unpredictable environments, adapting to learn to do the thing that you need to do, solving the problem, right?

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
These are all like problem-solving things related to what you need to do. Universal computation doesn't care about what you need to do. It's just a capacity to do anything, right?

Hessam Akhlaghpour
Well, yes. The question here is it's not about, "Okay, can I learn to universally compute?" The question is when I'm learning, what is learning about? It's about picking a function in the space of all possible functions. Okay. It's about, you have a bunch of examples and you want to find the function that solves these examples.

Now again, the same question comes up, what's the space of learnable functions in your system? Is it the same space as all computable functions or are you just leaving out a ton of functions? For example, I don't know, if you're thinking about addition, let's say you have a bunch of input and output test cases. you could solve that benchmark for addition with a lookup table.

If your learning algorithm, if your learning method is restricted to lookup tables, you're going to find a lookup table that's going to solve that benchmark, but you're not going to solve addition. If you want to solve addition, then I hope that the scope of functions that you're searching for, that you're learning in, includes programs and you might be able to stumble upon the program for addition and that program can solve things that are not in that benchmark. It can generalize.

Now, I just want to be clear, I'm not saying that current methods are lookup tables, but that was just an example to illustrate the point that the space of functions that you're learning in really matters.

Paul Middlebrooks
Shouldn't we be way better at math if we have universal computational abilities that is guiding our cognition? Sorry, it's a very naive, dumb question, but--

Hessam Akhlaghpour
I don't know. I don't know how to answer that. I guess we are good at math. There are people who are very good at math.

Paul Middlebrooks
[unintelligible 00:41:42] There are people--

Hessam Akhlaghpour
There are examples of people who-- yes, but they're not, they're not running on magic. There has to be some kind of way.

Paul Middlebrooks
Yes, they are, Hessam. It's magic.

Hessam Akhlaghpour
[laughs]

Paul Middlebrooks
No, okay. I understand that there are savants in many different areas. Maybe that's not the best example, but shouldn't we all be, right? Or is the brain-- are neurons in our way? If we could just get to the RNA computation, we'd all be-- the brains are slowing our universal computers down,

Hessam Akhlaghpour
No, I don't view it as like, okay, there's the neural network, and then there's the RNA, and these two things are very different things. That's not how I would--

Paul Middlebrooks
The neural network just won't listen to the RNA, who's trying to tell it,

Hessam Akhlaghpour
[laughs] Yes. Related to that, there's some, if anyone wants a primer for this small field within neuroscience, I would really recommend Sam Gershon's paper. The reason I brought that up right now is because it's an attempt to synthesize the view of synaptic-based memory and molecular mechanisms for memory.

In the second half of the paper, he lays out some model that would synthesize these views. How do you connect these, the idea that neurons are talking to each other through synapses with this idea that maybe memories are stored molecularly? The first part of that paper is the primer that I'm talking about, because it's the best intro to this field that I know of. It covers a lot of the conceptual reasons and the empirical reasons why the synaptic story of memory doesn't really hold up.

It also has a good summary of something that happened back in the 1960s, where there was a short period of time where a lot of people were working on what they called macromolecular engram theories, where they thought that memories could be stored the same way that we have genetic information stored within molecules, maybe memories are stored within macromolecules. RNA was one of the leading candidates for this.

Paul Middlebrooks
Let me just define engram for those listening is basically a physical trace of memory in your brain. However, that's instantiated. Some people think the engram is a certain set of cells that are associated with the memory. Some people think the engram is stored within the synapses and then a growing number of people perhaps think the engram is laid out physically within these molecular structures, macro or micro.

Hessam Akhlaghpour
Yes. Yes. Let me just clarify something that a lot of the times when I bring up the idea that molecules could be storing memories one common response I hear is that of course, molecules are storing memories. Everything is molecules. Synapses are also molecules. It's going to be molecular. The real point is, where is that information stored?

Maybe a better description of this would be an atomic theory of memory. How are genes stored? What's the mode of genetic information encoding? It's really about how atoms are arranged within molecules, not how molecules are arranged within the cell.

Paul Middlebrooks
You're taking a very reductionist approach. That was going to be my reaction to what you've said about the common response was like, "Well, yes, of course. It's also in atoms, and it's also in quarks. It's also--" You can do that all the way down.

Hessam Akhlaghpour
We don't say that about genes anymore. Right now, we say genes are stored as sequences of nucleotides.

Paul Middlebrooks
Right.

Hessam Akhlaghpour
There could be little, I don't know, tricks that organisms use to also carry information transgenerationally. There's epigenetics, there's also different ways that you can have inheritance of information across organisms. The main way that we conceive of genes being stored is in the sequences of nucleotides.

Paul Middlebrooks
It's about what the right level of emergence and emergent properties is the level that carries the most causal information about what we're talking about. Man, that was a mouthful. Sorry.

Hessam Akhlaghpour
Yes, that makes sense. There's a lot of parallels. I talked to you about this when I met you a few months ago. There's a lot of parallels between how this issue is being treated now, the issue of memory engrams, how it's being treated now versus how it was treated-- how genetic information was treated before the discovery of DNA.

People used to think that it's messy. It's like, there's not going to be a clean story for it. It's in proteins. It's like, every cell has a different protein composition. Proteins are rich in information because they didn't use that word information, but they're very rich. The protein composition of a cell of a turtle is going to be different from a human cell. That's what leads to a human being formed versus a turtle being formed.

