

Cristina Savin and Tim Vogels discuss how AI has shaped their neuroscience research

Not all neuroscientists use artificial intelligence in the same way or for the same purpose. Neuroscience researchers from different fields discuss the impact AI has had on their research and how it influences productivity in their labs.

11 October 2024 | by PAUL MIDDLEBROOKS

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

[music]

Paul Middlebrooks

This is “Brain Inspired,” powered by *The Transmitter*. Hey, it’s me again. I’m on dry land. I hope that you are somewhere. This is the second conversation I had while teamed up with [Gaute Einevoll](#) at a workshop on neuroAI in Norway called Validating models: How would success in NeuroAI look like? Gaute creates his own podcast called “[Theoretical Neuroscience](#).” Go listen to it. I linked to it in the show notes, along with a handful of other relevant links to the good people you’re about to hear. That’s at braininspired.co/podcast/196. Thank you for supporting “Brain Inspired,” and thank you to *The Transmitter*.

We are back. We’re still on a boat. I’m here with Gaute Einevoll again. We still have our, was it sea legs? Is that what they’re called?

Gaute Einevoll

[chuckles] Absolutely. I think we all thought after when we got off the boat, it’s interesting because of what we call vestibular system, balance system is a little bit out of whack. Even back on land, we still have some oscillations.

Paul Middlebrooks

That last day we had another couple of sessions, or one last session, in a conference room off the boat. I was a chair of one of the panels and moderated some things, and I was standing up there, swaying still. That was fun. In the last episode, you heard from [Andreas Tolias](#) and [Ken Harris](#). In this episode, Gaute and I had a conversation with [Cristina Savin](#) and [Tim Vogels](#). At the end of our conversation, we’ll wrap things up with [Mikkel](#) again, who helped to organize, he and [Konrad Kording](#) helped to organize this workshop.

I’ll just start by saying we’re not going to give huge introductions here, but Cristina gave a talk and researches more the theoretical side. She’s right up your alley, Gaute, for your “Theoretical Neuroscience” podcast. She uses recurrent neural networks to study how learning works in a very theoretically driven way. Do you want to say something about Tim?

Gaute Einevoll

Yes. Tim, he has been working on synaptic plasticity for many years. I know some of his work back from when he worked with [Wulfram Gerstner](#), who has educated many of the people working on synaptic plasticity in Europe. There it is, some really interesting work, I think, on how networks get into this balanced state by self-tuning and by inhibitor plasticity. Now he’s using these AI tools, or at least optimization tools, to not just actually explore the whole space of possible synaptic plasticity rules.

That was all, but his group in Austria is doing really excellent work, I think, when it comes to exploring synaptic plasticity in its many facets.

Paul Middlebrooks

I think even more so than the first one, the audio quality in this is in and out. There’s a lot more creaking and noises. This was in the night, and the sea was angry that evening, it seems.

Gaute Einevoll

It’s actually, when I talk to the people on the boat about it, say this because we pass Folda. It’s a well-known stretch of ocean where it often gets rough. It’s going to last only an hour or two, they said, which is true.

Paul Middlebrooks

Thankfully, our discussion didn’t get rough.

Gaute Einevoll

No, it didn’t. It was late in the night. That was also another reason for bumpiness maybe or--

Paul Middlebrooks

Anyway, final thoughts. This was just an excellent workshop, excellent amount of people, excellent people. It was just a lot of fun and I learned a lot.

Gaute Einevoll

Actually, in that final comment there because I think what people really liked that was really was a wide variety of people there. They had really different backgrounds. Also, one of Mikkel's worries, as I mentioned, I think, was that people were not going to be able to communicate, have some common ground to discuss, but that was not the case at all. Interestingly, he also said that some of the people he invited was actually he had heard on the "Brain Inspired" podcast that got sort of-- Maybe that's also you can take some credit for the excellent selection of researchers who was invited.

Paul Middlebrooks

Thankfully, I can take credit instead of be at fault, because it worked out.

Gaute Einevoll

That's true. Absolutely.

Paul Middlebrooks

OK. Enjoy our second discussion here.

[music]

Gaute Einevoll

We're going to ask you some general questions or many questions about the relationship with neuro and AI. We're going to start being a bit personal. It's a bit late in the evening and the boat is waving even more than when we did the other recording. Anyway, so in what sense has neuroAI changed the way you ask questions or do your science?

Paul Middlebrooks

Do you want to start, Cristina?

Cristina Savin

I can go first. I would say that we were doing conceptually what we're now doing with neuroAI many years before the term was invented. There are different flavors of neuroAI. We kept talking about it throughout the week. There's the AI to neuro in data-science terms and in theory terms. There's the other way around. We do a little bit of everything. Before there was AI, there was machine learning. We did exactly the same things with a different set of tools. From my perspective, our fundamental structure of our approach hasn't really changed, but the tools have gotten better.

Paul Middlebrooks

How about for you, in particular, you personally, on a personal level?

Cristina Savin

No, this is specifically about the research that we do.

Gaute Einevoll

Maybe you should just briefly mention what your research in the group is.

Cristina Savin

Sure, why not? My group is sort of like fundamentally interested in understanding principles of computation in the context of adaptive behavior. We're interested in sort of like normative/mechanistic interpretation. We make circuit-level models of learning, of memory, of task-dependent adaptation, attention, things like that.

Gaute Einevoll

With normative, you mean that you ask, how is this helpful for the animal? Is that what you mean by normative?

Cristina Savin

Exactly. Normative, we mean that we think that these are fundamental computation for the animals. Through evolution, development, etc., they've been optimized to do them well. Then looking at the machine-learning optimal solution of the same kind of problems should give us indication about the essence of the computation that the brain has to do. I think that it's really important for those computations to be very important for the animal. This wouldn't happen for everything, but that's the approach.

Paul Middlebrooks

Deep learning itself has not-- You're saying the machine-learning tools that existed before deep learning were sufficient for what you do?

Cristina Savin

I'm saying that the structure of the approach, how we ask this question, hasn't changed. Before deep learning, we would formalize our normative descriptions of the task in probabilistic terms, and you would use tools from Bayesian machine learning to say what the optimal solution looks like and try to make the maps to circuit function. Now we have a richer set of tools.

Gaute Einevoll

Bayesian machine learning, that's like using machine-learning techniques to find these probabilistic functions?

Cristina Savin

Yes. As I said, we have more tools and we have more powerful tools, but the way we approach the questions hasn't really fundamentally changed.

Tim Vogels

I think I would agree with that.

Paul Middlebrooks

Wait a second.

Cristina Savin

This is for what we do in my group. I'm not saying that this is a general statement.

Gaute Einevoll

What about you, Tim?

Paul Middlebrooks

Based on your talk, I was thinking that you would have a different answer.

Tim Vogels

No, I think--

Gaute Einevoll

This is for the listeners, so we should--

Paul Middlebrooks

That's what the introduction is for.

Gaute Einevoll

OK.

Paul Middlebrooks

Do you want to just give a brief roundup of what that was?

Tim Vogels

My lab is similarly interested in circuit dynamics and the interaction between network-level activity and plasticity rules. How do plastic synapses change the dynamics and how do the dynamics change the synapses? I think I spent a lot of my time as a Ph.D. student and as a postdoc tuning spiking networks at times for months. Now I don't have to do that anymore, I guess in part because I'm not the one doing the programming, but also in part because the tuning part is being taken over by machine-learning methods. How does AI change the way I approach a question? I don't think about how painful the tuning is going to be anymore.

Paul Middlebrooks

Parameter fitting.

Tim Vogels

Yes.

Paul Middlebrooks

Looking back, does that feel silly that you spent that much time, or was that valuable?

Tim Vogels

I think it was valuable. I think, actually, I had a blast doing it. It was frustrating, but it was also rewarding. I don't regret having tuned for six months.

Gaute Einevoll

Also, when I read one of the papers that you wrote with [Henning Sprekeler](#) and also in the group, [Wulfram Gerstner](#), was this really cool thing where inhibitory plasticity-

Tim Vogels

Does the tuning for us?

