Kenneth Harris and Andreas Tolias explain how artificial intelligence has informed their neuroscience research

Modern AI models have shaped how the pair thinks about our brains and minds, asks research questions and views scientific progress and productivity.

8 October 2024 | by PAUL MIDDLEBROOKS

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

Paul Middlebrooks
This is "Brain Inspired," powered by *The Transmitter*. Hey everyone, it's Paul. This is the first of two less usual episodes. I was recently in Norway at a neuroAI workshop called Validating Models: How Would Success in NeuroAI Look Like. What follows are a few recordings I made with my friend Gaute Einevoll. Gaute has been on this podcast before, but more importantly, he started his own podcast a while back called "Theoretical Neuroscience," which you should definitely check out. Gaute and I will introduce these conversations we had with a few of the invited speakers at the workshop and one of the main organizers.

I'll link to everyone's information in the show notes at braininspired.co/podcast/195. I hope you enjoy these conversations that we had on a rather large boat you'll hear about in a second. Enjoy.

[music]

Paul Middlebrooks
Two podcasters on a boat in Norway.

Gaute Einevoll: How can that go wrong?

Paul Middlebrooks
[chuckles] Hi, Gaute.

Gaute Einevoll
Hi, Paul.

Paul Middlebrooks
Why are we doing this together? What's happening?

Gaute Einevoll
You have been making "Brain Inspired" for how many years?

Paul Middlebrooks
I think five, maybe six.

Gaute Einevoll
Five, six years, exactly. I've been podcasting for five, six years, but this new "Theoretical Neuroscience" podcast, that started in October. I remember I talked to Konrad Kording, and "Brain Inspired" was an inspiration for me to make this what I call academic podcast. It's not really about- It's actually for people in the community. It's not a popular science podcast, really. Which are of course really important too, but it's a different thing. In some sense, when I talk to Konrad Kording about this, he said, "Oh, your podcast is 'Brain Inspired'-inspired?" That is actually true [chuckles].

Anyway, we are a little bit some sense like sister podcast. Would you say? We are both sort of— Yours are neuroAI and mine is a little bit more into the more other aspects of theoretical neuroscience. Maybe more studying the physics tradition, because that's where I come from. Anyway, when we both were going to this very nice workshop on neuroAI up on the coast of Norway, then of course we have met before. You visit me in Oslo once. Then we decided why don't we pool the resources and make a joint podcast.

Paul Middlebrooks
Yes. What came out of it is just a few discussions that we had with some of the invited speakers and then also with one of the organizers, Mikkel, to

frame the workshop. You'll hear from Mikkel a little bit about how this workshop came about. This will be two episodes. Then at the end of the second episode, which will be our second discussion, Mikkel summarizes and wraps up as well.

Gaute Einevoll
Anyway, the workshop, just to introduce that a little bit more. That was then on the neuroAI, it was Mikkel Lepperød in Oslo and I think Konrad Kording who is most centrally working out the program. There were also other organizers who I think secure the funding, like people in Oslo, Anders Malthe-Sørenssen, Marianne Fyhn and Tone Skramstad. I like to give credit to people.

Paul Middlebrooks
Of course. They did an awesome job. It was an awesome workshop.

Gaute Einevoll
It was an awesome workshop. This was quite unique. I never taken—it was a coastliner—along the most beautiful part, beautiful coast of Norway from Tromso down to Trondheim. It took three days or two and a half days or something. It was a good weather, and it was really, really excellent in all ways. As we mentioned, the first clip we present here is with Mikkel, main organizer. Where we actually discussed with him why he made that podcast and what he wanted to get out of it. It was done in the—

Paul Middlebrooks
You said podcast. You mean workshop?

Gaute Einevoll
Workshop, exactly. [chuckles] We call a Freudian slip. It's not something— Yes, I meant workshop. This was done actually at the last day of the workshop in a luggage room of a quite fancy hotel in Trondheim, the Britannia Hotel which is— I used to study in Trondheim. This was the places you don't go when you are on a student budget, to be sure.

Paul Middlebrooks
It was fancy.

Gaute Einevoll
Yes, it was very fancy. We're in the luggage room and so the sound is actually better, but they also had people walking, coming into the luggage room, like our friend John Krakauer, for example, who attended the workshop.

Paul Middlebrooks
Yes. He makes a boisterous entrance, but it's brief. He was the main interrupter or whatever, so we say hi to him and you'll hear other things in the background and other people coming in and out.

Gaute Einevoll
We link to the homepage of the workshop.

Paul Middlebrooks
This is divided into two episodes, so in the episode that's going to come out in just a couple of days, you'll hear Cristina Savin or Savin or Savin, I'm not sure how to pronounce her last name, and Tim Vogels, and we'll talk a little bit more about them in the next episode. On this first conversation that we have, we have Andreas Tolias and Ken Harris. Gaute, do you want to talk about Andreas?

Gaute Einevoll
Yes. Andreas Tolias, he has just moved to Stanford by the way, and at least he's done many things. Obviously, he's doing monkey physiology and mouse physiology. Some of the work that I'm particularly interested in, is this foundation models he has made for the visual cortices or visual cortex of mice. Where he has trained these deep networks to essentially predict calcium responses when the mouse is shown different kind of visual stimuli. I think in terms of predictability, this is like really, really the state of the art when it comes to making these kind of models that predict things. Of course, in terms of interpretability, which we'll also discuss a little bit in the podcast, it's a bit different. It's harder to interpret. Then there is Ken Harris.

Paul Middlebrooks
Ken, among other things, he's interested in how large populations of neurons sense the world and convert that sensation into action. One of the things he's been known for in the past few years is just the immense recording capacity. In the past few years, our data has skyrocketed. The computational power has skyrocketed. Speaking of cutting edge, he was one of the first to use a super high-density recording electrode to record thousands of neurons at the same time. That has not been done before, so we're still grappling with what to do with all these neurons. We don't actually talk so much about their research in these conversations.

Gaute Einevoll
No, because he actually at the—

Paul Middlebrooks
Workshop?

Gaute Einevoll
- workshop exactly. Not podcast, at the workshop, he had been challenged to talk about actually what it means to understand something more on the philosophy of science.

Paul Middlebrooks
That's right. I think that's a fair enough introduction. All right. Enjoy our discussion and see you in a few days. Mikkel, you had this idea to put this workshop together that we're now at the end of, but we're going to go back to the beginning and ask you how this workshop came about, and what was it supposed to be about?

Mikkel Lepperød: I had the training in both computational and experimental neuroscience. I've been working the past kind of, I don't know, four years on neuro—

John Krakauer: I need to check in now, right?

Mikkel Lepperød: NeuroAI models.

Paul Middlebrooks
It's OK.

Mikkel Lepperød: It's OK.

Paul Middlebrooks
It's just John Krakauer, guys.

[laughter]

Paul Middlebrooks
He's number 15 on my podcast. It's 15th time you'll hear on the podcast.

Mikkel Lepperød: He has a characteristic laughter?

Paul Middlebrooks
Yes.

Mikkel Lepperød: Which we have enjoyed all through the workshop.

John Krakauer: I hope you mean it.

Paul Middlebrooks
I do. Are you leaving? Are you out?

John Krakauer: I'm going to check in.