Paul Middlebrooks
Then the central dogma came about DNA, genes, RNA, the proteins. It turns out it is messy, just in a different way.

Hessam Akhlaghpour
Well, actually, I don't even know if it is messy.

Paul Middlebrooks
We can't even define a gene, right?

Hessam Akhlaghpour
Yes. The messiness that we see right now still might be a result of us not understanding the system correctly.

Paul Middlebrooks
The way we need to do that is through universal computation. Is that what you're going to--

Hessam Akhlaghpour
Well, I think so. There's a whole debate right now, well, over the past 20 years there's been a debate over the non-protein-coding portion of the genome.

Paul Middlebrooks
The junk.

Hessam Akhlaghpour
Is it functional, is it not? Yes, it used to be called junk. Nowadays nobody really calls it junk, but one end of the spectrum believes that it's not functional, most of it, and the other end of the spectrum thinks that most of it actually may be functional. Actually, when I stumbled upon this literature it was very exciting to read. It's one of the most heated debates that I know of that's out there in papers that you can read.

Paul Middlebrooks
Like the junk versus non-junk?

Hessam Akhlaghpour
Yes, over the past 20 years, I guess. The main proponent of the idea that the non-coding DNA is functional or one of the main proponents would be John Maddock. If you look up his publications you can find the trace of that debate, but the idea is that, hey-- Well, there's a couple of arguments here.

The people who say that most of the DNA is non-functional, they usually rely on things like conservation. If you use conservation as a criteria for what's functional or not, you come up with a upper bound of let's say 20% of our genome would be functional.

Paul Middlebrooks
What do you mean conservation? Sorry.

Hessam Akhlaghpour
It's what portion of the genome is observed across species.

Paul Middlebrooks
Oh, that kind of conservation. Got you. It stays the same, it is the same across species.

Hessam Akhlaghpour
Yes, exactly.

Paul Middlebrooks
Sorry. I'm trying to just make sure--

Hessam Akhlaghpour
Yes, that makes sense. The other end of the debate they would argue that conservation is not a good criteria. There's many other criteria that you can use for hints for functionality. One of the arguments that John Maddock actually has brought up is that you see that the non-protein coding portion of the genome, the ratio of non-protein coding to protein-coding increases as a function of organismal complexity. In single cells, it's a lower proportion and it just increases as you go into multicellularity. The criticism towards this is that, well, what is complexity exactly? How can you assign complexity to organisms? That's a fair criticism. What happens is if you sort animals, if you sort species based on this criteria of what's the proportion of non-coding to coding, it just looks intuitively like it's an increasing complexity.

Paul Middlebrooks
Do you know if you do the same thing with relative brain size, it's the same?

Hessam Akhlaghpour
It's not. I think ants are above us or something.

Paul Middlebrooks
Okay. On the logarithmic scale of brain complexity size to body mass-- All right. I'm not going to look it up at the moment, but my--

Hessam Akhlaghpour
Yes. I get the point here. It's like we're looking for some sign, some indicator that puts humans on the top of the chart. That's a weird thing to do. It's a very human-centric, the Earth has to be the center of the universe kind of approach. Nevertheless, there is this problem of, okay, what is all this non-protein-coding DNA doing? Is it just transcriptional noise? Because a lot of it's transcribed.

That's basically what happened 20 years ago is that we realized that these portions of the DNA that don't encode for proteins are being transcribed. There's a lot of specificity within the cell. You can see that a lot of them are localized in very specific ways, and we don't know what they're doing.

You can find correlations with certain traits and diseases, and then you would see that those non-coding RNAs that are associated with a certain trait are actually expressed in tissues that are relevant to that trait. If it's some neurological problem, then it's also expressed a lot in neurons. There's a lot of little hints like that say that, okay, there's a story about genetics that we don't understand, getting back to your point about, okay, look, genetics was actually messy.

The thing that wasn't messy was how proteins are encoded. There's a very clean story to that. There's a codon for every amino acid. There's a lookup table that the system uses, and it does a very simple translation of RNA molecule strings to amino acid strings, which become proteins. It's a remarkably elegant and clean system to encode proteins. Still, the story of, "Okay, how does this actually encode for an organism," is messy. Maybe that's because we just don't understand the system well enough. In humans, less than 2% of our DNA ends up in messenger RNA. Actually, half of messenger RNA, on average in humans is untranslated. It ends up being less than a percentage point of our DNA encodes for proteins, has sequences of nucleotides that encode for sequences of amino acids. Right now, the story is, okay, the rest is like there's a lot of regulation, and it's all about how is protein synthesis regulated across different cells.

It seems to me that when you take a step back, you see that there are these molecules within cells that resemble strings of symbols. Strings of symbols that come from a very small alphabet of four letters. They also fold up into these tree-like structures that would be very useful for doing computational stuff. Could it be the case that these molecules are involved in something more than just regulating the synthesis of proteins across cells? Maybe something else is going on. Maybe some deeper explanation would actually make it make sense.

Paul Middlebrooks
Hold off on that because I want to ask you how you even came to appreciate the combination of RNA and combining it with combinatorial logic. However, let's just--

Hessam Akhlaghpour
Yes, combinatory logic.

Paul Middlebrooks
What am I saying?

Hessam Akhlaghpour
It's combinatory logic.

Paul Middlebrooks
Combinatorial.