Gaute Einevoll

-does the tuning for us.

Tim Vogels

That came out of us.

Gaute Einevoll

Yes, but that was before-- This sort of, if you had the good AI tools then, maybe you wouldn't have thought about this.

Tim Vogels

No, I think I still would have thought about it, but we may not have hand-tuned the rule.

Gaute Einevoll

OK, yes.

Tim Vogels

Yes, I think we would still.

Gaute Einevoll

So the absence of AI didn't help you to prevent--

Tim Vogels

I don't think so. I think there's other ways. You asked the other question, how has AI negatively affected your--

Paul Middlebrooks

I was about to just go ahead and bring that up because, yes, I thought, go ahead, you're going to answer it anyway.

Tim Vogels

I think there is a push, there is some pressure to use ML tools as a scientist. If you don't, you're not considered interesting.

Paul Middlebrooks

Is that ML tools or is that specifically deep learning?

Cristina Savin

Deep learning.

Tim Vogels

I don't think it has to be deep learning. I think it has to be some flavor of large amounts of compute. If you don't-- You don't agree?

Cristina Savin

No.

Tim Vogels

I think if you don't put in your CV somewhere that you've used 600,000 hours of CPU time, you're not available.

Cristina Savin

I don't put that on my CV. That's bullshit.

[laughter]

Tim Vogels

I'm saying this facetiously, but there is truth to that.

Gaute Einevoll

Now I know why I haven't gotten the grants lately. This has not been in my CV.

Cristina Savin

I do think that the community has gotten more machine-like in the sense that it's driven by fashions, and so the particular brands of deep learning have become fashionable, and it's easy to publish certain things and very difficult to publish things with other tools. For me, the consequence of this fashion-driven research enterprise has been a reduction in the entropy of our approaches.

Tim Vogels

I agree. There are some people that are-- Sorry to interrupt. There are some people in our community that are very clearly very deep thinkers and very theoretically minded scientists that are our seniors that wouldn't get a job today because they're not using ML to further their thinking.

Gaute Einevoll

And a reduction in entropy, meaning that there's more things--

Cristina Savin

Basically, I don't think there is one way of doing research that solves all of the questions. I think that there is strength in diversity in the community for the approaches because for different kind of questions, different approaches are better or sort of like make more sense. It's also sort of longer-term. We want to preserve knowledge about lots of different ways in which to do things because they might become relevant again.

One of the things that I learned when I was an undergrad from one of my professors, I did computer science as an undergraduate, was about this sort of like really old ways of doing memory storage.

Paul Middlebrooks

They made you learn this?

Cristina Savin

Yes. We had to learn this in school. It's like, OK, like why are we learning how--

Tim Vogels

Grad school?

Cristina Savin

Grad school. No, undergrad. Sorry.

Gaute Einevoll

She studied in Romania, so she's hardcore.

Cristina Savin

Hardcore--

Paul Middlebrooks

Hardcore what?

Tim Vogels

We had to learn this in kindergarten.

Cristina Savin

It was sort of like we were learning in the computer science classes in undergraduate about the history of different operating systems and how they handled like memory things and things like that. That was like, "OK but we have better computers. We know how to do this better. Why are we learning these things?" The moment when mobile phones came on, which had like very different resource constraints, all of these old tricks that were completely irrelevant for a number of years became all of a sudden like super relevant and super important again.

I think this is the thing that we also want, as a community of scientists, we want to explore sufficiently many different things to be able to do sort of like long-term selection, this cross-breeding of ideas. It's an echo chamber. Everybody thinks exactly the same way. We have lost something really fundamental about the process of doing science.

Tim Vogels

I think there used to be a period about 10 years ago or so that there was a bit of maybe a snobbery towards people that were purely numerical, and that has flipped. Now there is a snobbery towards people that are not purely numerical. That just theory will simply not get you there.

Paul Middlebrooks

Neither of those are healthy.

Tim Vogels

No, clearly.

Cristina Savin

That's exactly the point. Entropy is the answer.

Gaute Einevoll

Exactly. Many approaches. Actually, also in physics, it used to be that people who just do too much numerics, they were like number crunchers and all that sort of thing.

Tim Vogels

Not theorists. Are you a theoretical physicist, or are you even a physicist?

Gaute Einevoll

Exactly.

Paul Middlebrooks

In your talk, you have used machine learning as a tool to explore the space of possible parameters that could tune the plasticity, right? You were alluding to that earlier as well, but so for you, it's really just changed the way you approach things, but as a tool.

Tim Vogels

Right. My students can do things that I would have never been able to do. They maneuver vast landscapes of parameters that I could only dream of. They have the means to not only just travel through them, but actually find meaningful combinations.

Paul Middlebrooks

That's the whole point, right? Because you got to explore such a large landscape of possible combinations of parameters, you found the ones that actually work. The backstory of this, of course, is that there is a, I'm just going to repeat the term, I know it's said over and over. There's a zoo of "plasticity rules" in terms of the duration between pre and post-spikes that then lead to strengthening or weakening synapses. It used to be that, the Bliss and Lomo, sorry, that's a technical sort of very specific thing, that was the rule, but since then, there have been lots of rules that have been found. What you'd use machine learning for was to explore the capacity, essentially, the possibility of the rule space.

Tim Vogels

We added another 10,000, basically.

Gaute Einevoll

It's sort of interesting. It has some analogy, analogous to sort of like when we used to do, like biophysically detailed neuron modeling, right? Then there was only a handful of models that people used, like [Mainen and Sejnowski](#). There was like a few cells that everybody used, essentially. Suddenly, you got this automated way to make these neuron models, right? You got like the Blue Brain Project produced a lot of neuron models, and also the Allen Institute. Suddenly, you go from three or four, a handful of neuron models to like a whole [crosstalk]. Has it made life more complicated?

Tim Vogels

I think for morphologically plausible modeling, certainly. I think choosing what model you want to use is not gotten easier.

Gaute Einevoll

Also for you, too, in this synaptic plasticity rules, right? You showed us--

Tim Vogels

No, I think for us, it's gotten a little easier. We have different questions, but they're certainly more satisfying. Because when you found a single rule that worked, it was almost certain that you were wrong. The experiments were incredibly arduous to do. Now we're still wrong, but--

Cristina Savin

You have a space of hypothesis.

Tim Vogels

Yes.

Cristina Savin

I think that's actually one of the good ways of using these more powerful tools that we've inherited from the deep-learning revolution, is to explore options that you wouldn't have thought of otherwise.

Paul Middlebrooks

That are not so expensive to explore as well, right? Time-wise.

Cristina Savin

Yes. It's sort of like practical for a Ph.D. student to do that and get the Ph.D. in a reasonable amount of time.

Paul Middlebrooks

What I'm hearing from you both, and I think that everyone in neuroscience would agree, is that the new brand of deep learning, machine learning on steroids and AI, is great as tools. However, it seems like you both agree that there's something lost in terms of the knowledge of the other spaces of possible solutions to things and approaches.

Tim Vogels

I don't know if it's lost for us, because we already have a job.

Cristina Savin

Yes, but we're old enough to know these things. I do worry about the incoming Ph.D. students who have been trained, know how to train convolutional neural net or whatever, they know how to run some deep learning, all of the diff-- They come to their Ph.D. and expect that to be the essence of what--

Paul Middlebrooks

You do worry about that?

Cristina Savin

I do worry about that. It's like we're talking about sort of negative impact, I think.

Tim Vogels

I don't know if we are just old, right? We just sound old.

Cristina Savin

We're old and grumpy.

Paul Middlebrooks

I was just saying that you guys are the "young investigators," quote, unquote.

Cristina Savin

Let me finish the thought. I think with really powerful tools, they're as good as used, sort of like your uses of them are. It has the potential of making things substantially better, but in the wrong hands, it could also make things much worse. My worry is about the ability of our students to critically think they use these tools in a reasonable way.

The focus now, if you're thinking about a junior person trying to get into this field, it's not how you use these tools. You can go to the internet and learn how to do that in a week. That's not the educational component. The educational component is how to think hard about the problems and the use of these tools in a meaningful way. That's hard. Because the objects are increasingly complicated, so reasoning about them is hard.

