THE TRANSMITTER | BRAIN INSPIRED

Grace Hwang and Joe Monaco discuss the future of NeuroAI

Hwang and Monaco organized a recent workshop to hear from leaders in the field about how best to integrate NeuroAI research into the BRAIN Initiative.

4 DECEMBER 2024 | by PAUL MIDDLEBROOKS

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

Joe Monaco
Maybe this interdisciplinary process, whatever we're calling NeuroAI, can help dislodge some of these stale debates and arguments and help us get to a new conceptualization of what's really going on in brains.

Grace Hwang
There are some people that think NeuroAI is about using AI to understand neuroscience, and then there are people that think NeuroAI is about using insights and principles from neuroscience to improve AI. We're really more interested in the convergence and the reciprocal aspects of NeuroAI.

Joe Monaco
You can ignite it, essentially, and then the connections are such that the nonlinearities all line up, and you get self-reinforcing, self-supporting, self-sustaining activity. That's the basis. That's the base computational unit.

Grace Hwang
I turned out to be the only student that had the right answer.

Paul Middlebrooks
Is that because you used a for loop and just did it the careful, slow way?

Joe Monaco and Grace Hwang
Yes.

[music]

Paul Middlebrooks
This is "Brain Inspired", powered by *The Transmitter*. If you're a longtime listener, you may recognize that I added back in Chopin. That is at the demand of one of my guests today, Joe Monaco. Joe berated me, you could say, for removing it a long time ago from the introduction. He took it upon himself to play it, record it, and send me the recording. It's back. I hope you're happy, Joe. Joe and Grace Hwang, the other voice you just heard, will introduce themselves in a moment, but I will introduce them now as co-organizers of a recent workshop that I participated in, the 2024 BRAIN Neuro AI Workshop.

You may have heard of the BRAIN Initiative. In case not, it is a huge funding effort across many agencies, one of which is the National Institutes of Health, or the NIH, where this recent workshop was held. The BRAIN Initiative began in 2013 under the Obama administration with the goal to support developing technologies and implementing those technologies to help understand the human brain so that we can cure brain-based diseases. The BRAIN Initiative just became a little over a decade old now, with many successes under its belt, like the recent whole brain connectomes you may have heard of, and discovering the vast array of cell types, and many others. I'm not going to list them here, but I'll point to a reference for you to learn more.

Now the question is how to move forward. One area they are curious about that perhaps has a lot of potential to support their mission is the recent convergence of neuroscience and AI, or what has been recently coined as NeuroAI, for better or worse, as we discuss. The workshop was designed to explore how NeuroAI might contribute moving forward, and to hear from NeuroAI folks, people doing the NeuroAI research, to hear how they envision the field moving forward. You'll hear more about that in a moment. That's one reason I invited Grace and Joe on.

Another reason is because they co-wrote a position paper a while back that is, among other things, an impressive synthesis of lots of concepts in the cognitive sciences and neurosciences, and history. We talk about that. It also proposes a specific level of abstraction and scale in brain processes that may serve as what Joe calls a base layer for computation. The paper is called "Neurodynamical computing at the information boundaries of intelligent systems." All right, you'll learn more about that in this episode as well. Okay, I don't want to yammer on here. Let's get you to Grace and Joe.

There are lots of show notes in this one, to workshop-related stuff, and to many of the papers that Joe and Grace reference. Those are at braininspired.co/podcast/200. I forgot. It's the 200th episode. That is awesome and amazing. What a fitting way to bring in 200, talking about a NeuroAI workshop, among other things. Patreon supporters, you are the best. Let's have a live chat before the Christmas holidays, if you're up for it. I'll be in touch about that. Go to braininspired.co to learn how to support the show on Patreon and join in on fun stuff like that and get full episodes all the time. All right, here are Grace and Joe.

[sound effects]

Yes, we're starting. We're starting.

Grace Hwang
Okay.