Hessam Akhlaghpour
There's two things. There's combinatorial logic and then there's combinatory logic. Combinatory logic is the one that--

Paul Middlebrooks
What have I been saying?

Hessam Akhlaghpour
No, you just pronounced it differently. You said combinatory.

Paul Middlebrooks
Sorry. Combinatory.

Hessam Akhlaghpour
It's combinatory.

Paul Middlebrooks
Combinatory. Geez. All right.

Hessam Akhlaghpour
Yes, combinatory logic.

Paul Middlebrooks
Now, I'll just have to go back and edit myself the whole-

Hessam Akhlaghpour
Sorry. I think that this [crosstalk]

Paul Middlebrooks
Combinatory?

Hessam Akhlaghpour
Yes. That's right. The emphasis is on the first syllable.

Paul Middlebrooks
Aglapore.

Hessam Akhlaghpour
Yes, exactly.

[laughter]

Paul Middlebrooks
This is sitting here, by the way.

Hessam Akhlaghpour
Oh no.

Paul Middlebrooks
Yes, of course, because it's just me pronouncing things like an idiot. I wanted to hang on to for just a minute because I want to ask you about how you came to appreciate this or how you came to this idea. The regulation story. What I've come to appreciate through works like Alicia Juarrero, Terrence Deacon, the autopoietic like Varela, Montemurro, the whole, it is like in life systems, like if you have in a system, the contextual things, the regulation things are way more important than the thing that we think is doing the thing, right?

You drop, so in like a water maze, the walls, or let's say a river. You don't have a river without the banks, right? The river affects the banks and banks affect the river, but when we talk about rivers, it's like the river is the thing, but the river is not the thing, it's the banks and the river, and they're affecting each other. There's top-down causation, bottom-up causation. Then I think, so Alicia Juarrero wrote this book, I think it's called *Context Changes Everything*, and she has argued for this strongly, that these processes are all affecting.

You have to appreciate the context of whatever process is happening within which a process is happening, just as much as what you consider the process. Then, and unfortunately, something like DNA to RNA to protein, that manufacturing system, if you want to call it that, could be like a huge bureaucracy where we have all this super unfortunate regulation that seems somehow necessary for a giant democracy to get nothing done, but maybe just the very little that we can get done. I'm sorry that's a bad analogy, but of course--

The other analogy I would make is people who work on cognitive architectures have come to appreciate that, all right, you have a working memory module, you have a long-term memory module, you have an executive module. Getting those things to work by themselves, not that hard. Getting them to work in concert, that's the really hard part. It's how do you regulate, how do you make these things act together, and it turns out that the regulation part of it is a huge part of it. I just wanted to linger on that for a second to say, yes, there does need to be a lot of regulation.

Maybe it's not all regulation, right?

Hessam Akhlaghpour
Now that you describe regulation like that, I think I would agree with all what you said. I was imagining that as computation, right? There's decisions that need to be made of how do I direct-

Paul Middlebrooks
A protein to the membrane.

Hessam Akhlaghpour
-the cell's function. Yes, exactly. A protein to the membrane, the cell function, where does the cell actually go in space in the body plan? What kind of proteins do I need to,-- The issue is there's this concept of gene regulatory networks, and that's I think, what people mean when they say gene regulation. If what you mean is something more broad, then I think I would agree with you because gene regulatory networks, at least the way that they're conceived of right now, are as powerful as finite dimensional dynamical systems.

They're just the same issue that comes up with neural networks, comes up with gene regulatory networks, where you have this gene is inhibiting this gene, and this other gene is promoting this other gene. You have a big network of gene interactions that define, okay, which genes are expressed and which are not. It's very similar to a neural network kind of model of computation. If you think that's how it's being regulated, I would take an issue with that because of the computational capacity of the system that you're describing.

It's very likely that if biology can or has achieved universal computation at the molecular scale, it will use it for development, for implementing a body plan.

Paul Middlebrooks
What about cognition?

Hessam Akhlaghpour
The cognition too, yes. the thing is there's so many domains of life in which computation is not just useful but essential.

Paul Middlebrooks
What is your blog called? Life is Computation?

Hessam Akhlaghpour
Yes, that's the name of the blog. Life is Computation.

Paul Middlebrooks
Everything is computation in neuroscience, and I'm reacting generally over the past few years to that, over the course of a lot of conversations I've had on this podcast. Like let's say RNA has this universal computational capacity, and, man, I don't want to just blabber on here, I want to get to the story, but if life is not computable, and I don't believe it is, because the universe is open, and so there are only solvable problems in closed domains. I'm not using rigorous mathematical terms here.

Hessam Akhlaghpour
I don't know if that's what people mean when they say non-computable in classical theory of computation.

Paul Middlebrooks
Okay, that's your thing.

Hessam Akhlaghpour
You just said it's an open system, it's not a closed system, and in fact, that's how you get universal computation. A system that is able to recruit more dimensions to store its state, that's the key ingredient for arriving at universal computation. If you think of a program-- Yes, go ahead.

Paul Middlebrooks
No, I didn't mean to interrupt. I'll come back to it, go ahead.

Hessam Akhlaghpour
Okay, so just to illustrate by that, if you think of a program that's running on a computer, it occupies a certain amount of memory and it has the instructions that would enable it to expand in memory, to recruit more memory if it needs to. It might run out of memory on your computer, which you can conceive of as a closed physical device. That program, you can't describe it as a closed system because its progression in time requires it to interact with the environment and recruit more space to actually go on with the computation.