Paul Middlebrooks
OK, cool.

John Krakauer: I'll have the dinner. Sorry.

Paul Middlebrooks
No, you're good. You've been studying computational neuroscience.

Mikkel Lepperød: I got really excited when I learned that we could— When I did my work in my Ph.D. and training. I was working on grid cells and spatial representations in rat brains foraging.

Paul Middlebrooks
It's a very Norwegian thing to do.

Mikkel Lepperød: It is a very Norwegian thing to do, I guess. I've always been interested in the "why" question: Why do we have these representations? What are these cells used for? My training from mathematics always drove me into comparing experiments with models. That is how I envisioned science would be for me. I tried to do that in some of the experimental work I did. I tried to perturb neurons with optogenetics. I perturbed the medial septum area while recording neurons around the cortex, and trying to see if I could say something about validating models back then, like the classic computational models of grid cells. I think this notion of validating models has always been a core part of my vision of how science should be done.

Paul Middlebrooks
How did the neuroAI thing come in?

Mikkel Lepperød: I started working on neuroAI models of grid cells, if you will, or navigation when there was a model coming out from DeepMind. Also in parallel, there's a paper from Cueva and Wei that did a parallel discovery, that if you train recurrent neural networks on path integration tasks, you give them velocity input and you train them to output like a position based on that. Then in some models, you can see a similar pattern that you see in real brains.

This seemed to be very, and I think it still is, it's a very interesting way of modeling these phenomena, because you don't put that many assumptions in. In the classical mechanistic models, you build in what you're seeing, and this is a slightly different approach. My hope with these models was that without putting in too many assumptions about what the system does, you could have these patterns emerge or become part of the computation that the RNN is doing. Then you could probe it afterwards and ask "why" questions. What are these cells doing, how are they interacting, and so on and so forth, because that's something you can't do in experiments very easy.

Working on these models, I realized it's not that easy. You can train a model and you can get a lot of different results. There was a big discussion whether or not this grid pattern would occur consistently. It doesn't.

Paul Middlebrooks
You mean over different training regimes and different architectures and stuff?

Mikkel Lepperød: Exactly. Yes. Even just the initialization, at least in the early models. Looking at that, I started worrying about, how are we supposed to relate these models back to neuroscience? That was my motivation, at least. I think it's the motivation for many that, you can use these models to say something about how the brain works.

Paul Middlebrooks
For the workshop, did you think, oh, I need to make some new friends here? Have some people to—

Mikkel Lepperød: One thing is that, this work sparked the notion in me that, we really need to do this in the right way in order to say something that has a scientific rigor. We can't just, generate lots of models and say that every model that has, sufficiently, gritty or whatever pattern you're looking for is a good model. That's sparked the scientific question. In terms of the workshop, I think many of these problems are— They would occur in many other systems as well. Looking at other types of neuroAI models that are being used, that are somewhat similar where you can train a vision model to recognize images. By linear probing, you can compare their activations with real neurons or just do this similarity analysis.

Paul Middlebrooks
Lots of different ways to compare the models with the neural activity.

Mikkel Lepperød: Exactly. I think in many of those cases, similar problems can occur where you can generate lots of different models. Even though they look similar to what you study in the brain, and they even look similar to between the models, but that doesn't necessarily mean that they're doing the same.

Paul Middlebrooks
Even if the output is the same, the way that they're doing it might not be the same.

Mikkel Lepperød: Exactly. It's kind of this algorithmic-level question of whether or not they're implementing the same algorithm.

Paul Middlebrooks
You gathered together a wide variety of people from different fields using different types of models to study different brain systems, but also using AI as a tool to analyze their data. There's even a philosopher here, one philosopher. That's enough for the conference. It was a nice group of people.

Gaute Einevoll
It was a fantastic venue you chose.

Paul Middlebrooks
Oh my God, yes.

Gaute Einevoll
There's this coastline that shows this— the coolest thing of Norway from nature point of view are the fjords. Even in the memory of "The Hitchhiker's Guide to the Galaxy," this guy got the planetary engineer who got an award for the fjords of Norway. I think that was a well-deserved award. Now with this trip going from the north to the middle of Norway or Trondheim on the boat was really showing of the best of the nature. I guess that was probably intentional.

Paul Middlebrooks
I've also heard multiple times, this is the coolest conference/workshop that I've ever been to. I'm sure you did too multiple times.

Mikkel Lepperød: I think most of the people I've talked to have said they really enjoy themselves. I'm very happy.

Paul Middlebrooks
You just made the bar very high for people to go to other workshops.

Gaute Einevoll
Now again, this is going to be at the start of the podcast. Now we're going to have, I would say, the main material would to— first one in this episode, we're going to have one pair of participants. Then in the next episode, we're going to have the other pair of participants. Then after that we want to come back to you, and then at the end of the second episode, to maybe think a little bit ahead.

Paul Middlebrooks
Through the magic of podcasting, let's get to the first episode, and then we'll revisit Mikkel.

Gaute Einevoll:[laughs] Yes, exactly.

[music]

Paul Middlebrooks
We're on a boat in Norway, Gaute and I have stolen you away to answer a few very broad, unfair questions, maybe, Gaute?

Gaute Einevoll
Really cool question. Unfair? I don't think they're unfair. Then maybe haven't thought so much. I don't know. I think if I had to answer to these questions without having thought so much about it, I don't—. Anyway, you are better than I am.

Paul Middlebrooks
These are broad questions. One of the questions that we had for you both was in parentheses, whether neuroAI, and we can talk about what neuroAI is, how has it changed the way that you ask questions or approach your scientific questions?

Andreas Tolias: That's a really good question, but—

Paul Middlebrooks
Let's try and maybe first—

Mikkel Lepperød: This is Andreas speaking.

Andreas Tolias: This is Andreas.

Mikkel Lepperød: If you hear slight Greek accent then it's Andreas.

Andreas Tolias: For me, neuroAI is defined in two ways. One is using the modern version of AI, which is deep learning with large data and large compute to build models of the brain. The way that I think it's impacting the research that we do, one way that is impacting, is essentially it has changed our thinking of embracing high-entropy data. Then using the tools that have been developed to feed models—

Paul Middlebrooks
What is high-entropy data?

Andreas Tolias: Meaning you sample, let's say, natural images without being very hypothesis-driven. You do naturalistic behaviors, and you record without trying to control the behavior a lot. They build a lot of tools, a lot of them are engineering tools like GPUs, libraries, PyTorch, and stuff like that, that enables us to feed this large-scale data and basically use them as ways to extract statistical structure from the data. That's one way that it's been impacting us.

The second way is that this field has been developing tools, once you have a neural network that predicts something, they are trial-developed tools to try and get some interpretability, the whole field called mechanistic interpretability, how it works. We're incorporating their tools. These are the more going from AI to neuro. On the other hand, of course, as neuroscientists, we're always interested to build intelligent systems. That's a much harder thing, but it also has helped us think about what are the type of tasks that may be important and what is the advantages of brains versus AI, for example. This stuff about generalization, robustness, adversarial robustness. I think it has been a fruitful interaction between AI and neuroscience.

Paul Middlebrooks
Can you contrast that with the way that you used to do science?