Gaute Einevoll

I remember I read this Linus Pauling. He was the guy who found the structure of protein. He said that when he did this X-ray, he had to do so much manually. He had so much time to think about things. He was worried that the new tools would get this X-ray spectra out automatically.

Tim Vogels

Wasn't it Socrates who said that he worried that writing would make us dumb?

Paul Middlebrooks

Yes.

Gaute Einevoll

I was just thinking from this paper that you mentioned, where you tune this inhibitory plasticity to get to the balanced state, right? In this process of tuning by hand, you thought about it much more.

Tim Vogels

Yes, but for the ISP paper, we didn't tune. That was the beautiful thing.

Gaute Einevoll

The ISP, what is it?

Tim Vogels

The inhibitory synaptic plasticity paper. We tuned for the balanced paper the two years before that.

Gaute Einevoll

OK. Then you tuned that. Then did you learn much about it? Did you get the dynamics of the network under your skin?

Tim Vogels

I did. Yes. I still do, in fact.

Gaute Einevoll

You do that also with the machine learning techniques?

Tim Vogels

I don't know.

Gaute Einevoll

Or your students?

Tim Vogels

I think my students do.

Gaute Einevoll

Yes. They sort of have like an intimate--

Paul Middlebrooks

Is that for the students?

Tim Vogels

Yes. I think so.

Cristina Savin

I keep thinking about this sort of like this, was it an Arthur Clarke quote? About "any sufficiently advanced technology is indistinguishable from magic" kind of comment. Sometimes I feel that these transformers and some of these like really big LLMs and things like that feel like magic to me. I wonder how many of the neuroscience users of that technology treat it a little bit like magic. That you are so detached from how it does things, what are its limit cases, what are the things that it could possibly go wrong, then you take for a given that whatever that thing spits out is the truth. I think that's another thing that we're losing with all of this complexity. That our ability to sanity check the process becomes very, very hard.

Gaute Einevoll

I guess you both say that it has improved the methods for doing what you already were interested in some sense, right? That's like an improvement of tools. I guess these tools will lead to new discovery, hopefully, new discoveries in the future. As I said, you already have some examples of where it has already changed the way you think about the brain and cognition.

Tim Vogels

We've been talking about the tools, and the other side of that is this continually growing people are very excited about using deep-learning models as a proxy to model brain areas and they understand brain areas and/or cognitive functions better. That's a different facet of neuroAI. That's where I was thinking like, well, has it changed the way that you thought about how brains function? Not from the tool front, but from the using them as models of brain function.

Cristina Savin

When I said normative, we are actually doing that to some extent. There is one example where I was like-- OK. I started my faculty job saying that we're going to do interesting mathematically tractable things and this is going to not be a deep learning lab. It was sort of like forcefully. Then students came into the lab and they really, really, really wanted to do it, then they had to--

Tim Vogels

To do deep learning?

Cristina Savin

To do deep-learning methodology. We tried to think really hard about how to do that in a way that's not stupid and trivial. I couldn't stand by my original statement, but I was pleasantly surprised in a couple of occasions.

Gaute Einevoll

They were doing deep learning after you had left for home or something.

Paul Middlebrooks

That must be all students these days. Doesn't every student want to come in and do deep learning these days?

Cristina Savin

This is the problem, partly.

Tim Vogels

Not to my lab, actually. No. I do spiking networks.

Paul Middlebrooks

I know, but I just, I don't know. Everyone wants to apply it to everything, so I just couldn't imagine.

Cristina Savin

It's going to happen sooner or later. There is one example where there's a problem that we've been thinking for a very, very long time. Using tools from deep learning to ask those questions gave us qualitatively different solutions. That's sort of like trying to think about how the brain infers from noisy observation, things that are important for, like latent states in the world that are important to drive behavior.

We and others have a cottage core. There's an entire niche in computational neuroscience, engineering, probabilistic representation. How would neurons go about encoding beliefs about the state of the world?

Paul Middlebrooks

You're on that paper with [Ralf Haefner](#).

Cristina Savin

Yes, for my things. It took us I don't know how many years to write. Those are something very constructive. Me, as a theorist, I go about to say how would I go? What do I think are the things that are important? I know math. These, I think, are the things that are important according to math, and this is how maybe I could map them into neurons and neural activity. In a recent project where we're trying to understand the behavior of animals making inferences about changes in the context.

Gaute Einevoll

You have to take your hands off this. It's going to keep creaking.

Cristina Savin

Sorry about that.

Paul Middlebrooks

That's all the creaking is?

Gaute Einevoll

Yes.

Cristina Savin

That's me. We went a different route because the task was sufficiently complicated that it wasn't obvious how to apply the traditional approach. We try to train some deep reinforcement learning agents to do the task. At the level of behavior, so when we're analyzing this recurrent neural network the same way as we do the animals, their behavior looks close to optimal probabilistic. If you open the box and look exactly what they're doing, it's nothing like any of the solutions that we have imagined empirically. This taught us a big lesson that this perfect mathematical elegance in the map to the neural activity was probably a futile endeavor altogether.

There are kind of these not-so-obvious ways in which to achieve functionally the same thing. We won't have ever come up with those kind of solutions on our own without the use of these technologies. I think those examples are still rare.

Tim Vogels

Are you thinking, how does the neuron implement backprop? Some people think that, right?

Cristina Savin

So many people talk like that. Unfortunately.

Tim Vogels

I don't care at all.

Paul Middlebrooks

You don't care at all.

Gaute Einevoll

No, in terms of finding parameters.

Tim Vogels

No, I just I'm like, meh. I don't know. I'm exhausted.

Cristina Savin

If you care about learning at the circuit level, I think you do need to care about how does task-relevant information shape synaptic plasticity in a way to drive behavior towards good states. Backprop to time, in this case, are mathematical tools to formalize that problem in precise ways.

Gaute Einevoll

In your project, you were just interested in getting a good solution. I say in war, love and optimization, everything is OK.

Cristina Savin

Yes.

Tim Vogels

Not everything is global. There's a lot of local learning that is not backprop.

Gaute Einevoll

Absolutely. We're talking about biology right now.

Tim Vogels

[Carson Stringer](#) just published a paper in the arXiv that says at least 50 percent of what they see in terms of learning effects can be explained with local changes and local rules.

Gaute Einevoll

That's a question of what's happening in real race.

Tim Vogels

Yes, exactly. How is backprop implemented is not that interesting a question for me.

Cristina Savin

Tony Slater would have totally liked this answer. I still believe that there's a sufficiently large amount of goal-directed learning happening that we need to figure that out.

Paul Middlebrooks

Tony is going to be on my podcast. He's been on it before, but I don't know if it'll come out before this chat or after. It depends on how long.

Cristina Savin

They keep inviting us to the same neuroAI meetings, and they keep putting his talk before mine. He gives an entire amazing talk about how the brain does very little plasticity, and it's all something evolutionary inductive. Then I go say, "I'm going to give you a talk about plasticity."

Paul Middlebrooks

Are you going to that this year? It's at the end of this month. I was asked to go, and I was like, I don't know.

Cristina Savin

No, we've already done a UCL one this year.

Gaute Einevoll

What meeting is this?

Paul Middlebrooks

It's the Cold Spring Harbor--

Tim Vogels

Spiking Networks. Ironically called Spiking Networks.

Paul Middlebrooks

No, this is a neuroAI.

Cristina Savin

This is a neuroAI.

Paul Middlebrooks

Something subtitled.

Tim Vogels

Yes, subtitled Spiking Networks.

Paul Middlebrooks

To make me stop, yes, there you go, yes, but anyway.

Tim Vogels

Ironically, no Spiking Networks.

Paul Middlebrooks

It sounds like you don't give a damn about it. No, I'm just kidding. It doesn't sound like this deep-learning revolution has changed the way that you think about brain function?

Tim Vogels

No, I don't know that it has. I'm still interested in the same things. I'm interested in local changes in plasticity rules and neuromodulated changes to activity that will then produce local changes in plasticity.