In all of these, the reason I say that that's the key ingredient is that in a lot of these abstract systems that stumble upon universal computation accidentally, they happen to have this ingredient. If you think of *Conway's Game of Life*, you can think of some pattern of on and off cells and it's always finite in size, but it can expand in the surrounding space. If you're implementing it on a piece of paper, you might run out of space on your piece of paper, but you know that you fail to stick to the rules of the system at some point, and you need to add more paper to it. You're describing a system that's essentially open.

Paul Middlebrooks
We have not that much time left, and I want to, because we haven't talked about combinatory logic yet really, and its connection to RNA. You have this problem. All right, let me just sum up here. You go through this disillusionment in graduate school, you stumble across Gallistel's book, you start thinking about computation and you return to the idea, "Maybe I can think of the brain as a computing system or a computer." How did you come across the idea that RNA has universal computational capacity? Then how did you connect that with combinatory logic?

Hessam Akhlaghpour
Yes, so RNA was in the spotlight already, just from all these other people that are in the field that have put it forth as a candidate for storing memories or for computation. It was in the spotlight for me, like since grad school.

Paul Middlebrooks
That's because it's stable enough that it could last like not a protein that will get degraded over a day or whatever for other reasons.

Hessam Akhlaghpour
Yes, we've learned that it's stable recently. There's a paper that came out earlier this year that shows that you can have RNA strands in the nucleus that last for years, for the lifetime of the animal. I can send you that paper if you want to link that in the show notes or whatever. A lot of people think of RNA as being short-lived and transient. I think the reason why RNA, there are many reasons that it has come up, but I think just there's something very appealing about it being a string of symbols.

It was basically in the spot. In the 1960s, it was the main candidate for these macromolecular theories. They would claim that we purified RNA from a planarian and injected it to another one, and it worked. This is related to your interview that you had with David Glanzman. He's also landed on RNA methylation as a very promising candidate for memory.

Paul Middlebrooks
An epigenetic mechanism. Hang on, let me ask you. Was it also in the '60s? When did the idea that RNA may have preceded DNA as like the original life molecule? People used to think DNA is the original molecule, and then RNA came out of that to produce proteins.

Hessam Akhlaghpour
I don't know if there was a time where people thought DNA was, the first molecule. I know that the idea that there was an RNA world, that there was a world where before DNA and proteins, there was RNA. That comes from the realization that RNA both has a encoding capacity, in the sense that you could have a reader, like a ribosome, actually read its content and translate it to something. It has an enzymatic capacity, so it can serve as an enzyme to chemical reactions. These two capacities, people would argue, are the main molecular components in life.

RNA has both of them, although not as efficiently. Proteins are way more catalytic, and DNA is way more better at storing information. Then these two things evolve later. I think that's the main argument for the RNA world hypothesis. I don't know actually how serious people take it today. You know what else? DNA involved for better storage, proteins involved for better enzymatic activity, and brains involved for better cognition.

Hessam Akhlaghpour
[laughs] Yes, that evolved later too. As long as you're willing to concede that RNA had the computational capacity to begin with and then--

Paul Middlebrooks
Sure. Yes. Capacity is great. It's just like whether it's implemented and stuff. I'm all about capacity.

Hessam Akhlaghpour
Yes. Sure. Where was I?

Paul Middlebrooks
I interrupted you and asked about the RNA progenitor story with respect to DNA. The RNA was the early storage thing, but you were about to-- You were talking about in the 1960s, that was the molecule that people thought could be used as a symbolic string of symbols, right?

Hessam Akhlaghpour
Yes. A lot of people, I guess, based on just the discovery that genes are actually encoded in molecules, maybe that was an inspiration to the theory that cognitive memories might actually also be stored in molecules. Some people thought it was proteins. In fact, the earliest proposals were in 1950, before the discovery of the double helical structure of DNA. Those were based on proteins. Those were in a time where-- Yes.

Paul Middlebrooks
A time when people thought proteins were more stable? Is that what you're going to say?

Hessam Akhlaghpour
A time that people thought proteins were genes, that genes were proteins, actually.

Paul Middlebrooks
Oh, okay.

Hessam Akhlaghpour
The majority, at least according to this book, *The Age of Creation*, until the discovery of the double helical structure of DNA, most biologists thought that genes were proteins.

Paul Middlebrooks
All right.

Hessam Akhlaghpour
Yes. Anyway, the point is that in that era, in the 1960s, there are dozens of different papers that are working on this idea or doing different kinds of experiments on this idea of memories being encoded in molecules. They used different approaches. Not of all of it was like these crazy feeding planarians to other planarians that James McConnell is known for. People, for example, showed that you have changes in RNA composition with learning. They would do experiments where, I don't know, you put a rat on a tightrope, and then you would show that, okay, the ratio of RNA nucleotides changes in this brain region.

That field died out in the beginning of the 1970s. I guess that one of the main critiques to that theory to that subfield was that how do you know that the changes that you're seeing are actually encoding memories? You could extract some kind of chemical from a learned animal and inject it to some other animal, but that could just be like a hormone for fear or something. How do you know that that's actually encoding the content of the memories? They really didn't have the tools at the time to really study that question deeply.