Paul Middlebrooks
You guys just blew my mind. This is the beginning. I've been interacting with both of you. I had no idea that you were a married couple because my interactions with you were very professional because I just came back from this NeuroAI BRAIN initiative workshop. Then 30 seconds into us speaking to each other, Grace or one of you said, "You do know that we're a couple, right?" No, I did not. First of all, hi, Joe, hi, Grace. Thanks for being on the podcast.

Joe Monaco
Hi, Paul. Thanks for having us.

Grace Hwang
Hi, Paul.

Paul Middlebrooks
This is a little bit different than the way I normally do things. Could you very briefly state your name and occupation, or not your name but your occupation? Grace, we'll start with you.

Grace Hwang
I'm Grace Hwang. I am a program director at the NIH, at the National Institutes of Neurological Disorders and Stroke. I support the BRAIN Initiative full-time.

Paul Middlebrooks
Joe?

Joe Monaco
I am Joe Monaco. I am a scientific program manager, and I am a contractor for The Office of the BRAIN Director at the NIH. We are housed under NINDS. I'm there with Grace, but I work with the brain director, and I work with all of our internal brain teams.

Paul Middlebrooks
I just worked with both of you, among many other people who worked with both of you because you both put in an absurd amount of work to organize this recent BRAIN Initiative NeuroAI workshop. Now I don't know where to start because your partnership has gone back a long time. Grace, I have recently learned that it was an interesting way that you guys met originally and came to form an intellectual partnership. Maybe we could just start there.

Grace Hwang
Well, we met originally in 2004 at Brandeis University when we both were taking a computational theoretical neuroscience class, under Larry Abbott back when he was there.

Paul Middlebrooks
Yes. That's a famous class. I feel like a lot of people matriculated through that class and have fond memories of it.

Grace Hwang
Indeed. Joe and I started our first collaboration. this was when I had just gone into computational neuroscience and was learning how to use MATLAB for the very first time and had to do a homework problem. I was the only student that had a different answer than everybody else because I had not figured out how to unroll my loops. I was literally writing all these loops.

Paul Middlebrooks
Nothing wrong with a for loop, although it's better to not have them if possible.

Grace Hwang
That's right.

Joe Monaco
That plays into the story.

Grace Hwang
Because I was the only person who wrote a for loop, I had a different answer, and I thought I had the wrong answer. Even the TA had a different answer than me. I sat in the computational annex for days trying to find the bug in my code. Then at 1:00 or 2:00 in the morning before the homework was due, comes in a student very quickly writes his code and is about to leave with this correct answer. I asked him, "Can you just look at my code?" He looked at it and he said, "There is nothing wrong with your code." An hour later, he found a bug in his code that every other student who were super programmers made. I turned out to be the only student that had the right answer.

Paul Middlebrooks
Is that because you used a for loop and just did it the careful slow way?

Joe Monaco
Yes, because all of the other, computer nerds in the class, they know about vectorization and MATLAB is very slow when you write manual for loops. Everyone got the same wrong answer. They thought it was the right answer. Looking through Grace's code, she sequenced the order of operations absolutely correctly. Then I was able to figure out when you vectorize it, it actually changes the order. We were doing classical conditioning. The order in which you update the parameters of the synapses really mattered.

Paul Middlebrooks
You guys were already married at that time, right? No.

Grace Hwang
No, no, no, no.

Paul Middlebrooks
You were just saying that--

Grace Hwang
That was when we met.

Paul Middlebrooks
Yes. Yes.

Grace Hwang
We've been married for 10 years. It took us a long time to get married in 2014.

Paul Middlebrooks
We had to get our PhDs first. Backing up even further, Grace, I understand that the only reason you went into computational neuroscience was because you developed mouse allergies. Is that correct?

Grace Hwang
Yes. I initially went to Brandeis University to study biophysics and structural biology. I was building a single photon molecule microscope trying to study the translocation of HIV protein. That required that I work with protein chemicals and rotate it through a mouse lab. I had really bad allergies where I was in hives and had an EpiPen. It was not a sustainable lifestyle. In my third year of graduate school, I changed labs into a computational memory lab where I joined Michael Kahana back when he was at Brandeis University.