Cristina Savin

I wouldn't say the same. For us, sort of like what's changed is the scale of our ambitions. With this set of tools, we attempt to understand adaptive behavior of much higher complexity than we would have without it.

Paul Middlebrooks

That doesn't speak to like the way that you sort of internally think about how brains function.

Cristina Savin

No. This is why I said that their methodological approach only has gotten richer, but it's like structurally hasn't changed.

Gaute Einevoll

You are sort of in like, what you're doing like--

Paul Middlebrooks

He's pointing at Tim, folks.

Gaute Einevoll

What?

Paul Middlebrooks

He's saying you.

Gaute Einevoll

Yes, Tim. Tim, you are sort of working off with the spiking networks of like integrated fire type and it's not like multi-compartmental modeling, but it's still quite biophysical, right? You think that--

Tim Vogels

It's AI.

Gaute Einevoll

Yes, it's sort of like physical, right? There are real things that [crosstalk]

Tim Vogels

Yes.

Gaute Einevoll

Do you think that this focus on AI will sort of actually suppress the activity?

Tim Vogels

No, I don't think so. Look, that's part of the appeal of I think the questions that we're asking in my lab is that we're not competing with the big companies. We're not competing with DeepMind and etc. They don't give a shit about spiking networks because so far they haven't been proven to be computationally viable. One of the questions in my talk was, can't you find a function or a task that is actually computationally interesting? All the neuro tasks you have memorizing something are totally boring for someone who's doing AI.

Gaute Einevoll

You still have this potential for providing, well, requiring much less energy. Isn't that the whole idea?

Tim Vogels

That's an argument that people make for spiking. I don't know that's making spiking networks the water carrier for big AI. That's not my interest. I want to understand how the brain works.

Gaute Einevoll

Cool. How do you see, well, these are exciting times in AI, how do you see the relationship between AI and neuroscience, can you speculate that in the years to come?

Cristina Savin

Long-term future.

Paul Middlebrooks

By the way, sorry, Cristina, I think you almost got the quote verbatim. It's Arthur C. Clarke. It's "Any sufficiently advanced technology is indistinguishable from magic." Sounds like you got it verbatim.

Gaute Einevoll

Yes.

Cristina Savin

Yes. I think about that a lot.

Paul Middlebrooks

Sorry to interrupt.

Cristina Savin

Thank you. The future of the interaction between AI and neuroscience, I think that's actually where at the moment where this is like maximally unclear. One reason for that is--

Tim Vogels

Maximum entropy. Isn't that what you like to say?

Cristina Savin

We do write papers about maximum entropy models, but this is an accidental happening. On one side, AI is in this phase of young enthusiasm and exuberance. You blink and the entire set of those fanciest architectures, fanciest trick has already changed. You can't keep up with the literature, the things change so fast. It's hard to say what like AI would be about in six months' time, not to mention 10 years. That's a big source of uncertainty because we don't know why, like where that's going. Presumably as this thing matures, you're going to see the same things that in maturing other fields. It's going to be less about changing our mind about how we want to approach this every few months. More about converging to a set of good solutions and trying to build a foundational understanding of why they work. We're not there yet, but—

Paul Middlebrooks

So you think that this recent exuberance will die down in the next couple of years?

Tim Vogels

I think we're just going to switch course. I think general artificial intelligence is not far off.

Paul Middlebrooks

Oh, geez, hot take.

Cristina Savin

I'm going to not say that.

Tim Vogels

At least in parts, we're going to have a more powerful--

Gaute Einevoll
Partial general?

Tim Vogels
Yes, so what do you find is general, right? Is it creativity and--

Gaute Einevoll
Creative.

Tim Vogels
I don't know that we're going to be using these units for a while. Because creativity is--

Paul Middlebrooks
Not part of your definition of general.

Tim Vogels
Yes, I would say, like transferable skill is fine, creativity and--

Cristina Savin
I'm not even sure that that's true. I think LLMs do something good, but it's not general intelligence.

Tim Vogels
But we're closer to it than I think we can project. We can't project two years down the road. What the relationship is between AI and neuroscience is entirely dependent on how, for the next two years, it's going to be entirely dependent on how quickly tools from AI are going to become ever more powerful to understand relationships in neural recordings that we can't even fathom yet.

Cristina Savin
Yes, we're going to become users of this technology, that's for sure, but--

Paul Middlebrooks
We already are.

Tim Vogels
Yes.

Cristina Savin
To an even larger scale. It's going to be the bread and butter that you need to know how to do these things.

Tim Vogels
NLPs are spiking up.

Cristina Savin
I do wonder what purpose does basic research in neuroscience serve if you have a functional model of general artificial intelligence? Because a lot of the original motivations of why I got into this field was, oh, this is the most intelligent system that we know. If you want to understand principles of intelligence, looking at the brain is a good idea. If we have an artificial model of that that we're satisfied with to some degree, I think it's not clear exactly how to assess that.

Assuming that would be the case, then I think that the computational neuroscience community would have to have some really serious soul-searching about what are our questions, what purpose do we serve now. What Ken and Andreas were saying earlier about, sort of like, maybe this becomes more about circuit level and molecular going more low level. Systems neuroscience would not have that much of a purpose in basic research terms, but you're going to have to go down to get to the clinical applications part of the process. Might be one way this would play out. That would make me very miserable, and I might have to change fields. Who knows?

Gaute Einevoll
One thing that we have seen is, for example, talking about Andreas, that he's already making these foundational models, which are extremely good, and based on deep networks, extremely good at predicting, in a neural activity, neural responses based on visual input, but it's difficult to interpret. This thing of getting increasing predictability and losing, the older models were sort of not as good at predicting things, but you can sort of think about things like receptive fields and Bohr functions. Are you more comfortable, maybe? Are you comfortable about sort of having less interpretable models, but more--

Cristina Savin
I think interpretability is a concern. I think what the modern deep-learning tools are providing us, which is going back to the general intelligence

discussion, are extremely powerful statistical descriptors of large quantities of data. There's a very big difference between very good statistical descriptions of a dataset and a process model that describes the causal relationships that generated that data. Those could be completely different solution classes. I think ultimately we want to understand how things work, and we need process models. It can't be the end of the story because it just summarizes exceptionally well a large amount of data.

Tim Vogels

It also won't work to replicate what computer science has done, which is to create systems that we then don't really understand. Our goal is to understand the brain, so it won't help to simply replicate the brain in silico.

Cristina Savin

There's another famous quote, "I don't understand what I can't create." I think the fact that I can create something doesn't necessarily mean that I understand it.

Paul Middlebrooks

Yes, you said that, right? Feynman is wrong.

Cristina Savin

The converse is not true.

Gaute Einevoll

Yes, it's like a new thing. Also, I think if you have like this, at least if you have like a biophysical model, a network model, which of course is just as complicated as a deep network in some sense, at least if you're able to make this model predict experiments, then it's like a white box, you can start playing with it. In that sense, it's like a beautiful research animal in some sense.

Cristina Savin

It is a beautiful research animal; I agree with that.

Tim Vogels

A one-to-one map of the world is going to be as--

Paul Middlebrooks

Is the world.

Cristina Savin

The best model of a cat is the cat, preferably the same cat.

Tim Vogels

I don't know, is it?

Cristina Savin

Yes.

Paul Middlebrooks

Is that Schrödinger? No?

[laughter]

Cristina Savin

I don't remember who that is, but this is the quote from the very first summer school ever in computational neuroscience that I attended.

Tim Vogels

Gaute.

Cristina Savin

I keep collecting these.

Paul Middlebrooks

Someone else. I don't have any good quotes.

Tim Vogels

I like the other one better with a cat, which goes, "Science is like looking for a black cat in a dark room, and it's not even a cat."

[laughter]

Tim Vogels

I think that's what we're asking. What's the relationship between AI and neuroscience? It's not even a cat.

Cristina Savin

Going back to the original question, what's the future of the relationship between AI and neuroscience? I don't know, but I'm excited to find out. I think this is going to be fun. Whatever crazy [crosstalk]

Gaute Einevoll

I agree with you.