Andreas Tolias: Yes, it's very different, because the way we used to do science, we were always limited by data, both in terms of, you, let's say, wanted to record from neurons in the brain. Even if you could record from any neuron, you had, let's say, an hour or two hours to learn an experiment, and you were developing some hypothesis that you were testing. Whereas now it allows us to control, do more like non-hypothesis, more data-driven science. It has changed from, I would say, hypothesis-driven science to data-driven science.

Paul Middlebrooks
Maybe, Ken, let's just ask the same question, because you have a different perspective on this, I think.

Kenneth Harris
Yes. I'm a bit more old-school, perhaps. We certainly use AI technology, for example, to do video processing, deep lab cut software, and other things like that.

Paul Middlebrooks
As a tool.

Gaute Einevoll
As a tool.

Kenneth Harris
As a tool, to let you do science by doing video processing that wouldn't have been possible a few years ago and that sort of thing, certainly. In terms of AI for informing scientific questions and conclusions, for me, I'm less up on the more—. We certainly read stuff like what Andreas publishes using the most recent techniques. For me, the really valuable concepts at the moment are still those from a few years ago, such as kernel machines and variational Bayesian inferencing. Things like these we just use very fundamentally in the way we think about things.

Gaute Einevoll
Why do you stick with them, because you understand them better?

Kenneth Harris
Exactly. Because these are the things that you can understand. Because a kernel machine, you know exactly what it is doing. You know how it works. A deep network, it's found a solution, but I don't really understand how it works. Maybe the point is that's the same with the brain. Maybe we're never going to understand how the brain works in the same way that we don't understand how a deep network works.

Paul Middlebrooks
By the way, you guys interrupt each other and argue with what you're doing.

Andreas Tolias: I agree that there's pros and cons here. Right now, we're basically putting a lot of emphasis on building a very accurate model of the brain in silico that's differentiable and stuff like that, that then we hope we can then analyze to understand it. What Ken is saying, which is another approach, and of course, it has other advantages, is like you start with a model that is already building interpretability. Then if you feed the data, then the understanding falls out of it. Whereas the deep-learning approach that is more data-driven, you basically emphasize the predictability first, and then you hope that by looking inside the— The other thing is very good to compare the two, because if, for example, you use an interpretive model that only explains half of what the deep-learning model does, it also says that there's really a lot that we don't understand.

Gaute Einevoll
Because I remember I've read some of your papers in terms of how you train these deep networks to predict these two-photon calcium responses to all these visual stimuli in mouse. When you compare with the previous approaches, were like, Gabor filters and whatever. In terms of predictability, nothing compares to the predictability of these train models, but that maybe comes at the expense of being not quite sure how it works. Maybe with the interpretability— What's the interpretability of your models, Ken? Of these kernels?

Kenneth Harris
The kernel is just a measure of similarity between how similarly a population of neurons responds to any pair of stimuli. It's just a number for every pair of stimuli. It tells you how similar the population responses to those two stimuli. There's a whole field of machine-learning theory that went out of fashion about a decade ago that did a lot of very useful work in understanding what representations those give you. You can use all of that to understand what's going on in the neural code of the brain.

Paul Middlebrooks
You made a point earlier on vaguely, that it's changed your science in terms of making it more data-driven. I hear a lot that what we need is more theory with big compute, with big models as tools, with big data. How do you think of that? We're in this weird space, where you have to explore to then generate theory, maybe even. I don't know. How do you think about it?

Andreas Tolias: No, it is something that I think a lot, and it bothers me in a way, but also it's like we are at this stage where the way we are doing science and engineering with deep learning, it's a little bit non-classical scientific. In that we figure a way, a hacky way, to basically build models that can drive cars around, but we don't really understand them in a classical scientific way how they do it. To some extent, we understand the loss function, we understand something about the architecture of the network, but we don't have that algorithmic understanding in a more classical way. That is an issue. That's a problem.

Now, the question that I hope, and it's a hope, is that by building very highly predictive models that can generalize very well, and I think generalization is key here. Then the fact that they are differentiable. The fact that you have a model that you can do any experiment you want that is actually a neural network, which is like the brain, even if it's not implemented in the same way, but it still has synapses, synaptic weights, activities. Then we just have to develop tools and the AI community is already doing that because they care about that, too. This is one of the issues about safety and robustness and channelization. These are key things in AI.

That then we can leverage what they're doing, what we develop to try and gain some more understanding. I think once we have that understanding, then it may be possible to then think about more interpretable models that are going to be simpler, that then we'll fit the data, and then we'll get to bring it back to what Ken is talking about.

Paul Middlebrooks
Do you believe that, Ken?

Kenneth Harris
I don't know. If you look in the history of science, there's some cases which are a lot of cause for optimism. For example, if you were an astronomer in the days of Tycho, but before Kepler, you might have thought, there's no way all of this data is ever going to end up being simple. A few hundred years later, you have Newton's laws. You've got one equation that can explain everything. Same if you're a chemist before the periodic table, you would never have guessed it was going to be that simple.

On the other hand, if you're a biologist before the genome, you are then confronted with the fact that there's 20,000 genes, and they're all different, and you're never going to know what they all are. We don't know for neuroscience whether—

Paul Middlebrooks
There's a difference. Those weren't complex systems. The DNA code readout is not a complex system, whereas, with brains, we're dealing with a complex system. Does that difference make a difference, do you think?

Kenneth Harris
The solar system is a pretty complex system.

Paul Middlebrooks
It's complicated, but it's not a complex system. Well, I guess ... three-body problem.

Kenneth Harris
What's the difference?

Paul Middlebrooks
Complicated just means hard, lots of parts. Complex means interacting parts where there's emergent properties [crosstalk]

Gaute Einevoll
It's the stable. You can predict the orbit of Venus 500 years from now.

Paul Middlebrooks
Oh, you mean chaotic?

Gaute Einevoll
Yes, chaotic. I think complex systems are often chaotic. Not always. That's true, we are not chaotic completely.

Paul Middlebrooks
That's one often property of complex system.

Gaute Einevoll
That's true. Anyway, yes. Also, maybe that's—

Paul Middlebrooks
Your point was things have looked impossible over and over in the history of science, and maybe this is one of those things.

Kenneth Harris
Things have looked impossible over and over. Sometimes they weren't. Sometimes, so far, they still appear to be.

Paul Middlebrooks
We're in that regime.

Kenneth Harris
We don't know.

Gaute Einevoll:I think if I can come back to what you said, Andreas, that now you're training these deep networks, difficult to understand, and you hope to get them. That is all forms of basis for more interpretability. Then, of course, we have better tools for understanding how the brain works. At the moment, given what you have at the present stage, has it told you anything new about the brain or cognition that you—?

Andreas Tolias: I think that's a really good question. I think it's still early to know if the way to the future is like, this is going to be the standard. There are cases where, for example, starting contextual modulation, in the visual cortex, people have been studying it with the gradients and they found a specific relationship between center-surround interactions. Then when we followed this image synthesis using these deep-learning methods, we got something that was different. The nice thing here is that then you can verify it back in the brain, you can run this closed-loop experiment.

Paul Middlebrooks
Oh, absolutely. Inception loops, right? That you've been doing.