Paul Middlebrooks
I remember now where we were because I was asking you how you came to the linkage between combinatory logic and the RNA story, and you started to talk about the '60s and how this is an old story that went out of favor. Then how did-- Go ahead.

Hessam Akhlaghpour
Our RNA was already in the spotlight for me just because there's many people that have brought it up as a good candidate. Then since reading Randy Galloway's book, I had been thinking about, okay, how would a really actually universal computation system look like at the molecular level because you could imagine like, okay, it's cool. RNA are strings of symbols, are you going to treat it like a Turing machine tape? With a Turing machine, that goes in and edits this sample and moves next.

That can't possibly exist in cells. You would see it in the ribosome, which does something very simple, which is just translates every triplet of nucleotides to an amino acid. That's huge. It's visible in electron microscopy. Trying to say that something like a Turing machine might exist isn't very plausible. What could a computation system look like? This has been on my mind ever since. I think the connection between combinatory logic and RNA happened when I realized two things basically.

One was I just learned that RNA has a secondary structure, meaning that the same way that DNA strands can come in and fold and form these double helices by you bring in two strands of DNA and they connect to each other and form a double-stranded helix. The same thing can happen within the same RNA molecule. An RNA molecule is a strand, is a string of nucleotides. You can have one segment of that strand base pair with another segment of that strand. That can happen--

Paul Middlebrooks
It has to fold on itself and the pairs have to be the right pairs that would match and it can do that-- I'm sorry, I'm trying to just be crystal clear.

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
It can do that.

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
Let's say it's like 100 pairs long and if the first and last four pairs are the ones that match, it would fold in on itself and then you'd have a big loop of all these pairs that weren't matching, and then those four that just connected at the bottom like a tiny lollipop.

Hessam Akhlaghpour
Yes, but usually, there's a lot of matches. It's not just the--

Paul Middlebrooks
I wanted to go simple. Yes.

Hessam Akhlaghpour
Yes, exactly. People do this. You can study a certain RNA strand and study its secondary structure. Usually, the maps look really pretty. They're very intricate, there's many layers of essentially parentheses. When people want to represent the secondary structure, they use parentheses because if you think about it, one part of the strand coming to another part of the strand is matching a nucleotide to another nucleotide. Usually, it turns out it has a parenthesis structure. Sometimes you get these things that they call pseudoknots which deviate from a balanced parenthesis structure, but usually, it's a balanced parenthesis structure.

Paul Middlebrooks
Sorry, this is an aside, but in a given, I don't know, 1,000 base nucleotide RNA sequence, there are going to be lots of places that could attach. What is the possible number of secondary structures that a piece that long could form?

Hessam Akhlaghpour
That's a great question and the answer is a lot. In fact, we know that a lot of RNA strands have many different confirmations that they can take and it's not always static. It's not always a single structure.

Paul Middlebrooks
Are you saying a single thing can fold up one way and then in a given solution relax and then fold up a different way?

Hessam Akhlaghpour
Yes, absolutely. In fact, that's how ribosomes work. There could be an element within an RNA strand that folds differently depending on temperature. That can be used as a sensor for temperature. That's a very important feature of RNA strands. Then just out of curiosity, I was like-- I don't know actually. I can't remember exactly how I stumbled upon combinatory logic.

Paul Middlebrooks
Did you see Lambda calculus first? You probably knew about Lambda calculus.

Hessam Akhlaghpour
Yes, I think I saw Lambda calculus first. I don't even remember the first time I was confronted with it, but I was trying to learn how it works and then at some point, it just clicked in. I think I remember the moment that it clicked in.

Paul Middlebrooks
What was that? You got to tell that story. What was that?

Hessam Akhlaghpour
There's nothing special about it. I was sitting at my desk and I was on a Wikipedia page, I think a Wikipedia page for combinatory logic. [laughter] Then I was like, wait a minute. It's very similar to how we represent RNA strands. There's these parentheses that you use for representing a secondary structure. You have a limited alphabet of combinators usually it's S and K or B, C, K, and W. That's it with that limited alphabet, you can express any computable function. The rules for running that function are very simple. They're very local.

Paul Middlebrooks
Wait, how long was this moment?

Hessam Akhlaghpour
Probably a minute. I don't remember. [laughs]

Paul Middlebrooks
Where you had all these thoughts?

Hessam Akhlaghpour
Yes. I guess just noticing the parallels between the two, and immediately afterwards, I was sure that somebody had written about this.

Paul Middlebrooks
Because it's so obvious once you [crosstalk]

Hessam Akhlaghpour
Yes. I was sure there had to be some paper or something. I kept on searching. I couldn't find anything which made me more and more excited.

Paul Middlebrooks
I'm sorry. I'm interested in this. You had this moment at the computer and you realized it, and then how soon? That must've been like, "Oh my God." Then immediately you're like, "Oh, this must have been done." It's a downer on the moment, right?

Hessam Akhlaghpour
I don't know. I guess I was preparing myself to be very disappointed that somebody else had brought this up.

Paul Middlebrooks
That's right. You'd been in academia for a while and in experimental neuroscience.

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
It all makes sense. Sorry.

Hessam Akhlaghpour
I don't know. Then we had this journal club with my advisor here Gaby Maimon. He's also very interested in how RNA might be involved in computation. He shows up a lot of the same ideas with this growing group of people that think molecules might be involved in computation more so than is appreciated. We had this journal club with two other neuroscientists Abbas Rizvi and Jeremy Dittman. We were reading papers and just thinking about how molecules could be involved in memory and computation. That's where I had first tested the waters with these ideas.