Paul Middlebrooks
The interesting thing about that is I've only physically been near Joe a couple of times now, but I think I developed an allergy.

Joe Monaco
I don't know how to take that.

Paul Middlebrooks
No. What you just co-organized was a NeuroAI workshop. Then how did you guys end up coming together in that steed?

Joe Monaco
Out of the two of us I'm the one who was a neuroscientist from the beginning. I was a computational neuroscientist and a theorist. I did my PhD on grid cells and place cells and modeling how they might be related through remapping. It's very important computational transformation and

hippocampal studies. I brought that through to expanding outwards and thinking more about the complexity of behavior throughout my postdoc. I joined Jim Knierim's lab at Johns Hopkins where he had a wealth of experimental data of freely moving but on track rats, basically navigating in a clockwise circle a number of times.

You can very closely track the behavior, where's the position of the head and body. Then you can track the emergence of place field activity over time. That was the basis of me getting into a complex neurobehavioral analysis and thinking about from an organismal perspective, what's actually going on here. Do we have all these just internal computational representations? How does that translate into what this interesting little animal is doing on a moment-to-moment basis? That's where I started thinking more deeply about complex temporal dynamics and behavior. Eventually Grace was off doing other things in other fields. She's a polyglot of science and technology.

Grace Hwang
After I graduated with my PhD, I went to the MITRE Corporation, and I was developing optical biosensors to detect pathogens from exhaled breath. I did some working in government at IARPA as a CETA contractor and also DARPA as CETA contractor. It wasn't until 2015, when I went back into "academia". I joined Johns Hopkins University of Applied Physics Lab as a program manager to run their applied neuroscience program. Initially, I was an assistant program manager.

That was the first time since 2005, did I start to pay attention to neuroscience again. It was so exciting because neuroscience had just accelerated. I started going to site for neuroscience meetings, but I wasn't following Joe's research at all. It wasn't until the 2017 Society for Neuroscience poster session did I learn about Joe's discovery of phasor cells. This came out of his research with actually Kechen Zhang and Tad Blair. I was so intrigued by phaser cells because unlike hippocampal place cells that maps phase-

Joe Monaco
-they're mapped asymmetrically to the traversal of a place field. Very famously in hippocampus, the pyramidal neurons there, the place cells will start firing at a later phase of the theta rhythm. Then as the animal moves through the place field, each spike will become earlier and earlier within the theta cycle across subsequent theta cycles. If you plot the distance through the field against the theta phase of spikes, you get this nearly monotonic decrease in the phase, this advancement. That's called phase precession. That's a very robust field. It's thought to be related to sequence learning and very important things that like CD3 is doing in hippocampus.

That's the highly recurrent sub-region. With Tad Blair, I had a collaboration where we were looking at the one synapse downstream into the subcortex. We were looking at lateral septum. The septal bodies are very interesting. People don't record from them very often, but we were looking at what other phase codes might there be? It's a theta rhythmic brain area. I went looking for a different code, one that wasn't locked to a particular trajectory, but one that was locked to space. I was looking purely for spatial information in the timing code relative to theta oscillations and found it. I coined the name phaser cells. There's a couple other models out there called phaser.

Grace Hwang
This blew me away because, you were able to directly map phase to place. I was working at the Polyphysics Labs in the Intelligent Systems Center surrounded by roboticists. My immediate thought was, well, if we could use this internal phase code to do self-localization and mapping, wouldn't that be cool? On that very same day, there was this other paper that came out. It was titled “Swarmulators.” These are swarmulators that can sync and swarm using an internal phase variable, which came out of Kevin O'Keefe and Steven Strogatz's lab. The two things got me thinking, wow, maybe there's a there there for using Joe's discovery of phaser cells and controlling multi-agent robotics. That was how our collaboration reinitiated back in 2017.Then AI became pervasive. We now just begrudgingly call it neuro AI because people tend to know what that means. There are some people that think neuro AI is about using AI to understand neuroscience. Then there are people that think neuro AI is about using insights and principles from neuroscience to improve AI and AI could be hardware, software, or a combination of the two. We left it open-ended. We told people for the purpose of the NeuroAI workshop, we're really more interested in the convergence and the reciprocal aspects of neuro AI and not in the feed forward using AI for any science approach.