Andreas Tolias: Inception loops. Yes. That's one example where it's not a circuit-level mechanistic model, but the description of how the center and surround interact, we gain some new understanding that—. The other thing is like once we do that, then you can design experiments in a more classical way. Because now you develop the hypothesis, then to test it and then do exactly what Ken has been talking about. In fact, we did build a model like that. We built like a more Bayesian model, like trying to explain center-surround interactions based on natural image statistics and priors.

Kenneth Harris
In the primary visual cortex.

Andreas Tolias: Yes, in the primary visual cortex. I think that's one example where you start with a data-driven model that there's no interpretability, it's just sort of an engineering task. You analyze it, then you derive a principle or some level of understanding, you test it back empirically to see if it's correct. Then you build a simpler model that is more based on more classical stuff, like if this case was hierarchical Bayesian inference. And say, can a model of hierarchical Bayesian inference trained on natural images predict at least qualitatively the same type of effect? I think that's an example.

Paul Middlebrooks
You started your talk today with the clip of 60 minutes of Hubel and Wiesel serendipitously discovering the edge that— so they were moving a piece of paper, they were trying to test the visual neuronal response to dots. They happened to move the transparency or whatever it was off the screen, and part of that transparency, there was an edge where that transparency ended. That's how they discovered edge cells, right? That's a data-driven approach, or an exploratory approach, but it was also serendipitous. It made me think, right when you were showing that, are we past that stage where we're going to— How does serendipity play a role these days? Ken, you're shaking your head.

Kenneth Harris
All the time, almost everything we do is serendipity. You get your data, normally you have a hypothesis in mind, and when you get your data, you realize, oh, actually, wait, that was never going to work anyway. Then you notice something a bit odd in the data, and it's probably a bug, but you chase it up, and, well, it doesn't seem to be a bug. Then you follow it up a bit more, and then there's something you don't understand, and you don't understand why you would see this, and then you try and figure out why you're seeing it.

Paul Middlebrooks
It's the analog of, oh, what's the famous quote when science progresses by saying, huh, that's funny.

Kenneth Harris
Exactly.

Paul Middlebrooks
Does anyone remember who said that?

Andreas Tolias: No.

Paul Middlebrooks
I'll have to look it up later. OK.

Gaute Einevoll
Is there any way that AI has set neuroscience back? Or put some or many or all neuroscientists on the wrong track?

Kenneth Harris
It's probably led to some blind alleys, but the problem is we don't know which ones they are yet. We will in the future.

Paul Middlebrooks
I worry that in a lot of, especially young researchers' minds, might confuse the map with the territory, in terms of people thinking of the brain as a transformer. Substituting the model for the real thing, and then that conceptual framework then frames their research questions, and I wonder if that has had a deleterious effect at all.

Andreas Tolias: Yes, it's still early to know where this thing is going to progress right now, because there's also this two-way interaction. Some people are working, it's like using tools from AI or as a model of the brain. I do think that there is some danger here of basically using it as an engineering tool that is sort of the end game, versus it's just the beginning of the end. I do think it's important to just see it as a tool right now, and not take it as just sort of like, OK, if I just take a model and I build it and it works great, I'm done, I should move on to something else. I think there is a little bit of that.

Gaute Einevoll: It's a lower threshold for starting doing it, right? If you have tools like whatever, PyTorch, TensorFlow, you can sort of quite easily train networks to do something, so it's maybe it's sort of you- If you do traditional statistical analysis, it takes more effort to get into it, maybe. Certainly compared to the physics type modeling, which is the threshold for getting into it, is more basic tools. I guess low threshold is both good and bad. It means that it's easy to test, but it's also easier maybe to do things which are not that high quality.

Paul Middlebrooks
Is there a danger that it introduces less critical thinking from the beginning?

Andreas Tolias: There is definitely that danger from AI in general. There are people that are trying to build, let's say autonomous AI scientists to analyze the data. I don't think it's going to happen around the corner, but there is—

Paul Middlebrooks
It might replace me. It probably won't replace you guys.

Andreas Tolias: It is a niche, it's easier to— But on the other hand, that's good. I think historically, people were worried about when calculators came around, people were not—. I don't know. I think it's important to just use it as a tool, but make sure that people are educated and they're doing the critical thinking.

Kenneth Harris
Every new technology introduces new ways of doing things wrong.

Paul Middlebrooks
You don't think this is different at all, in that respect?

Kenneth Harris
Not really. No.

Paul Middlebrooks
More is different.

Kenneth Harris
Well, it's easier to fool yourself when the thing trying to fool you can speak English.

Mikkel Lepperød: Fluently.

Gaute Einevoll
That's true.

Paul Middlebrooks
By the way, just as an aside, I have to commend you, because you worked on a talk and I think against your will, the talk that you gave, you really put thought into it and tried to address-

Andreas Tolias: The talk was on? What was the talk on just to get a sense?

Kenneth Harris
The talk was on what would it mean to say that a system such as a deep network is a good model for the brain.

Paul Middlebrooks
Did that drive you crazy creating that talk?

Kenneth Harris
Well, I thought at first I was going to say it's impossible and ended up after thinking about it, thinking that it could be possible. I don't think it's been done yet, but it's not actually impossible. What it would need is an interventional experiment, where you have a mapping between the artificial system and the brain. That not only shows they have a similar representation of information, but also if you perturb them, you perturb the network that causes a change in the rest of the activity in the network in a way that maps onto the way a perturbation and brain activity would. Then you can say there is a mechanistic similarity in how they're computing information, not just a similar representation.

Paul Middlebrooks
You also said that a model is something that shows what is not possible.

Kenneth Harris
Yes.

Paul Middlebrooks
How does that fit in?

Kenneth Harris
Oh, I think that's just what science is.

Paul Middlebrooks
That's falsifiability.

Kenneth Harris
Falsifiability, exactly.

Gaute Einevoll:I was trained as a physicist, and there, sort of, the uh ... but there it looked a little bit down upon by active physicists to talk about philosophy. It was all this idea, shut up and calculate.

Kenneth Harris
Instinctively.

Gaute Einevoll: Yes, but are we talking too little about these questions? What does it mean to believe in explanation? This is some sense of being our job for the last decades to do science, which is about this. Still, it is we don't really formulate it too often to ourselves, right?

Paul Middlebrooks
Yes, well, that's why I was asking Andreas, too, if we're in a weird space, because you have to explore with all these new tools.

Andreas Tolias: To me, we are in a more weird space than I've ever experienced. It feels different, something collectively different than before.

Kenneth Harris
Why do you say that?

Andreas Tolias: I think the reason is because we are capable of building systems and models that seem quite intelligent, very intelligent, have capabilities. I bought a Tesla, and I drive it, and I'm really impressed by it. You can drive for hours in a very complex environment, and it's really impressive. Then, at the same time, it's a neural network that nobody really understands in a classical scientific way. I do think it's a technological advancement of these two things. We have data, I'm talking not just in neuroscience, but generally in biology, in medicine, the way you have cameras everywhere collect data, the internet. Basically, we've created a society where we made it easy to collect a lot of data, and then we have computers that have become very fast. We put these two things together with neural networks and we build all these very complex systems that are capable of doing complicated things, but again, we don't understand them. I do think there's some danger here.