Paul Middlebrooks

Do you have an AGI take?

Gaute Einevoll

No.

Paul Middlebrooks

Ken gave his. Cristina gave hers.

Gaute Einevoll

No, I don't have anything-

Paul Middlebrooks

It's OK.

Gaute Einevoll

-worthwhile to say about this.

Tim Vogels

I think it's telling that we can—all three—First of all, I think that we have a misconception of what intelligence is, but we all have a different definition. We have to operationalize it, etc., etc. What I think of AGI is way far away. We can all three disagree, and none of us knows anything. That is exciting.

Gaute Einevoll

I should mention that I used to do condensed matter physics back in days before I got-- That was at postdoc. I took a Ph.D. and was even a postdoc doing this. I switched at a late stage. This was at a time when I would say we solved the Schrodinger equation and we-- There was no mystery at the end of the rainbow. Then you come to neuroscience, and we understand so little. It's fantastic.

?Tim Vogels

There's no rainbow.

[laughter]

?Paul Middlebrooks

There's no rainbow, yes. We're trying to invent the rainbow.

?Gaute Einevoll

I think it's fantastic. At the end of there, somewhere far out there, there's consciousness. Right?

Paul Middlebrooks

Yes, sure. There's no rainbow, but there are leprechauns everywhere.

Tim Vogels

And the stars.

Gaute Einevoll

That's fun though. I agree with Cristina. It's fun that we are-- it's a big privilege to really be at the frontier of this big unknown.

Paul Middlebrooks

We already talked about what they believe neuroscience can learn from AI. That's really in the form of tools. It's just mostly a tool.

Cristina Savin

I wish there was more the other way around. We have-

Paul Middlebrooks

Do you mean that what AI could learn from neuroscience?

Cristina Savin

Yes. Ways in which knowledge from neuroscience informs architectural choices, for instance, in deep learning or other algorithmic course. I think it's something like there is a transfer the other way around, but it's subtle. and it's not one thing that made a humongous impact. It's like in subtle ways it affects a lot of things that are happening in deep learning.

Paul Middlebrooks

Like what? Can you--

Cristina Savin

Like neurons and attention being the immediate obvious things, but also the way they approach interpretability of their trained recurrent neural networks is essentially by treating it like a brain in doing neuroscience experiments. They do ablation experiments.

Paul Middlebrooks

It helps.

Cristina Savin

They do in-silico mapping of receptive fields and things like that. I remember in earlier days, five years back when deep learning first started to be really, really successful. I was looking at what they're doing with the network trying to understand their properties like, "Haha, you're using experimentalist tools to try to understand something really complicated and failing in exactly the same way as the first of the neuroscience community."

Paul Middlebrooks

That hurts on the inside.

[laughter]

Cristina Savin

This is the point. It's not that that there's no flow of information. It's just that it's hard to pinpoint this is one thing that completely made the difference. In subtle ways, we're influencing how the process works in a lot of ways.

Paul Middlebrooks

It's not subtle how AI is influencing neuroscience. Someone had to argue that the brain is not a transformer today. They actually spend a whole talk arguing that we're not transformers.

Cristina Savin

If he stated it. I don't think he was trying to-- I don't think that was a position difficult to defend.

Paul Middlebrooks

No, but the fact that it was defended for that period of time.

Gaute Einevoll

There are also neuroscientists to make models for how transformers can be implemented both at the neuron circuit level and at neural area level.

Cristina Savin

We do have models of contextual modulation of visual processing, for instance, that are circuit-level models. They're not transformers in the details, but they are transformers in the spirit. In the sense that I have top-down information that decides what kind of things from my input stream is task-relevant, what kind of bits are not. I preferentially transfer the bits that I care about. The dynamics of that is what we have circuit models for and related to actual neural data. Again, it's what counts as similar enough. Is it similar in spirit inspiration or is similar in spirit but not in the details, something to be ignored. I don't know.

Gaute Einevoll

Changing the topic a little bit. You, Cristina, just got tenure. Congratulations.

Paul Middlebrooks

Oh, I didn't know that.

Gaute Einevoll

Very good. Which is a milestone, obviously, if you want to stay in academia. You have a permanent job, but only for a few years.

Tim Vogels

Four years now.

Gaute Einevoll

Exactly. You're rather newbies in terms of tenure. That also means that-- because this question, what does it mean to be productive in science? From an operational point of view, if you want to stay in academia, being able to qualify, getting tenure is a good thing. You have to show up with papers and grants and stuff to get to that. At least you have just been there for a few weeks, and you have now time to-- you can do different things with your career. You can be involved in many projects and this spread your thing. You can give [crosstalk]

Cristina Savin

We were always doing that.

Gaute Einevoll

Yes, exactly.

Tim Vogels

I think actually, the time before tenure for me was probably much more diverse in what I was doing.

Paul Middlebrooks

Is that because you were trying to find the thing?

Tim Vogels

No, also I was doing things like the Imbizo or the summer school in South Africa that I started to co-direct, or World Wide Neuro. These things were not necessarily directly career-relevant, but they were fun. They served as an outlet for what is otherwise a relatively-- computational neuroscience doesn't have a direct impact on many things.

[laughter]

Tim Vogels

Those additional hobbies serve that purpose.

Cristina Savin

If you're a tenure-track American professor, you don't have time to have hobbies. My entire career has been about doing too many things at the same time. That's a bad strategy, they tell you not to do it, but it was just happened to me. It's what works with how my brain works. I have lots of spread-around interests. I don't care about one thing, but about a range of things. We had lots of collaborators, so that mushroomed even more projects. It just happened.

Gaute Einevoll

Did you think that you do your best work in that way, or do you think now that you can focus on fewer things?

Tim Vogels

I find myself being swallowed up by administrative chores since I have tenure. It is now said that the time of the juniors has to be protected, and so here is the administrative load increased.

Gaute Einevoll

Oh.

Tim Vogels

What I used to spend on the Imbizo, which I have just retired from, or World Wide Neuro, I now spend on hiring committees, and grant committees, and all kinds of nonsensical, less-

Cristina Savin

I would say that--

Paul Middlebrooks

You've always--

Cristina Savin

[crosstalk]

Tim Vogels

-immediately relevant things.

Paul Middlebrooks

Wait, before-- were you going to ask about productivity?

Gaute Einevoll

I think I'm asking about it now.

Cristina Savin

Yes, let me finish the answer to the question. Basically, because I was putting up my documents for tenure not so long ago, and I had to think retroactively, it's like what is it exactly that we've achieved, where is this going, things like that, there was a lot of soul-searching involved. I think I came up with the conclusion that doing fewer things better is something that I would like to try in the incoming years. They're practical things, so you can trim down as much as you want, potentially, but--

Paul Middlebrooks

I just wanted to ask, immediately, and then we can come back to this. What advice would you give to people going in now given-- reflecting on your own, are you going to say, "Don't do what I did?"

Tim Vogels

No.

Cristina Savin

I do tell people that on a regular basis, but this is not what I'm going to tell today.

[laughter]

Cristina Savin

I think taking the time to find a good question before you jump into doing frantically things is something that I try to encourage my starting Ph.D. students to do.

Paul Middlebrooks

How do you know what a good question is when you don't know anything?

Cristina Savin

This is the second immediate advice that I give to those people is read literature broadly. Maybe this is like old gramps talking, but I do feel that in graduate school I spent a good fraction of my time reading papers. There used to be this 10 advice plus-- what was it called?

Paul Middlebrooks

10 rules [crosstalk]

Cristina Savin

10 rules for blah blah blah. There was 10 rules for becoming a great writer. One of those was read 10 papers a day. Nobody can read 10 papers a day. What are you talking about? These papers are getting so complicated that it takes a week to read. I think a milder version of that advice is very good. That you need to know what the field is about. People tend to read very narrowly. Exactly in the niche of what the project is about, but missing a lot of really important connections because they just don't have [crosstalk]

Gaute Einevoll

Of course, now you have two excellent podcasts in the field.

Paul Middlebrooks

Yours, and who's?

Gaute Einevoll

"Brain Inspired" and mine, sort of.