Joe Monaco
I wouldn't say begrudgingly, I would say this is where the opportunity is. In this interview, we both have to be clear about when we're speaking from the perspective of the BRAIN Initiative, when we're speaking from our own scientific perspective, from the perspective of the BRAIN initiative, there's a lot of interesting opportunity here. In my opening remarks, I showed a figure from, Brad, his review paper, of four or five main threads of how AI has evolved from different ways of computing through learning with data. The question is, how does all this come together? What do more brain-like forms of this type of computing by learning with data?

What does that look like going forward? There's been a few major inspirations from neuroscience from the brain, over the decades, since good old-fashioned AI in the '50s where, you had these symbolic approaches, Newell and Simon, coming out of cognitive science. Then, we're all familiar with the history of the back and forth there, the AI winters, connectionism, rose in the '80s, with the advent of neural networks and back propagation, for updating weights and, it's only in the last 10 or 15 years did the scalability come into play with a hardware that enabled the amazing advances that we've seen in the last 10, 15 years in what we now call AI computing or AI technology.

The convergence right here is really ripe. I think we should not be arguing about definitions necessarily, because people in cognitive science and neuroscience and artificial intelligence are very good at arguing about terminology and definitions. We could bring in the consciousness researchers if you really want to go.

Paul Middlebrooks
Leave the philosophers out, because they're the best at arguing about semantics, right?

Joe Monaco
Right. We do need people worrying about this. I think it's at this early stage and this exciting convergent period-- There's a lot of decades of thought and research going into all these different threads. They've all hit fairly related roadblocks, it seems. Cognitive science didn't become that fully encompassing research program that started in the '60s and '70s that Miller anticipated. Neuroscience has gotten wrapped around certain ideas of attractor dynamics and population geometry. We're trying to figure out if this is the right way to go or not and how to incorporate large-scale data. Maybe coming together, we could solve all these problems simultaneously.

Grace Hwang
I just want to defend my personal position a little bit about being-

Paul Middlebrooks
Begrudgingly.

Grace Hwang
-begrudgingly with the word AI, because this would be the second time that the word NeuroAI was put in a program or in a concept that I'm working on. When I was at the National Science Foundation back from 2020 to 2023, I created a program, a topic that was part of the Emerging Frontiers in Research and Innovation program called Brain-Inspired Dynamics for Engineering Energy Efficient Circuits and AI. The original name of that program did not have AI in it. Because AI was so hot, my former NSF leadership said AI's got to be in the title. That was where I was coming from, Joe, is we had to use the word AI because it's gotten so in vogue these days.

Joe Monaco
I think it's useful to use the words that people are using, right? It's a new term. I think it's not fully well-defined. That's what this workshop was about. Let's bring together-- Grace and I have been going to NeuroAI conferences and workshops, a lot of them over the last year, as I know you have as well, Paul. You see a lot of themes emerging. BRAIN is interested. Let's explore what potential roles look like. Is there a piece of this where BRAIN can help, where it fits within BRAIN's mission to go forward? We want broad input from the community for the workshop. We wanted broad input from the community. I think we got it about helping us identify what those opportunities are to be considered further.

Paul Middlebrooks
Yes. What I wonder is, so I hear when you have a new hot term like NeuroAI, it makes me think of, I'm not sure what hats you want to put on if you respond to this, speaking scientifically or from the NIH, but you have to put the word mechanistic in all of your papers now, because mechanisms became the hot thing and computational neuroscience is dominated by mechanisms. I've heard the sentiment I'm about to express expressed among grad students just in the past few weeks, of everything now is NeuroAI. If you want funding for anything, you just call it NeuroAI.