Kenneth Harris
You're not talking about science, specifically, you're talking about the state of the world?

Andreas Tolias: I'm just talking about the state of, yes, basically the whole— If you look at, let's say AI in general. A neuroAI is just neuroscience using AI. You can think of, let's say, the people who predict the stock market using AI, people that are trying to build autonomous driving. They were doing it before.

Paul Middlebrooks
Neural networks.

Kenneth Harris
Basically, it's just engineering.

Andreas Tolias: Yes, all these areas. Even in physics, they're using AI.

Gaute Einevoll
Yes, because Feynman analysis said that, "If I can't build it, I don't understand it, or if I understand it, I can build it." He didn't mean physically build it, but meaning make a physical model, physics-type model that in a sense is going to—

Andreas Tolias: [crosstalk]

Gaute Einevoll
Now, we're— Let me finish. Because now we put in these learning models and so we build things we don't understand.

Andreas Tolias: Yes, exactly.

Gaute Einevoll
Which is also a new thing and which is very fascinating. Now, we've made this fantastic large language model, and we're just thinking about what it can achieve. We don't understand it. It's like a research project, you do research on it just like we do on the test animal, right? Which is really fascinating. We've built something, it's man-made, but we don't understand it.

Andreas Tolias: My impression is we went through that before, somewhere during the Industrial Revolution. The first steam engines were built without really understanding, let's say, the dynamics.

Kenneth Harris
That's true.

Andreas Tolias: In fact, they were not safe. They used to blow up all the time.

[laughter]

Andreas Tolias: It only became, after people started doing measurements, especially after discovering law of thermodynamics, temperature, and pressure, things became so. Maybe we're going through that phase as well.

Paul Middlebrooks
Well, at least we know that AI is safe and doesn't affect society.

Gaute Einevoll
Exactly. Nobody worries about that.

Paul Middlebrooks
No worries.

Gaute Einevoll
That's really interesting because the electromagnetism there, the revolution came after Faraday and Maxwell, like these people. It's true, the thermodynamics came after the steam engine. They made things like entropy design term and energy conservational stuff.

Paul Middlebrooks
If we are in a weird space right now, just scientifically then one of the questions that we were going to ask you was how long are we going to be in this space? Will AI still have a role in 50 years? Or will we have used it to solve lots of brain problems? I'm not going to say solve the brain or solve intelligence, but will it have been a key factor, or is it a fad? Will it go away?

Andreas Tolias: I'll tell you what. I think about this a lot. I don't think I've never been in a situation in my life that I'm really curious what's going to happen in 5 or 10 years. Are we going to really have a GI that's going to basically solve not just very difficult questions for the brain and neuroscience, but let's say health, I don't know, climate change, or is it going to be like the iPhone? When the first iPhone came, it was amazing. Now we have, I don't know, what is it? Model 13. It looks similar. These better apps, but it's large language models, all this thing is just a fad, as you said, then in 10 years, we're going to be like, yes, you know.

In fact, in the last two years, I wouldn't say it's impressively amazingly getting much better. If you look at ChatGPT 3.5 versus 4.0, it's a little bit better. It does some things better. Or it's new applications. It can do video now, but it's not a collectively new type of thing. We're getting incremental changes. Now that doesn't mean it will remain like that, but I think it's very interesting because it's like— In some days I wake up and I feel like, "Oh my God, this is it." Others are like, "Yes, not much is going to change."

Kenneth Harris
It depends on [crosstalk]

Andreas Tolias: One thing is that this stuff relied a lot on the scaling loss. I do believe that there's a big problem that these big companies are facing now. They're just running out of data to train the models.

Paul Middlebrooks
Scaling in terms of computers or data?

Andreas Tolias: Both. The scaling loss, it's both larger computers. Basically more energy and more data. They're synthesizing data, but it is a niche.

Gaute Einevoll
A little bit of the ancestors.

Andreas Tolias: Yes, exactly. A lot of the stuff they're doing now is data selection, because there's a lot of— they've shown that if you select the type of data you train, let's say a large language model or a video model, it may make a big difference. It does feel like to me that there could be things like robotics with VR, people wearing cameras and collecting their body movements and imitation learning, but it's possible that we are going to have incremental improvements, or maybe it will be a new revolution, I don't know. I think that's going to impact us too.

Kenneth Harris
I think you're right, even if there isn't another set change though, just the one we've already had.

Paul Middlebrooks
We don't even know how to deal with it yet.

Kenneth Harris
It is creating a lot of changes to societies in ways that are going to be unpredictable like the Industrial Revolution. For example, a friend of mine works in classical music, she has a lot of friends who are composers who used to write for film scores and now they're out of work because that's something that AI can do pretty well, is write classical music for film scores. Who would have predicted that? Society is going to change in a lot of ways. There was, what was it called, there was something that had an AI girlfriend app, I think that's gone away.

Paul Middlebrooks
Japanese thing, I think, that's popular.

Kenneth Harris
Yes, but there's all sorts of ways that society may really change that we just can't predict, even if there isn't another revolution.

Gaute Einevoll
They'll go back to the science because I'm doing physics-type modeling, sort of a little bit like, sort of like extensional Hodgkin-Huxley and multi-compartment models and networks, and so on. It seems now that, and computation neuroscience has always been a small subfield of neuroscience. Now I also feel that maybe computation neuroscience is, because you have all this AI, if you want to go into computation neuroscience, maybe many people go into this using these AI tools, so that in some sense, is physics going to be traditional computation neuroscience? I had this hope and now this idea when you push this learning roots and stuff, at some point to understand more, we have to get back to the biophysics of neurons and stuff.

Kenneth Harris
It'll be important in the end.

Gaute Einevoll
Yes, sure.

[crosstalk]

Kenneth Harris
This is an important point that what physical brain neurons do is so much more complicated than threshold ReLU units. That's surely important.

Paul Middlebrooks
For what though? Maybe not to explain cognition, maybe it's important for— For what question is it important?

Kenneth Harris
OK, if you took your brain and you replace every neuron with a ReLU unit, I would predict your cognition would be severely impacted.

Paul Middlebrooks
They still get plasticity though?

Kenneth Harris
Yes, you just don't get any iron gate voltage-gated iron channels and intrinsic oscillations.

Paul Middlebrooks
My brain again might not be that different.

Kenneth Harris
Yes, I think that's an important point. The physical brain, for whatever reason, has all these different cell types. The cells are very complicated, maybe it didn't need to be like that to have an intelligent system. We don't know, but it is like that. Neuroscience, whose goal is to understand our brains rather than just to come up with an intelligent system, that stuff is going to matter.

Paul Middlebrooks
Isn't there a limit to the biophysical detail that you would need to implement and can't you just figure out if there are a thousand different types of neurons in the brain, give them a thousand different activation functions and other algorithmic profiles, for example.

Kenneth Harris
You're saying you don't need to simulate every single channel.

Paul Middlebrooks
Where is the bottom layer that matters?

Andreas Tolias: I agree. I think that it is a very interesting question. I don't know the answer, but one possibility is the following. To really understand cognition and behavior, you really don't need to go down to the ion channels and the nonlinearities in the dendrites at that level, because just the bridge too far is too complicated. Just understanding it at the representational level and manipulating it, maybe as you said, Ken, at that representational level.