Tim Vogels

And ChatGPT.

Paul Middlebrooks

That's not a podcast.

Cristina Savin

That's a thing. Tim's new app to summarize thoughts.

Gaute Einevoll

What is the new app?

Tim Vogels

[TopoTopic](#).

Gaute Einevoll

TopoTopic, yes. Tell us about it.

Paul Middlebrooks

I'll put it in the show notes.

Cristina Savin

No, I just want to say that the challenge there is because the field has mushroomed a lot more productions, so basically, there are a lot more papers to read in general. That's maybe one of the places where deep learning might help because if you have summarization tools, you can get at least a superficial breadth.

Tim Vogels

Tweeprints, actually, I think is hugely interesting way of--

Gaute Einevoll

Tweeprint?

Paul Middlebrooks

Tweeprint is what [crosstalk]

Tim Vogels

Tweeprint, like the preprint on Twitter.

Gaute Einevoll

Oh.

Tim Vogels

Now on ...

Cristina Savin

X.

Paul Middlebrooks

It offers a summary.

Tim Vogels

Yes, quick summaries on Twitter. I really enjoy those.

Gaute Einevoll

I should maybe start doing those.

Paul Middlebrooks

What would be your advice to--

Tim Vogels

I think drink from the firehose. I don't know. I really don't know. I find the term work-life balance really problematic--especially for young PIs, but also for grad students and postdocs. I think the only work in your life is that you have to sleep enough. That's the only work time. Everything else is life and probably--

Paul Middlebrooks

Wait a second. I'm trying to understand this. The work is the--

Tim Vogels

Sleeping. I find sleeping is my worst chore in the day. I set an alarm clock in the evening so that I go to bed at 12:00.

Paul Middlebrooks

Otherwise, you're going to keep working, maybe?

Tim Vogels

No, I just keep staying up. That's really the only thing that I violently dislike in my day, going to bed.

Gaute Einevoll

Have you ever tried putting the cocaine away?

[laughter]

Tim Vogels

No, but I think what would be my advice? Do what you're interested in. Run as hard as you can. Don't take prisoners. Drink from the firehose. I don't know.

Paul Middlebrooks

Do you think work-life balance has gotten out of control?

Tim Vogels

I think I'm going to get slaughtered if I say that seriously.

Paul Middlebrooks

Things can--

Tim Vogels

I think what is considered work has been a little bit corrupted in a way. You can't be at the same time a student and a worker, in my opinion. You have to decide whether you want to be a graduate student and take on what you get as a privilege of being taught something, or you decide that everything you do after you reach your lab is work. Then it's a 9:00 to 5:00 job. You can't have the cake and eat it, I think. If you decide that you're actually a graduate student or a scholar of some capacity to whom society gives relatively large amounts of money for very little productivity compared to a baker, or a builder, or various other jobs [crosstalk]

Paul Middlebrooks

You're talking of graduate students?

Tim Vogels

Postdocs, and PIs as well.

Paul Middlebrooks

Who get a relatively large amount of money?

Tim Vogels

I would say so.

Paul Middlebrooks

OK, we're through. Yes, all right.

Tim Vogels

I certainly got a lot less money than what grad students are getting today.

Gaute Einevoll

As a grad student?

Tim Vogels

As a grad student, yes.

Gaute Einevoll

When you were in the U.S.? You're now working in Austria. I think it's partially a European-U.S. thing, right?

Tim Vogels

Yes. I don't know. I think it's a life choice you make at the end of the day, and you are going into a field that is, a, very competitive, but, b, very privileged in that we're sitting on a ship talking about science.

Paul Middlebrooks

It's the first time for me that I've been on a boat at a conference.

Tim Vogels

Yes, me too. Same. A friend of mine, [Guillaume Lajoie](#), always says, "I fucking love science." He fucking loved science always, even when he was being paid very moderately. To keep that in mind, to call that back into your own memory, into your own every day, that you're not, in fact, doing something because you're being forced to but that you're doing it by choice, is an incredible privilege.

Paul Middlebrooks

Cristina, you look like you're chomping at the bit there.

Cristina Savin

No, I think I agree with this. I think that being in graduate school is very intense and very hard. If your heart is not really in it, and you're doing it as a job, then probably there are better ways of getting the same amount of money with better benefits and better life. I think sometimes people get into graduate school either as inertia. They were doing well-

Paul Middlebrooks

Oh, that's maybe 70 percent of people, I would guess, or used to be.

Cristina Savin

-or because of the social pressures that they want the title, but they don't really enjoy the process.

Paul Middlebrooks

Those are overlapping populations.

Cristina Savin

There might be. I feel sad about those people.

Gaute Einevoll

You are clearly, well, both extremely motivated and have made it also in the sense you've got excellent academic jobs. There are also people who they're not that full health. They have some health limitations or maybe some family obligations. Are you saying that--

Paul Middlebrooks

Maybe they're just lazy.

Gaute Einevoll

Yes, lazy.

[laughter]

Paul Middlebrooks

[points at Gaute Einevoll]

Gaute Einevoll

Really?

Tim Vogels

That's impossible to do.

Paul Middlebrooks

That's an affliction.

Gaute Einevoll

There are, should I say, if you cannot put in--

Cristina Savin

I'm not saying you need to work 12 hours a day. I'm just saying that the hours a day that you work, you need to be a 100 percent in it. Those need to count.

Tim Vogels
Exactly.

Paul Middlebrooks
Whereas what's a job that that would not apply to?

Tim Vogels
Most jobs for me are work from 8:00 to whatever, 9:00 to 5:00. If you work in a shop or maybe a--

Cristina Savin
I'm sure the--

Paul Middlebrooks
I don't know if this is airable.

[laughter]

Cristina Savin
This is not airable, but basically there are plenty of things that you can do at 80 percent and get away with it. I don't think you can do science at 80 percent and get away with it as a career. You might get the Ph.D. that way, but you're not going to be very successful.

Tim Vogels
To be fair, at the postdoc level, I think there is a lot of people who are putting in 100 percent and 120 percent, and they're not getting jobs. They're not getting invited, and their interviews go poorly. The bottleneck is getting a PI position, and that's incredibly tough.

Paul Middlebrooks
Trying to get into a firehose that won't turn on.

Tim Vogels
Yes, no, or that is passing them.

Gaute Einevoll
I say to people who ask me that taking a Ph.D. is where you learn how to program code or whatever. That's a safe investment regardless. Going on a postdoc, maybe that's a little bit-- If you're aiming for academia, that's a bottleneck. There are many more postdocs than PI jobs, permanent jobs.

Tim Vogels
I was ignorant to the cliff that I was standing on when I started my postdocs. Until I had kids, I actually had no risk-management plan. I was just--

Gaute Einevoll
You didn't really need to.

Tim Vogels
No, I didn't need to.

Gaute Einevoll
You were young and immortal. [crosstalk]

Tim Vogels
Exactly.

Paul Middlebrooks
Go ahead. Sorry.

Cristina Savin
I did my Ph.D. in Germany. The German system has a lot more insight into this because it's so difficult to get a permanent academic position in Germany. People start the Ph.D. largely with the expectation they were going to work into industry. The transferable skills and what you're learning being useful for lots of different things is part of a memo. It's not an afterthought or plan B. It's the default. If anything else happens, that's good, but it's not really expected.

In the U.S., we're still selling the academic path as the default, and everything else is plan B, although the numbers are really not-

Tim Vogels

Substantiating that claim.

Cristina Savin

- in any way reflecting that.

Paul Middlebrooks

I'm slightly afraid to ask you a productivity question, given the last one.

Gaute Einevoll

The slow productivity thing.

Paul Middlebrooks

Yes.

?**Tim Vogels**

-that I talked about? Yes, I mentioned that as I read this book about slow productivity.

Paul Middlebrooks

We have just a few more minutes. If you guys are good for just a couple more minutes, and then we'll [unintelligible]

?**Tim Vogels**

I think, yes. Just as I said, I read this book on slow productivity. It's like this. The basic idea that it's easy to measure productivity if you're a farmer or producing industrial things, or-- the knowledge work is not so easy to measure, right?