Therefore, you don't really have to worry about how what you're doing fits into NeuroAI. If it's ill-defined or not defined at all, then-- Here's the worry is that then there's going to be this surge of grants and everyone's going to use the hot term NeuroAI, and that's going to bolster their chances of getting funded. I'm not sure if there's a solution to that, because that's just the name of the game. I don't know if you have a response to that.

Joe Monaco
It's always something like that. You can't control the community. You can't control people. Essentially, you have to communicate in a way where people need to be clear about what they're doing when they're applying for funding. Obviously, the NIH and NSF and other funding agencies have scientific review processes in place to discern who's maybe following hype and over-claiming or overusing words versus where the real advances are. It's obviously not a perfect system.

Paul Middlebrooks
Right, well, of course.

Grace Hwang
I'll say, from my experience, when I was at the NSF, for the BRAE topic, we had very strict solicitation-specific criteria that would filter out people who are just using the buzzwords to try to get in. I think having strong review sections would avoid these kinds of problems. I do want to say that there is a true inflection point here in that the technologies that's been enabled by the BRAIN Initiative is allowing us to observe circuits in animals across many different spatial and temporal scale. There is a real opportunity. I think just putting the word NeuroAI on any project is going to be easy to filter out.

Paul Middlebrooks
Perhaps so, yes. I think that that little light bulb went off in my head when Grace used the term begrudgingly, because that sense that I get from these handful of people, it's sort of an eye roll, oh, that must be NeuroAI, because that's just a hot term or something. I'm wondering, man, I already feel the backlash of this emerging field.

Joe Monaco
I originally wanted to call this workshop the BRAIN NeuroAI and Theory Workshop, because I wanted that convergence point to be focused on advances in theories, theoretical frameworks, and theory-driven models, because I think there's so much there. Personally speaking, that's where I think a lot of the obstacles have been. There are a few unifying theoretical frameworks in neuroscience, and you've discussed them on this podcast over the years with a lot of people who have made major contributions to those theories.

I don't see that Kuhnian process of the field, testing confirmatory hypotheses, falsifying different theories, and making progress. I think maybe this interdisciplinary process, whatever we're calling NeuroAI, can help dislodge some of these stale debates and arguments and help us get to a new conceptualization of what's really going on in brains, and which, by the way, are fundamentally embodied and inherently integrated as biological systems, which is different from AI.

Grace Hwang
I also think being a little loose with our definition is okay because we were able to bring in the neuromorphic community.

Paul Middlebrooks
Yes, it was a huge-- Maybe let's talk about the reason that the workshop existed and how you decided how to frame it, how to organize it. Because the neuromorphics was a larger part of the workshop than I had anticipated it might be if one just said we're going to have a NeuroAI workshop, and that's all I had to go on.

Grace Hwang
For me, coming from the NSF to the NIH, I was shocked that there was very few investment in neuromorphics at the NIH. If you go into the NIH Reporter, you type in the word neuromorphic, you'll probably get 60 returns and under $20 million in investments since the late '80s, whenever NIH Reporter started compiling information.

Paul Middlebrooks
Is that because neuromorphic is inherently slow? Is that the reason?

Grace Hwang
It has been as slow as we heard from the last two days, but I think the other part of it is just that the neuromorphic engineers and the neurotech and biomedical engineers go to different meetings. They don't really talk to each other. For me, it was really important to bring, to close the loop between neuromorphic and neuroscience so that we can better benefit brain health and not just brain, but health in general. That was very critical for me. I actually co-hosted a workshop with my colleagues at the NSF with co-funding from NIBIB and NINDS in late October in Baltimore. It was a workshop called the Neuromorphic Principles in Biomedicine and Healthcare.

It was important to capture the health focus for neuromorphic and neurotech at this workshop. That was why the second day was designed to be more like a tutorial. The first session of the second day was intended to teach everybody what neuromorphic means, both large-scale digital computing style of neuromorphic as well as the small-scale mixed-signal analog computing and neuromorphic sensing.

We had Jacques-Malon Devery, as well as Ralph Etienne Cummings, who are, who really, I think, helped teach the audience the differences between the different kinds of neuromorphic technologies and how they may or may not be useful in healthcare. Then the second session was to really bring it home and have clinicians talk about the value of neuromorphic technology.