You have, let's say, an AI system that represents information, and you have the brain, and you study that on this representational level, and you manipulate them on this representational level, and you try to understand them, it's enough to understand the science of intelligence or the loss of intelligence and build systems that are intelligent. Maybe they are computational neuroscience in a traditional way, maybe in trouble, but there is another way that is very important, which is diseases.

In order to intervene, and maybe that's what Ken was saying, in our brains, let's say, psychiatric diseases, neurological diseases, they're not ReLUs. They have ion channels, they have molecular pathways, they have dendrites. If you look at autism, the spines are different, and there is where maybe people are doing more, let's call it physics-based or more classical biophysics, computational neuroscience should maybe pivot in that direction, and that field probably needs people, like computational neuro.

Also, that's where probably, if you look at it, even from a practical point where the funding is, and what let's say NIH cares about, is to cure diseases more than to really understand the algorithm of perception. I do think there is an opportunity.

Gaute Einevoll
I'm going to quote you in my next application, like the great Dr. Tolias.

[laughter]

Paul Middlebrooks
Even those details, I think that's a good point about diseases and how those lower-level reductive small molecules, dendritic shapes, and spine sizes, etc., matter at that scale. Then just like an artificial unit, if you can abstract what's important about how that affects communication, that might just solve it.

Kenneth Harris
Yes, you won't need to know every single potassium channel, but I think there will be, to make an accurate model of how the human brain works, you will need to incorporate things like dendritic nonlinearities and oscillations in some simplified form. That's what I would [crosstalk]

Gaute Einevoll
Also, if you think still we haven't really figured out what mind or whatever conscious of this feeling of mind, which is different that we are not only looking at statistical relationships when we infer things. We have this first-person perspective. It's more to our intelligence that just the things that are picked up by AI, I would think. Saying that maybe you need these kinds of things too.

Kenneth Harris
Well, it's funny, I asked a large language model the same question that said the same thing.

Gaute Einevoll
Really?

Paul Middlebrooks
Going back to—

Gaute Einevoll
Did they say that they had a mind or that they didn't have a mind and we need it for better brain model or something?

Kenneth Harris
No, I'm just saying that the fact that you just said all of that, any agent can say that.

Gaute Einevoll
Yes, but you believe that I'm conscious, don't you?

Kenneth Harris
I don't know what that word means.

Gaute Einevoll
Really? But they—

[crosstalk]

Paul Middlebrooks
Oh, no, let's not do it, let's not go on there.

Kenneth Harris
OK, one sentence on the topic of consciousness. The moment that you actually define the word it becomes quite a boring question.

Gaute Einevoll
I think it feels like something to be Ken, I think it feels like something to be me and Andreas and Paul. I'm not quite sure if it feels like something to be a large language model.

Paul Middlebrooks
I don't think it feels like anything to be Ken. Just kidding, Ken.

Kenneth Harris
No, that's correct.

[crosstalk]

Gaute Einevoll
Anyway, we went off to the consciousness.

Paul Middlebrooks
Yes, and we're at 35 minutes.

Gaute Einevoll
Yes, but let's switch to what?

Paul Middlebrooks
I just don't want to take up everyone's time and everything, so I'm not going to switch to the—

Gaute Einevoll
Oh, they have plenty of time. Don't worry.

[laughter]

Gaute Einevoll
They're on the boat, where should they go?

Paul Middlebrooks
OK, go for it.

Gaute Einevoll
We have one question. I just read this book called "Slow Productivity." I think it's one of these. It's almost the book that you pick up at an airport, but it was interesting in the sense that what does it need to be an efficient knowledge worker, which is different from being an efficient farmer because then you measure the output. Then I thought, what does it really mean?

The book claimed that since it's not easy to define knowledge productivity in knowledge workers, you make these proxies to look busy. You work for a company, and they just look busy when the boss is coming. Maybe some of these things when you look at how we survive in science is having many papers. This has become a proxy for productivity, not only numbers but also the quality.

Paul Middlebrooks
Let me just add also, especially in the neuroAI space, one way to look really productive is just to throw a neural network at whatever problem that you have without considering the theoretical framework or questions or hypotheses. I just wanted to preface it.

Gaute Einevoll
Yes. I was thinking what does it really— I didn't think too much about it, but what does it mean to be productive in science?

Paul Middlebrooks
Besides paper—

Gaute Einevoll
It's like these mundane things of surviving and funding grants and getting a job.

Andreas Tolias: I think it's a problem. This is the way science is. The way that I like to be justified is that we humans didn't really evolve to science. We evolved to the other stuff. Science is something that we started doing in the last, I don't know, a few hundred years. It's remarkable how much we've advanced science as humanity. Whatever we've been doing, even if it's at any given point in time, any one of us, it looks like it's very incremental and we're not doing much.

As a species, we've done tremendous advancement. Now, does it need improvement? I think it does. One of what you are saying is maybe the way our reward system is, the classical academic system may need to be reformated to some extent. Some of these questions may require more teams working on it, where credit is. It's not about just publishing the next paper, but it's more about working on a project longer, maybe, and allowing people to work as a team on a project longer and having a way to reward them. Maybe allow some more risk.

Paul Middlebrooks
Do you mean personally productive?

Gaute Einevoll
Yes. I understand people who don't have permanent jobs and want a permanent job, that's a different thing. I'm a professor with a permanent job, so I can, in principle, do anything I want, but still, as a professor, you easily get involved in too many projects because you don't want to be— It feels good sometimes if I'm doing this as well.

Paul Middlebrooks
You get roped into podcast conversations.

Gaute Einevoll
Exactly, but I mean it's sort of this thing—

Andreas Tolias: There is a bit of a danger in this.

Gaute Einevoll
Yes, I think it's just something that my own sort of, the psychology of people makes us be very stressed sometimes. Ken is breathing here. Often, when you have several Ph.D. students and you have some feeling of responsibility, you maybe have to spend most time on the project which is working the least, so you have to try to salvage it. Sometimes I think, I would much rather work on the thing that this person is doing because that's all great. It's just something about the psychology of how we make choices.

Kenneth Harris
The best thing for my productivity was when I formed a team lab with [Matteo Carandini](). It's just something about that that when you got two people, I might have an idea I think it's great, but then he says, "Wait, but did you think X?" Then you save so much time just by that. You have a lot of collaborations as well and it's just a similar thing.

Paul Middlebrooks
Do you also feel the responsibility to be productive for the other person? Does having a collaboration make you more likely to do the work you need to do the work?

Kenneth Harris
Yes, because there's someone to remind you, "Hey, didn't you say you were going to—" "Oh, yes."

Paul Middlebrooks
Put your self-worth in someone else's hands.

Gaute Einevoll
But does that mean that you sometimes are just getting stressed and not able to think?

Kenneth Harris
Yes. When was life not like that?

Gaute Einevoll
No, but is it optimal? That sort of the—

Paul Middlebrooks
There's the Dodson-Yerkes curve. Is that the one?

Gaute Einevoll
I heard, for example, Francis Crick. After he sorted out the DNA and also this, whatever the other proteins—

Paul Middlebrooks
Codon.