It was very critical and just polishing the theory to get the feedback in this group. It was really nice to have this group of intimate four people sitting around and then critiquing that, "This part doesn't make sense," and a lot other things. A lot of the perspectives that I just recently described it for example, the ribosome is a very impressive molecule and it's doing something pretty simple and it's large and you can see it. That's something I specifically learned from Jeremy. Just that appreciation of the ribosome. That's where I had the opportunity to flesh out the details of this idea.

Paul Middlebrooks
I asked you about Lambda calculus because I know that combinatory logic and Lambda calculus share a lot of similarities, right? I've heard you mention that combinatory logic precedes Turing machines and Lambda calculus. What the hell? What is that story?

Hessam Akhlaghpour
Yes. I love that you brought that up because it's not appreciated that the first mathematical system, the first abstract system that humans came up with that had a universal competition capacity was combinatory logic. It was discovered by this mathematician named Moses Schönfinkel. That's the only thing that I know he did. I don't know if he had any other contributions, and then he was just forgotten. Then Haskell Curry rediscovered it again and then realized, oh, somebody else had worked on it and so he gave credit to Moses Schönfinkel.

I guess, this emphasizes the point. It was discovered twice independently before we had any other system of computation.

Paul Middlebrooks
By only two humans and if we have this universal computation, shouldn't it have happened much more? I am going to keep coming back to this stupid point.

Hessam Akhlaghpour
I don't know. Yes, maybe prehistoric times. [laughs]

Paul Middlebrooks
Four hearings, two prehistoric, and two post-historic.

Hessam Akhlaghpour
Yes.

Paul Middlebrooks
Then what's the connection? Why is it special?

Hessam Akhlaghpour
The point I'm trying to make about it being the first one is that it's simple. It's very simple. It's also very beautiful. It's a functional programming language. There's nothing that's not a function. Everything is a function. They call it combinators. Everything's a combinator. The inputs of these combinators are other combinators. There's no primitive data types.

Paul Middlebrooks
It's functions of functions of functions that take in functions and spit out functions.

Hessam Akhlaghpour
Exactly. Yes. Every function takes a single function and spits out another function. The way that you can get a function that takes two inputs, the way that you can build a function that takes many inputs is to say, okay, like let's say I want a function that does addition. Addition takes two inputs. The way I can do that is I can say, I'll define a function that takes a number and then spits out the function that adds that number to any new input that it gets. That's a technique called currying after Haskell Curry was the second person who invented a combinatory logic.

Paul Middlebrooks
There's no schönfinkelling?

Hessam Akhlaghpour
No, I guess this technique was actually, exclusively Curry's idea. I'm not sure about that actually. I might be wrong about that, but anyway, so it's very simple, and then it just also very nicely maps on to RNA biology. If you want to implement something like this at the molecular level, the main challenge is parenthesis matching.

Paul Middlebrooks
Regulation in other words.

Hessam Akhlaghpour
[laughs] Yes, I guess, we're even broadening the definition of regulation even more now.

Paul Middlebrooks
It's like housekeeping, right? It's just counting parentheses. It's not the sexy computational thing. It's like, I got to keep track of where I am in this nested series of functions.

Hessam Akhlaghpour
Yes. You could do that explicitly. If I'm actually evaluating a combinatory logic term on paper, that's probably what I would do. I would keep track of the depth of the parenthesis and just keep on going and then use that technique to determine what's a single term. Then you could do stuff with that term. In RNA, because of that intrinsic, secondary structure that RNA have, you don't need an explicit machine that goes in and does this parenthesis matching because matched parentheses are already in physical space. What that allows you to do is, you can implement every single one of these handful of application rules in combinatory logic with local operations.

Paul Middlebrooks
With local strands of RNA.

Hessam Akhlaghpour
With local operations on part of an RNA.

Paul Middlebrooks
Some of which is paired with itself in a certain section, and some of which is in this open loop.

Hessam Akhlaghpour
Yes, exactly. I guess most people will be listening to this, so there's no illustration, but I'll try to describe it in the most illustrative way.

Paul Middlebrooks
For people who are watching, I'll just put up a still of you giving a talk with the hairpin loop structures, right?

Hessam Akhlaghpour
Sure, yes. Maybe the rules or--

Paul Middlebrooks
I'll also [unintelligible 01:28:11].

Hessam Akhlaghpour
[laughs] I don't have to be in the picture, just showing how one of the rules can be implemented. One of the combinatory logic rules can be implemented. These application rules and the combinatory logic, they're just very simple operations, like you swap two elements, or delete an element, or add parentheses around two elements that come afterwards, or something like that. To implement that molecularly, you only need some kind of enzyme that detects the motif that encodes for that combinator first.

Let's say, we have three primitive combinators, and that's a key point. You only need a handful of primitives. Let's say, you have three primitive combinators, so you have three different codes for these different combinators. Some enzyme should detect that, "Hey, we have a motif here that encodes for one of these combinators, and that enzyme also carries with the rule for the application rule of that combinator."