?**Paul Middlebrooks**

[unintelligible]

?**Tim Vogels**

Then you get these proxies that all have many papers or that they look visible and they work a lot. Then you think about it when you read, for example, about the lifestyle of Darwin was, I would say, productive. He had a quite productive scientific life. He didn't work that many hours. He had in the morning and then he did some, but it was like this-- He was focusing on a few things and then doing high-quality work. I guess you feel pressure to getting grants and have students and this. Success is often measured now in how many students you have, and how many grants you get in, and how many papers you do.

Cristina Savin

Yes, I need to get enough money for my students to have jobs to be able to graduate. There are external pressures. It's not like you can take your time. I wonder like these historical examples whether it's fair to make those comparisons. I do think it's the better way of doing science, but the sociology doesn't work. The reason for that is a lot of the very successful scientists in that period were independently rich. They did this as a hobby for fun. They did it whenever they wanted. They had the time-

Tim Vogels

They had no competition.

Cristina Savin

-and the leisure to do that. It wasn't like, "I need to desperately get some stuff done." We're in much more externally driven-

Paul Middlebrooks

How do you think about productivity personally, then?

Cristina Savin

I think Darwin might be an extreme version, but I do think you actually make better science that way.

Paul Middlebrooks

Screw it. We're not in that system anymore, so let's make worse science.

Cristina Savin

Yes, but the sociology doesn't work. Until as a culture, as a field, we decide that we are going to change the incentive structure in a way to make that a feasible MO, we're going to have to do what the external pressures force you to do. I can't stop doing certain things. It's just not going to work.

Tim Vogels

My postdoc adviser, [Wulfram Gerstner](#), when I joined his lab, told me three rules for his lab. There was four, actually, but three big ones. Show up in

the lab once a day. One paper per year with your name and my name on it. Doesn't matter what positions. If you want to go to a conference, you have to present your own work. On top of it, we couldn't speak anything but English in front of open doors. That was rule number four. That was it. I try to propagate this, and I think that's for postdocs a great rule. One paper per year is doable. Not necessarily as a first author, because sometimes they support grad students. If you're a postdoc for four years and you've got four papers out of it, that's really productive.

Paul Middlebrooks

Yes, but if that's the rule, then it says nothing about the quality of the work.

Tim Vogels

No, but it does. One paper per year is not a lot if it doesn't have to be a first author.

Paul Middlebrooks

It takes a lot just to write the paper and shape the paper and get the paper out the door. That in itself is a lot of work. I'm trying to think. If you're working to get a paper out, that's different than working to answer a question. Presumably you would want to be doing both.

Tim Vogels

Yes, no, but I think you front load with work and then you end up with papers. Maybe I'm speaking because of conformational bias because I was in Wolfram's lab for four years and I have four papers.

Paul Middlebrooks

It worked out. It worked out.

Tim Vogels

This plus minus one seems to work OK also in my lab. The postdocs seem to-- They're not slacking off. They're working, and at the end of the day they're ending with about a paper per year.

Gaute Einevoll

You were doing monkey physiology, which is the hardest thing.

Cristina Savin

Yes, monkey physiology doesn't work on that type of stuff.

Tim Vogels

This was a theory lab, right? This was a theory lab. Experimental labs are probably structured differently.

Paul Middlebrooks

Easy. What you do is easy.

Tim Vogels

Exactly. It's not real work.

Paul Middlebrooks

I have to clean monkey cages. It's hard work. Scrape that. Never mind.

[laughter]

Paul Middlebrooks

All right. Actually, I have one more just fun question.

Gaute Einevoll

Yes, please go ahead.

Paul Middlebrooks

I asked this of them earlier, and that is how do you know when or if you have a good idea scientifically without doing any work to vet it? You're in the shower, you're daydreaming, whatever, and you have this idea. How do you know if it's any good?

Tim Vogels

If it's coming back.

Paul Middlebrooks

If it's coming back the next day?

Tim Vogels

Yes.

Paul Middlebrooks

If you think of it again?

Tim Vogels

Yes.

Paul Middlebrooks

You don't write it down?

Tim Vogels

No, I write them down. I used to write them down religiously, but now--

Paul Middlebrooks

Then you can just read it and it comes back? Every idea is good? Every idea you write down is good?

Tim Vogels

No. Now what I usually do is I text them to a student of mine or a postdoc, and then they're like, "You're insane," or "You're stupid," more likely, or they just don't respond because it's--

Gaute Einevoll

You're not even wrong.

Paul Middlebrooks

They're busy writing that paper they have to write.

Gaute Einevoll

Exactly.

[laughter]

Tim Vogels

If it comes back, usually, I think a good idea will avail itself. If it's a good idea, you think about it for quite a while.

Cristina Savin

I sometimes discover-- I look at old notes a year before and stuff like that, and it's like, "Oh, I had this idea before. I completely forgot about it."

Paul Middlebrooks

You must write better notes than I do.

Cristina Savin

I write terrible notes, otherwise I would have done it the first time around and I wouldn't have waited for a year. For me, ideas that I find good, they don't have to be right, they have to be good, is ideas that I'm itching to find the answer. It's like I keep thinking about it, and I really want to do the numerics right now, I want to do the math right now, I want to talk my student into doing it right now. I get this vibe.

Paul Middlebrooks

Wouldn't it be cool if we could do this?

Cristina Savin

It's not that all of the things that we do are like that, but I think what keeps me going is the things that feel like that.

Paul Middlebrooks

Good. All right, guys, keep going. Thank you for your time.

Gaute Einevoll

Thanks a lot. Perfect.

Paul Middlebrooks

It's been fun on the boat, hasn't it?

Tim Vogels

Thank you. Yes,

Cristina Savin

Yes, it's been more wiggly than average.

Paul Middlebrooks

It's super wiggly right now.

Tim Vogels

Because we're outside, right?

[boat whistling]

Paul Middlebrooks

We're here with Mikkel again, the organizer, the brains behind the conference. I know you had help. Tone and everyone was a great help in putting it together. All right. Now you've had this thing. Was it a success? The title, Validating Models What Would Success Look Like in NeuroAI? The last thing that I brought up, and the very last thing that we did was this panel discussion about what would success look like in NeuroAI? There was a wide variety of responses, actually. We've had a lot of great talks and great discussions throughout the trip. Gaute, you can chime in here too, but I just wanted to get your reflection on how you think it went. A boat this year. Gaute thinks spaceship next year? Outer space?

Gaute Einevoll

Yes, you have to talk to yourself. You have to do like space station.

Gaute Einevoll

Yes, exactly.

Mikkel Lepperød

There's no arguing against that.

Gaute Einevoll

So that will also...

Paul Middlebrooks

Next year, we all get our own speed boats, and we all have headsets. We experience racing and giving talks.

Mikkel Lepperød

Yes, that's a good idea.

Paul Middlebrooks

What do you think? Was it a success? Was the workshop a success? How are you feeling that it went?

Mikkel Lepperød

I think it was a great success. I think everyone really enjoyed the conference or the workshop in terms of the scientific material, but also in terms of the social aspect. Just the trip has been really great. In terms of the science, I think there's been-- I had two big worries. One would be high winds and waves, and-

Paul Middlebrooks

Oh, the weather was [crosstalk]

Mikkel Lepperød

-bad weather so everyone got seasick. That would be terrible. That didn't happen. Yes, that's a success. The second thing I was worried about would be that all the talks were off target or that no one would discuss or talk, and it would be just like another science conference-

[crosstalk]

Mikkel Lepperød

-where everyone just gives all their data, and there's impossible to respond to it because there's so much details.

Gaute Einevoll

I think this format where you had first-- everybody had a 40-minute allocation slot, and then 20 minutes for presentation, and then 20 minutes for discussion, that was very successful.

Paul Middlebrooks

Yes, that was good.

Gaute Einevoll

For this meeting, it was perfect.

Paul Middlebrooks

Did you learn anything? Is it going to change the way that you approach anything in your own work? I'll start off by saying I made connections and had conversations that gave me new avenues of thinking about my own work. That's been super valuable to me.