Gaute Einevoll
Codon. Exactly. Then after he wanted to stay under-busy. I think he just wanted to have so few projects that he could always jump on something which was really exciting. Of course, then he's in a position where you don't have any social constraints or practical constraints. Is this thing that— No, I don't know. I've been thinking about that lately. What is our external pressure and what is just internalized? We make silly choices because we are just internalized, some kind of behavior.

Kenneth Harris
We surely do. The moment you try and come up with how to change the system, the things you're saying about these changes and incentives, it's all good, but these are minor. This is all fairly minor changes. The system isn't perfectly efficient, but when you try and think how to make it better, it's quite hard to come up with anything radically different.

Gaute Einevoll
Absolutely. I'm more thinking that even when people have a Nobel Prize, they should, for example, then try to be out of his hat.

Paul Middlebrooks
I told you that. I said that to you. I said, well, here's for the quick example, he's already famous. He can do what he wants to do.

Gaute Einevoll
That's true. That means I can do what I want also. I think that many people at some stage in their career—

Kenneth Harris
There's loads of examples. Maxwell, the physicist-

Paul Middlebrooks
James Clarke.

Kenneth Harris
-spent somewhere like 20 years trying to find a particulate explanation of electromagneticism. There's so many cases, but you don't know what's the blind alley at all.

Gaute Einevoll
What would happen if you decide to, now we're going to just focus on one particular mathematical problem related to neuroscience analysis. You wouldn't be fired, right? Would you, or not?

Kenneth Harris
I'd have a conversation with my boss at some point. She'd say you used to get a lot of grants.

Gaute Einevoll
[laughs] Exactly. What about you, Andreas? You have actually changed, you just moved to Stanford, so you're trying to make a good impression.

Andreas Tolias: It's the same. We can do whatever we want, but if you keep, especially in the U.S., the system is very competitive. It's true everywhere, but I'm saying it is a little bit like running a little startup, especially in medical school, so you have to bring in resources. I do also think what you said about Crick is interesting because I think we're still at the stage where, obviously, the most important thing is to choose the right question. It's not clear this is the question. There's always a lot of possibilities. You have to remain focused, but also remain open-minded enough to see what new things may be happening.

Paul Middlebrooks
Here's a different way to ask the question. What makes you feel unproductive?

Kenneth Harris
Web surfing.

Paul Middlebrooks
Web surfing, OK. I mean in science, not like this, and you cannot answer this conversation.

Andreas Tolias: I think sometimes you can be unproductive. It's hard to know. Sometimes you feel you're unproductive because you're like daydreaming or you're like [crosstalk]

Paul Middlebrooks
A year goes by and you feel either it was a productive or an unproductive year. What do you think made it productive or unproductive?

Andreas Tolias: I think it's complicated. There could be years where you may not be publishing as much, but you feel you are very productive because you are doing your experiments, you have new results.

Paul Middlebrooks
You're laying the groundwork.

Andreas Tolias: Yes, so I wouldn't say there is a metric.

Kenneth Harris
I've spent a lot of time over the last few years on statistical methods that a few of them I've written preprints. One of them I've even submitted a peer-reviewed paper. I do think it was worth it. Maybe I'm at that stage now, I can do what I want. I've spent a lot of time on these questions of statistical analysis.

Gaute Einevoll
Then it seems you actually educated yourself also.

Kenneth Harris
Oh, yes, absolutely.

Gaute Einevoll
I think for me, an unproductive project would be something I just joined, even though I'm not particularly interested, not particularly challenging, but it has some reward in terms of maybe getting on that publication. You spend your mental resources on the—

Kenneth Harris
Things you don't really find interesting.

Gaute Einevoll
It's not really building up to anything. I don't really need a new AI for tax laws or whatever.

Kenneth Harris
But this is an actual project.

Gaute Einevoll
Yes, I don't know. No, no, this is not an actual project, but I was thinking if—

Paul Middlebrooks
Oh, an example.

Gaute Einevoll
Yes, an example of something which is important, but not really.

Paul Middlebrooks
Let's just move on the next perhaps closing question. By the way, the quote was, it's from Isaac Asimov. I did use ChatGPT to look this up. Hopefully, it's correct. "The most exciting phrase to hear in science, the one that heralds new discoveries, is not 'Eureka!' but 'That's funny.'" That's it. Isaac Asimov. I'm glad I didn't say Einstein because [crosstalk].

Gaute Einevoll
We had this question on what we prepared yesterday. I've been thinking of this question, that if for some reason there was a moratorium, it was almost like the pandemic where you couldn't do experiments for a year so that everybody had to work on existing data. Would that be a good thing? Because now I have a feeling that it's all these things that you measured and you don't really understand it. Let's put in another mouse.

Kenneth Harris
Can I go beyond that and say that you weren't even allowed to analyze existing data? All you could do is read their literature.

Paul Middlebrooks
You got shut down and that's why everyone, like 90 percent of neuroscientists wrote their first books during the pandemic. There's a bunch of books.

Gaute Einevoll
Exactly. Would it be a good thing? I don't want to live in a society where this is happening.

Kenneth Harris
If all I could do was read papers, I'd be very happy.

Paul Middlebrooks
For how long though?

Kenneth Harris
Oh, forever.

Gaute Einevoll
A year.

Kenneth Harris
Yes. I don't need to write anything.

Paul Middlebrooks
Don't you want to do the science?

Kenneth Harris
You talked about Francis Crick. That's basically what he did in neuroscience. He just read papers and tried to figure things out and then he—

Paul Middlebrooks
He wrote a book.

Kenneth Harris
He wrote a book, he wrote review articles, but that's all he did. It was very valuable.

Andreas Tolias: It's true. I think different people, I think we need to do more experiments.

Gaute Einevoll
I'm more thinking about the moratorium to think about what we have.

Andreas Tolias: Yes, I think stopping and thinking about things is very, very important.

Kenneth Harris
Yes. None of us do it really now.

Paul Middlebrooks
It's so much easier to do the next experiment than to think.

Kenneth Harris
Exactly.

Gaute Einevoll
It's also when I do simulations and I change parameters and it doesn't really work. Then I, oh, let's try another parameter set because it makes me feel like I'm doing something. While actually, I'm just sort of trying to avoid to think.

Paul Middlebrooks
We shouldn't have another COVID? Is it that the take-home? Or we should?

Andreas Tolias: I don't think we should.

Gaute Einevoll
No.

Kenneth Harris
For several reasons.

Andreas Tolias: I think there's a lot of pressure to stay in the rat race. I think figuring ways to get outside the rat race and think, it's always important.

Paul Middlebrooks
I have maybe a quick question before our last question, if you guys have a few more minutes. How do you know when you have a good idea, scientifically? When you feel like, all right, this is a good idea, without much vetting of the idea?

Andreas Tolias: Before you ask your experimental collaborators.

Kenneth Harris
Now you just know.

Paul Middlebrooks
Is it just an intuition? Don't you have those intuitions that turn out to be bad ideas?

Kenneth Harris
Oh, yes, lots of times.

Paul Middlebrooks
Then how do you decipher if you—

Kenneth Harris
You wake up. They normally come in the middle of the night.

Paul Middlebrooks
Right, or in the shower.

Kenneth Harris
Yes. Then if you still think so that afternoon, that's a pretty good sign.