Paul Middlebrooks
Yes, just as an interjection that, I've done some conversing with people, like I just said, reflecting about the workshop. One of the things I said to another academic was, I was surprised that there was so much neuromorphics in the workshop. This person said, "What's neuromorphic?" I was like, whoa. They were a neuroscientist.

Joe Monaco
It's surprisingly not very well known. That's one of the exciting opportunities and why we wanted to have healthy representation from neuromorphic approaches.

Paul Middlebrooks
Wait, because we should, Joe, maybe just say what neuromorphics is, because I realized there's thousands of neuroscientists, I guess, who don't know what neuromorphics is.

Joe Monaco
Let's introduce it. This question was asked by an audience member at the workshop, what is the definition of neuromorphic? Nobody wanted to take that on. I admonished folks not to get into debates about definitions.

Paul Middlebrooks
Yes. I just need broadly, what are we talking about?

Joe Monaco
I'm saying, so Kwabena Boahen, who is also a leader in neuromorphic computing, he's put forward, it's computing that scales. It's scalable computing. In order to have fundamentally scalable computing, you need to be more brain-like. You need to have memory on compute. The closer you get to the brain, which is very, fundamentally a memory on computing system, then you break or bend some of the scaling laws that make it difficult to scale up conventional CPUs, GPUs on CMOS processes.

Neuromorphic, I would say, it comes to the Feynman quote that we see everywhere. There were two good Feynman quotes that came out at the workshop. The one that you see everywhere is, paraphrasing, what I cannot build, I cannot understand, essentially. The neuromorphic engineering community is folks who have been trying to build it. From the synapse level to the cellular level.

Grace Hwang
The original lineage, Carver Mead lineage of neuromorphic engineers, were people that were trying to emulate these channels on a chip and creating spikes and characteristics that are comparable to what you would measure off of a cat pyramidal cell. That was the original Nisha Mohawad 1991 paper that I think first popularized neuromorphic. That's what Ralph referred to as old school neuromorphic. Since then, there's been all sorts of development, which is why it's so hard to define neuromorphic. It's a word that, it means different things to different people. It could even mean principles. As Kai said, you want to have your device that operates with the statistics of their signals on the brain. Some people even think of neuromorphic as principles.

Joe Monaco
They could be physical and material principles as well.

Grace Hwang
Guiding principles. Whether or not you actually use a neuromorphic hardware, it's okay. There's people in the neuromorphic community that think of neuromorphic as hardware. There are others that think of it as design principles. It's hard to define neuromorphic, but for me, it's brain-like. It usually operates on spikes, but not always. Most importantly, it's energy efficient. Six orders of magnitude energy efficiency, and it's adaptive to the user. It's a system that can evolve with the user. Those are the four things I think that really stood out and was discussed at the panel. People went as far as saying, because it's adaptive, it's less hackable, because you don't actually need a computer to run it. The on-chip device is self-containing.

Joe Monaco
I'd say it's important to distinguish not all neuromorphic hardware is learning hardware. A lot of the test bed systems that have been developed--Intel had a program called Loihi. There were two generations of these Loihi chips. It's an academic research partnership program. You applied and you can get some of these chips to devise models for and test and run them.

You could have a large number of spiking neurons, essentially, you integrate and fire neurons, and you could just run that physically on these chips. It was very difficult to implement spike-based learning rules like STDP on those chips. There's other kinds of chips which are more amenable to implementing different types of learning rules. That's something the field is still working on and trying to figure out. It's one of the big questions going forward. How do you have these be adaptive in a safe way?

Grace Hwang
That's a really good point. There are some chips that are designed for inference only and others that are designed for learning.

Paul Middlebrooks
It sounds like the whole thing was focused on neuromorphics, which is not the case at all, but this is just a way into talking more broadly about the workshop. I totally understand why neuromorphics then would be something to put a large focus on. Then the question is what Joe was talking about earlier about needing to bring in theory and get all these people under the same roof to bring together these otherwise disparate ideas to work on the problems. One could wonder, well, what is that? What is what you guys were just describing? What does that have to do with anything related to theory? Did we make progress on that or is that an ongoing challenge?