Mikkel Lepperød

One thing, with just rethinking about the workshop topic, I think continuing probing the community and the people doing the science on how we should do it right, I think, has been a major insight that this really is an important thing to do, like to ask these critical questions on how if you take a step back, like what would it look like if your model were actually doing something like the brain is doing and how would you measure that or--

Paul Middlebrooks

What are your satisfaction criteria?

Gaute Einevoll

I was surprised when you were asking when you were leading the panel debate at the end, Paul, that how many of you feel that you know what success would look like?

Paul Middlebrooks

There were about eight hands that went up, so I [crosstalk]

Gaute Einevoll

Yes, just to repeat the question.

Paul Middlebrooks

The question was, so on a scale of 1 to 10, and if you're 8 or higher in terms of feeling that you know what success would look like, then raise your hand.

Gaute Einevoll

Yes. How many, what fraction of the participants--

Paul Middlebrooks

They're about 30 here.

Gaute Einevoll

Just like this?

Paul Middlebrooks

Almost a third. 27.

Gaute Einevoll

I would expect that number to be higher and that it would rather be that people have different opinion of what this success would be like.

Paul Middlebrooks

Everyone who did raise their hand did seem to have a different opinion. I would expect that if people were honest, that it would be about that number.

Gaute Einevoll

Really?

Paul Middlebrooks

I didn't raise my hand.

Gaute Einevoll

No?

Paul Middlebrooks

No, I can't articulate it.

Gaute Einevoll

No?

Paul Middlebrooks

That's a problem. I know that's a problem. This is a good venue to explore that.

Gaute Einevoll

That's true. That was a little bit surprising for me. I think sort of, uh ...yeah.

Paul Middlebrooks

You raised your hand, Gaute.

Gaute Einevoll

Maybe because I'm a little bit like an outlier here in participants in terms of I come from this physics side of modeling, and I do use physics-type modeling, the brain as a physical system. Not that I'm not that I'm not interested in what the functions are and the other models. Here, the success is a little bit clearer in the sense that you try to mimic physiological data.

Paul Middlebrooks

That's why you like Andreas' work so much.

Gaute Einevoll

Yes, exactly. I think maybe that's-- Maybe when you come from physics, the idea of what the success is, whether it's a good idea or not, this is more clear, it's more imprinted in us.

Paul Middlebrooks

That's why when I drew that awful diagram, it had a lot of different little lobes of success and how those lobes could maybe attach to the different goals from the way people use AI as tools or models and stuff. Mikkel, you also raised your hand, didn't you, when I asked that question? You don't have to articulate it, but don't I remember you raising it? I think the question was, do you feel like not can you articulate it? You're an 8 or above?

Mikkel Lepperød

Yes. That has to do with what I was doing and thinking about before I started this workshop. I was thinking a lot about it before I started the workshop. If you would add a confidence score as well.

Paul Middlebrooks

That is the confidence score. That's why it's a 1 to 10.

Mikkel Lepperød

No, it's a confidence that you know what it would look like, but it's not a confidence of how sure you are that it's actually-

Gaute Einevoll

Achievable.

Mikkel Lepperød

Achievable.

Gaute Einevoll

That is a different thing. I know what success would look like. I'm not quite sure.

Mikkel Lepperød

If it's possible, right?

Gaute Einevoll

Yes, if it's possible.

Mikkel Lepperød

I mean, you could say success would look like, success would look like us building a perfect brain, like that. No one would argue that that's not success. [crosstalk]

Paul Middlebrooks

No, I would argue that-- No, no, no. If you've done a perfect brain, then what you're left with is a brain not necessarily the understanding of its functioning, how it works.

Mikkel Lepperød

You would need to understand how to build a brain if you were supposed to build. I'm not talking about growing it from you're putting some genes together. That's not-- If you can build a robotic brain or whatever, then ...

Gaute Einevoll

Or even make a very detailed model. It's true that that would be very hard to understand. Then you could have that as a starting for probing, and it would be the perfect test animal as we discussed.

Paul Middlebrooks

The model has to behave correctly.

Gaute Einevoll

Absolutely. It has to fit the experiments.

Mikkel Lepperød

If it would be a perfect brain, that would be a success, I think. You could build it, but you wouldn't be-- That's--

Gaute Einevoll

Some people had this idea that it really success would mean that you could make a model of an individual. That's an individual brain, like Konrad's brain, right? That, I don't think, is very realistic. I'm thinking more in terms of some kind of average brains on general properties and maybe what is the difference between a healthy brain and maybe a psychotic brain or different kinds of brain states. More like the average. Actually, mimicking a particular brain would mean that you need to replay the whole history, probably, with environmental inputs.

Paul Middlebrooks

Shouldn't we start with a below-average brain like mine, though? Isn't that more feasible?

[laughter]

Gaute Einevoll

No. You know the Tolstoy saying that's like he never wrote about happy families because it's like a happy family is only happy in the same way. There's so many ways to be unhappy.

Paul Middlebrooks

There's so many ways, yes.

Gaute Einevoll

Maybe that's the same thing that unhealthy ways. There's so many ways to be below average.

Paul Middlebrooks

Below average. Yes, below average. Any parting thoughts? We have to leave happy. This conference is happy.

Gaute Einevoll

Yes, this is happy. How do you want to-- This has really been a great success. I talk to people even when you're not there.

[laughter]

Paul Middlebrooks

That's the criteria, right?

Gaute Einevoll

They are extremely so happy and learning. Everything was perfect. Do you have any plans for-- You don't have to do it today. Now we have still that knowledge room at the moment.

Paul Middlebrooks

The thing is, you have to keep it small to be useful.

Mikkel Lepperød

I think that was one of the success criterium, actually, keeping it small enough that we could become a small group of friends, basically.

Paul Middlebrooks

I know what you're going to say. You're going to say, because you said it during the panel, that you'd like to have more people from the computer science side. Right?

Mikkel Lepperød

Yes. I think that would be really interesting to see. I don't know how that would look like, though. It would have to be people from the other side or computer scientists that are genuinely interested in the topic. It couldn't be a bunch of people that were just hoping maybe that can get some cool ideas from neuroscience and just take it and build something. It would preferably be someone that wanted to be in the community. That would be really cool.

Gaute Einevoll

That's always a challenge, I think. On one hand, you want a broad set of perspectives. On the other hand, you want people to have some interface so they actually can communicate, right?

Paul Middlebrooks

Right.

Gaute Einevoll

If it's too broad, then you're not able to-- half of the audience doesn't understand what the other half of the audience is talking about, and then it's difficult to get to--

Paul Middlebrooks

Yes, that's it.

Mikkel Lepperød

One thing I really want to say, one thing that I think is really important and I think has also been a really big part of this success is to bring people with different backgrounds together. It's important that they have some common focus or else we'll just talk by each other. If you have a focus that will make people think in some parallel or same direction, it's great.

Paul Middlebrooks

You did that because there are people working on synapses and spiking. There's people working on [Kwabena's](#) neuromorphic, essentially. What you call-- you set it up as the implementation level. There are people at the representation level, the algorithmic level, the computational. We had just a wide variety. Like I said before, I'm repeating yourself now, but I think that you achieved that already.

Mikkel Lepperød

That's definitely one of the coolest things with the scientific interaction or conversation or this sociology of bringing together people from different mindsets and having them talking together. I think there's some magic that can happen there.

Gaute Einevoll

Great. Thanks a lot, Mikkel-

Paul Middlebrooks

Yes, thanks again for the invitation.

Gaute Einevoll

-on behalf of the participants and on behalf of the field, to be a little bit-

Paul Middlebrooks

Oh, wow.

Gaute Einevoll

-pompous. Yes, exactly.

Paul Middlebrooks

Pompous. Pompous Gaute.

Gaute Einevoll

Exactly. That's a new one.

[laughter]

Paul Middlebrooks

All right. Thanks, Mikkel.

Mikkel Lepperød

Thank you.

[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 Dunovan. Thank you for your support. See you next time.

[music]

Subscribe to “[Brain Inspired](#)” to receive alerts every time a new podcast episode is released.