Paul Middlebrooks
I want to stop you and say, we haven't figured out how this could be implemented. These are [crosstalk]

Hessam Akhlaghpour
Yes. This is completely hypothetical. I do want to stress that the point of this model is that it doesn't require extraordinarily complex molecules, and very different from what we already know to exist in cells. Our RNA strands are frequently modified after they're transcribed. The most commonly discussed modification is splicing. There's something called the splices zone that goes in and selects segments within this RNA strands, and excises them and then attaches the two loose ends. That's a way to delete certain parts within an RNA strand.

These kinds of operations exist and we know that they're within the reach of evolution of molecular biology. The point here is that I can imagine a system that's very simple that doesn't require huge molecules that do complex operations. A system like this could have gone undetected over the many decades of molecular biology but going back to how it would work, it's basically every enzyme would execute cleavage and ligation operations, so would cleave a certain part of the strand and ligate the different part of the strand.

Paul Middlebrooks
You cut it and put it back together.

Hessam Akhlaghpour
Exactly. Cut it and put it back together and the locations of the cuts and the connections are fixed relative to the motif because by virtue of the secondary structure, bringing the parentheses together, now you can say, okay, I only need to cut at this position and connect these two parts without caring about what's inside the parenthesis. I don't care about how many layers of nested parenthesis are inside this enzyme just goes in and mindlessly does this single operation. Through this hypothetical system to emphasize, it's just a theory. It's obviously going to be wrong in its details.

Paul Middlebrooks
It's wild. It's so cool but it's wild.

Hessam Akhlaghpour
Yes, but I want to acknowledge that it's obviously wrong, it's going to be wrong in its details. I had to come up with details that allow the system to work that I know that I just came up with. The point is that this is a proof of principle that you can imagine something like this happening at the molecular level implementing a computation system. At the same time, you have all this RNA and all this DNA, that we don't know what it's doing. We haven't attributed a function to most of the genome, at least in humans. It's just very-- I don't know. It's a very compelling--

Paul Middlebrooks
Intriguing. Compelling.

Hessam Akhlaghpour
Yes, intrigue. It's a very compelling problem.

Paul Middlebrooks
Oh, it's a compelling problem but it's a compelling hypothetical solution.

Hessam Akhlaghpour
Yes. I would call it, it's compelling me at least towards a research direction, towards at the end of the day I want to study things empirically. I don't expect anyone, myself included to believe this theory until we find evidence for it. We have to do the science and so it's a research direction that I'm arguing for. Not like a specific model and the direction being, let's figure out if RNAs are being edited in ways that can implement computation.

Paul Middlebrooks
First of all, you were disillusioned and now you're hopeful. Would you call yourself hopeful? How would you describe your outlook?

Hessam Akhlaghpour
Yes, I would describe it as passionate. I would say I'm very passionate about this problem.

Paul Middlebrooks
Sure are.

Hessam Akhlaghpour
I feel like it's extremely exciting. Sometimes I forget and then I remind myself of all of the hints that are obviously pointing at RNA. It's heartwarming, I guess, for lack of a better word to know that there are also many other reputable scientists that take these ideas seriously and that that circle is growing and I hope it grows. I hope it doesn't end up being something like that subfield in the 1960s and it really depends on us. It really depends on us trying to make this argument that this is a worthwhile research direction.

I don't want everyone to be working on it. I don't want you to pour 50% of all research budgets towards it, but at least in the spirit of diverse approaches, I feel like we should be able to maintain a consistent research direction along the lines of molecular computation and memory.

Paul Middlebrooks
You just need a little drip out of the firehose of funding.

Hessam Akhlaghpour
Basically, yes.

Paul Middlebrooks
Those reputable scientists that you mentioned often get labeled, like often don't feel fully respected in the field, and often get labeled, I don't want to say crazy, but out of the mainstream, out of the dogma. Do you ever think, "Am I crazy?"

Hessam Akhlaghpour
I don't know. Actually, I might regret saying this, but sometimes I think, do people look at me the same way that I look at Roger Penrose about microtubules? [laughs]

Paul Middlebrooks
That's more Stuart Hameroff but Penrose was on board with that. Or the way I look at that. I look at the microtubule things and I'm like, "Jesus Christ."

Hessam Akhlaghpour
Yes, exactly. I don't know.

Paul Middlebrooks
No, but do you feel crazy? Not about how you think other people view you. Do you ever think like, "Oh, am I crazy?"

Hessam Akhlaghpour
I don't think so. I think the arguments and the evidence that I am resting on are very rigorous. Again, I want to be clear, the threshold of evidence you need to believe a theory is much higher than the threshold of evidence you need to work on pursuing a theory. I think it meets the latter threshold.

Paul Middlebrooks
I've had lots of conversations with people. Right now, Alex Gomez-Marin comes to mind because he's studying things that are outside the norms of what the scientific community, especially in neuroscience would consider okay to study or get funding for or something. You mentioned these people in the 1960s, this is not a new idea as György Buzsáki says, there's no new ideas under the sun. These are all recycled things. He says it in a different way and how brains do computations which he doesn't mention RNA, but that's besides the point.

You mentioned these people from the '60s the idea of RNA came up and then it died down. What I want to ask you is have you reflected on what this tells you about how science progresses because most people like stay, they get into, let's say, experimental neuroscience and drosophila, and then that's their career is just studying this space of problems fairly narrow. Now you've done that and you've learned about an alternative framework for universal computation which is one of your interests and then you realized this was not new, this ebbed and flowed already.

How does this make you feel about the history and progress of science in general?