Joe Monaco
It's an ongoing challenge. Like Grace said earlier, there hasn't been enough conversation between these fields. That was very intentional on our part to bring together these different communities to start having that conversation. Even within neuromorphic engineering, there are folks who want to use large-scale neuromorphic computing systems to basically run large models of the brain.

They need large-scale brain data to build those kinds of models, connectomes and cell types and that thing. You can use that as a test evaluation system, a modeling, a simulation system to confirm or falsify theories about how certain aspects of these networks work. Then there's the small-scale side of it, energy-efficient analog or mixed signal devices that can be distributed to the edge to do brain-like, neural-like intelligent computing in a wider array of applications. That's more towards the translational end. I think it covers the full spectrum.

Grace Hwang
I want to go back to the theory. I don't think we talked a lot about it. Brad Eimony actually came to the NIH and gave a Wednesday afternoon lecture series talk on October 23rd, where he talked about how large-scale neuromorphic computing could inform new theories, how you can make observations at scale that you just don't see in smaller circuits.

Joe Monaco
It's much more difficult and particular to the types of models that you're building, how to scale them up. I think there's opportunities here to break some of those scale barriers within computational neuroscience and theory-driven modeling and take it to the scale where the things that we care about happening in brains can actually be studied in a principled way.

Paul Middlebrooks
We've been focusing on neuromorphics and then we delved into theory. I didn't say in the beginning, the backdrop of this is that the BRAIN Initiative is 10 years old now. Part of the driving force, and correct me if I'm wrong, with putting this workshop together was to figure out where the future roadmap should lead or what avenues are explorable and should be explored moving forward. Is that a fair assessment?

Joe Monaco
We wouldn't use the word should. We want to get the shoulds from the community. We have the neuromorphics. We've got the people doing metrics and benchmarks. We've got the people thinking about natural intelligence capabilities. How do they all come together? What are the main priorities and opportunities they see? We want insight about that.

Paul Middlebrooks
You got some shoulds, at least during my last session there. There were some shoulds from the audience, which was fun. Good.

Joe Monaco
Right, well, so from the BRAIN Initiative's perspective, we want to get all these pieces together in this jumbled puzzle and then figure out which of those pieces makes sense for the BRAIN Initiative, potentially, to contribute to lend the type of work that we do in the BRAIN Initiative.

Grace Hwang
Another thing I would say as a co-organizer is we invited a lot of other funding agencies to this meeting to be part of the workshop, because a lot of these problems aren't necessarily BRAIN's problems. It's like coming up with the next super-efficient computing system is a great idea, and it's great for humanity, and it's great for solving the carbon footprint, but that probably lies in a different agency's mission space. However, the knowledge the BRAIN Initiative collects and continues to generate is useful to that longer-term mission, and I was really hoping that would be clear from the funders' panel. I don't know if we hit the mark there, but this is really a collaborative effort where lots of agencies are interested for different reasons.

Paul Middlebrooks
Yes. I want to complete the circle on just what the workshop was about and how you decided what kinds of topics to bring in. When I think of NeuroAI, I think of like the... It's not traditional. It's new, but it's not new, but it's testing using some artificial intelligent models, neural network models as proxies for brain processes and then asking whether you get something brain-like representationally out of those systems. Early work with convolutional neural networks and modeling the visual object recognition system on those networks.

There's been a lot. Since then, there's been a lot of, recurrent network work like that. Then someone like Andreas Tullius, who works on what are called foundation models, was represented. That side of it was also represented, but maybe you guys could speak to what other kinds of things that you wanted to bring into the workshop.

Joe Monaco
Our personal scientific opinion and perspective at that time, which came out of the symposium that Grace and I co-organized for the 2020 Brain Investigators Meeting. That was a great panel. We had, Conrad Kording and Zach Pitko and Nathaniel Duh, Kanaka Rajan, and I'm forgetting, Brad Pfeiffer on that and we called it, Dynamical Systems Neuroscience and Machine Learning. That's what we were bringing together, but it's a prototype for NeuroAI, thinking these ideas through. The perspective that Grace and I wrote coming out of that was that a lot of the problems seem to be from this purely computational perspective.