Paul Middlebrooks
OK.

Andreas Tolias: I agree. Sometimes, yes, you think you have a great idea, and then you think more about it.

Paul Middlebrooks
Isn't that the best feeling in the world, though, when that thing hits you, and you're like, "Oh!"?

Gaute Einevoll
I had this great business idea in the middle of the night. I woke up and I had to write this down. When I looked at it in the morning, it was just-

Paul Middlebrooks
That's good.

Gaute Einevoll
-completely ridiculous, but that was I felt in the middle at 4 o'clock, 4 AM.

Paul Middlebrooks
But that's good because you could falsify it. Because I fail to write mine down all the time, and then it goes off. I think, well, I still come up with good ideas, but who knows? Because I've forgotten them.

Gaute Einevoll
Mine was not very good.

Paul Middlebrooks
So, intuition.

Andreas Tolias: Intuition is key, yes.

Paul Middlebrooks
OK.

Gaute Einevoll
Yes, so the last one. It's like the advice to young researchers.

Paul Middlebrooks
This is Gaute's question because I've stopped asking this on my podcast.

Gaute Einevoll
OK, so then it's—

Paul Middlebrooks
I'll include it, for sure. I like the question.

Gaute Einevoll
You are very established researchers, what is the advice to—

Andreas Tolias: My advice, especially I think this is generally true, to young researchers is spend a lot of time thinking about what is the question that you really want to address. Talk to a lot of people. Don't just do it just because you're going to get some training. Just sort of focus on the question more. It's OK to explore stuff and be thirsty for a good question. I often find people are being too practical and jumping into a project, how you're going to get training, I'm going to do this, but I think it's very important.

One thing that I didn't, in my experience, but also this is true in the history of science, it doesn't matter how motivated you are or how much hard work you put, of course, these are important. Francis Crick is a perfect example with the double helix, right? It wasn't in their project, but he knows it was a good question. He worked on it, did an amazing discovery. I do think that we maybe don't spend enough time, and this is something we don't get trained as undergraduates, or in high school. We don't think, OK, what's a good question? We're just taught facts.

Then, suddenly, when you start doing research, now you have to start thinking about questions. It's a hard thing to do. I would say, for young people, spend more time thinking about what is really a question that I want to do. This is one. Unless you need to have something that you're very excited about, you're very curious about, that's the best thing of a scientist. If there's something that you're really curious about, just follow your curiosity and what really excites you.

If you're trying to choose, you say you want to become a scientist and you're trying to say, OK, what should I work on? I'm excited about a lot of stuff or I can get excited about a lot of stuff. Spend a lot of time thinking about the question.

Paul Middlebrooks
Just linking that back to that productivity question, would you consider that productive?

Andreas Tolias: Yes, I would think of that as very productive even if it looks like you're not doing anything for-

Paul Middlebrooks
I think that's valuable.

Andreas Tolias: -a month or a year or a couple of years. I think because it's like a big boat, once it takes off and goes into the direction, it's harder to steer it. I think it's important, unless you have like something that is like, OK, I'm very curious, I don't know, about some problem in physics, or I don't know, like people are just curious about something, they can't sleep, they stay up at night thinking about it, then go for that.

If it's like you're trying to get a Ph.D. and you're interested, let's say, in neuroscience or molecular biology or genetics, there's so many stuff, there's so many labs, just put effort thinking about what is the question and educate yourself more broadly. Don't just go very narrow. You're going to have to go narrow and focus, for sure, but start broad, like what is the impact of what I'm trying to do? Where is this field going? I think it's important.

Kenneth Harris
Yes, so what I'd say is, you have to enjoy it and if you're not enjoying it, switch to doing something else that you do enjoy.

Paul Middlebrooks
Do you have to enjoy it 24/7?

Kenneth Harris
No, but on balance, it has to seem worth it and if you ever think, why am I doing this? Then think, how can I change it so I actually want to do it? If you can't change it so you want to do it, then do something else because there are other options and if you're not at the point of science, you're just supposed to enjoy it. If you're not enjoying it and you can't figure out how to enjoy it, then do something else.

Gaute Einevoll
I guess also there is some, obviously, you are sort of successful, it's sort of like a survivor bias because you actually, well, at least you made it to get a—

Paul Middlebrooks
So far, so far.

Gaute Einevoll
So far, exactly. I guess people, like young people listening to this, they think, well, I enjoy it, I enjoy being doing research, maybe get a position in academia or at a research institution doing basic research, but what are my chances for actually making it? There's always that thing also, investing many, many years and it—

Kenneth Harris
But it's not like you can't do something else. For me, I very nearly stopped at a certain point.

Paul Middlebrooks
Oh, I like this. Tell us more of the story.

Gaute Einevoll
You have a quite unusual story.

Kenneth Harris
Yes, very unusual. Well, I started doing physics, switched during my PhD, which funded me to do anything I wanted anywhere in the world. I thought about all sorts of things in—

Gaute Einevoll
That's a nice scholarship.

Kenneth Harris
I know, it was great. It was the fine print. Lots of American citizens have this, they just don't know they have it. I ended up doing neurorobotics in London, and it wasn't that much of a success, my Ph.D. really, but because—

Gaute Einevoll
Your bachelor was in mathematics.

Kenneth Harris
It was in math, yes. Because the lab I was in didn't have any continuation of funding, I had to fund myself. I got a job building a very early internet gambling website and very easily could have stayed doing that. It's only a few years ago, actually, that I finally earned more money than I was in that job. Yes, there's always other options. Pretty much everyone doing this job can code. You have other options. I could have been a stay-at-home dad. That would have been great.

Gaute Einevoll
Yes. Because I think, at least when my students ask me about this, I say, well, taking a Ph.D. and you're sort of learning how to, like scientific programming, coding and stuff, that's always going to be good. I wouldn't worry about that. It's more maybe if you decide to go on postdocs, then you are sort of, that's more of a bifurcation point.

Kenneth Harris
Do you think it cuts you off? I don't think so.

Gaute Einevoll
I don't really know.

Kenneth Harris
I've had postdocs go and work in data science; they've done great.

Gaute Einevoll
Yes, and they've even got good jobs.

Kenneth Harris
Yes, yes.

Gaute Einevoll
Excellent, yes. Maybe I'm too pessimistic.

Paul Middlebrooks
Maybe. Thank you, guys. Is there anything else?

Kenneth Harris
Thank you. I appreciate it. It's great.

Paul Middlebrooks
By the way, Cristina, first of all, you're cheating by being in this room.

Cristina Savin: Yes.

[laughter]

Paul Middlebrooks
Secondly, are you—

Cristina Savin: It was purposefully done.

Paul Middlebrooks
Oh, so you're not waiting for us, right?

Cristina Savin: No.

Paul Middlebrooks
OK, I just realized—

Gaute Einevoll
Because Tim is out hiking.

Paul Middlebrooks
Yes, I thought—

Andreas Tolias: Now you know what we said, you can say the rest.

Cristina Savin: You stole my thunder.

Paul Middlebrooks
OK, anyway, thank you, guys.

Andreas Tolias: Thank you.

Kenneth Harris
Thank you.

Gaute Einevoll
Great.

[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 advance 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]