Hessam Akhlaghpour
It makes me realize that we're in a fragile place. It makes me realize that it may very well be that people look back at this era as some time where

some idea just reemerged and it died out again. Now that's fine if this idea is wrong [laughs] but if 100 years down the road they realized this was all correct, but it kept on resurfacing and dying out, that's a possibility. I want to prevent it. I think there's no inevitability when it comes to these sociologically decided things and that's why I'm trying to get people to, if they don't want to work on it, at least agree that it's worth putting some resources into.

I just want to say that also the era in the 1960s, the ideas that were there were mostly around memory, were mostly around molecules encoding memory. I think this is not the same thing, really. It's about computation and it's about actually bringing in the insights from the theory of computation to these strings of symbols. It's obviously related to memory. Memory comes up in computer science all the time. It's a different perspective. Also, just all of the conceptual arguments are much more mature now.

Read Gallistel's work, the arguments are much more solid and convincing than anything that anyone wrote in the 1960s, and also, we just have better tools. It's a very different era, and it's a different idea but it's very related.

Paul Middlebrooks
All right. Sorry, I have to ask this. We've talked a lot about molecules and biology, and you have a computer science, computer engineering background, and then you got into experimental neuroscience. A lot of people who start off in neuroscience have this computational bent but then a lot of people who are in cell biology, for example, don't have that. That's where the stamp collecting began and speciation and things like that. Sorry. I'm trying not to bias your-- how has this altered your appreciation or lack of appreciation for the micro molecular world relative to your computational mindset?

I'm just leading you into the answer. Sorry, but I wanted your reflections on it.

Hessam Akhlaghpour
I maybe wish that in retrospect that I had studied molecular biology just because that seems like a very relevant field now for the problem I'm trying to work on. Just to say, it's not that molecular biologists or geneticists are completely foreign to these ideas or find these computational ideas foreign. John Mattick, for example, explicitly argues that these non-coding RNAs might be implementing a digital computation device. There are people who are definitely on-- these ideas resonate with their ideas. I guess, right now, not in neuroscience, not in molecular biology is anyone really trying to take universal computation seriously.

Paul Middlebrooks
Are you having fun?

Hessam Akhlaghpour
I don't know if that was the correct-- Oh, again, maybe I misunderstood your question.

Paul Middlebrooks
My question, I've had come to have a little more awe just how goddamn complex everything is. Like the world of the cell, that's a whole world. The brain's the most complex thing in the universe. Although, Terry Sejnowski pointed out to me that his wife said to him, "Actually, two brains is more complex than one brain." Two people talking which is true. [laughter] The story that you're working on, if it turns out to have validity in one form or another, just the capacity, the astounding results of evolution that continue on, oh, it's just amazing to me.

That wet biology part, when I got into neuroscience, it's all computation spikes. That's how they're doing it, information, blah, blah, blah, but then you look in the cell and it's like, "Man, that is messy and hot." It's doing just as awesome a job, whatever the job it has. It's just amazing that anything works in biology.

Hessam Akhlaghpour
It's amazing until you understand how it works exactly.

Paul Middlebrooks
It's simple. Is that what you're saying?

Hessam Akhlaghpour
Until you explain, it's a mystery. It's like, "How the hell is the single cell creating a human fetus with all the intricate body parts, and it looks like a miracle?" Now, if somebody writes a program that draws something on the screen, I wouldn't call that a miracle, because I know how programs work. Draw something, like some complex pattern, it's like that looks really cool and as complex as a human fetus or something. I don't know. Maybe that was a bad example.

Paul Middlebrooks
This gets back-- We're out of time. Another time perhaps.

Hessam Akhlaghpour
Sure.

Paul Middlebrooks
Because I've taken you over. I've gone over. Hessam, good to see you again. Thanks for doing this. It looks like you're having fun. You having fun?

Hessam Akhlaghpour
Yes. This was extremely fun, Paul. Thank you so much for doing this.

Paul Middlebrooks
No, not this conversation.

Hessam Akhlaghpour
Oh.

Paul Middlebrooks
I mean in these research questions.

Hessam Akhlaghpour
Oh, in life. Yes.

Paul Middlebrooks
Hey, am I fun? Am I fun?

[laughter]

Hessam Akhlaghpour
Yes. You are definitely fun, Paul.

Paul Middlebrooks
Yes. Thanks. It seems like you're having fun. That's a great place to be at.

Hessam Akhlaghpour
Yes. As long as I know that I can survive in academia, it would be the condition that I would add to that. [chuckles] There's that question too.

Paul Middlebrooks
I wish you survival. If we are on a boat again and you go overboard, I will throw you a life vest or something like that.

Hessam Akhlaghpour
Yes. Thank you. I appreciate that.

Paul Middlebrooks
Okay. Thanks. Awesome.

Hessam Akhlaghpour
Thank you so much, Paul. Thank you for your time. This was really fun. Take care.

[music]

Paul Middlebrooks
Brain Inspired is powered by *The Transmitter*, an online publication that aims to deliver useful information, insights, and tools to build bridges across neuroscience and advanced research. Visit thetransmitter.org to explore the latest neuroscience news and perspectives, written by journalists and scientists. If you value *Brain Inspired*, support it through Patreon to access full-length episodes, join our discord community, and even influence who I invite to the podcast. Go to braininspired.co to learn more.

The music you're hearing is *Little Wing* performed by Kyle Donovan. Thank you for your support. See you next time.

[music]