That's the perspective that's grounded in this almost traditional brain as computer metaphor that has pervaded all of these fields. It's pervaded cognitive science, AI, neuroscience. Neuroscientists talk constantly about neural encoding, neural decoding representations, talk about representations without, as these computational constructs. That once you have representation and it has, to the experimenter, some explainable relationship with what we think is going on in the animal. We put it in a task. Okay, so it needs to solve, go left at this point after seeing this cue and we see correlations to the right kinds of things. It's like, okay, we found the computational representation. That is the explanation.

We're done here. I think the 'we're done here' part of that has been the barrier because you're not actually done there because there's still behavior that needs to be in the loop. Behavior is this moment-to-moment dynamical coupling between brain and body, between body and environment. That's why we wanted to expand outward and bring in ideas from embodied cognition, the 40 literature and, not to sign onto that, but to say, hey, there's something here about these massive, distributed feedback loops through the environment that are a key part of what's going on in cognition and animals. That's where we took this.

Grace Hwang
Day 1 was designed to be all computational intelligence whereas day two was more of the embodied neuromorphic translation.

Paul Middlebrooks
Even during day one, I was impressed with how many people were reflecting on the importance of embodiment. It came up a lot.

Grace Hwang
That was not planned. It was a surprise.

Joe Monaco
I may have, I may have planned that.

Paul Middlebrooks
Did you plant that seed? Is that what?

Joe Monaco
For instance, in the same issue in which our paper came out, there was another paper called *Deep Intelligence* by Ali Manai, who's a electrical engineering professor at University of Cincinnati. I was aware of that work and he's done computational neuroscience for a long time now, since I started in the field in hippocampus. I was aware of him.

Paul Middlebrooks
He had a very evolutionary perspective on it as well.

Joe Monaco
He has a very holistic perspective on biology. Biological organisms are inherently integrated. They're integrated through evolution, phylogenetically. They're integrated through development, ontogenetically. They're integrated through learning and aging and experience. You have to keep coming back to that because that is well, at least that was his perspective that he was taking in the paper and in the talk that we invited him to give, comparing natural intelligence with AI. There's so many important distinctions that you can make, but I think that's one of the key ones.

Paul Middlebrooks
It just occurred to me that holistic neuroscience would be a great term, except that it would be associated with holistic medicine, I think, which the word holistic has some positive and negative connotations.

Joe Monaco
Right. It gets into the impulse for reductionism and the counter movement of looking at downward causation and emergence.

Paul Middlebrooks
I just meant the science of holistic medicine is sometimes questionable. To be a holistic neuroscience, someone might see that and think, oh, it's woo-woo or something like that. On the whole, that's a pretty good phrase.

Joe Monaco
I think it's woo-woo if you do ignore the internal computational representations. You can't ignore it. That's why across two days, we had the focus on, personal opinion, I was calling it the mainstream NeuroAI, let's figure out how to map these task-constrained AI models to what we see in the ventral visual stream. Like you said, a lot coming out of that. People are looking at dorsal stream and people are looking at motor system and other areas.

Paul Middlebrooks
Cognitive maps.

Joe Monaco
Cognitive maps.

Paul Middlebrooks
Yes. You name it.

Joe Monaco
Cognitive maps are maybe the clearest example of actual high level cognitive encoding in the brain, at least that's my personal opinion as a hippocampal, researcher. Hippocampal chauvinism.

Paul Middlebrooks
Applying neuro AI models to account for cognitive functions. Cognitive maps has been a big--

Joe Monaco
Absolutely. It ties into the dimensionality reduction, the, task based low dimensional manifolds. We're hurting hundreds of thousands or millions

of neurons now. There's no way to visualize that. If you just throw everything into a U map, you get some interesting colored splotches on your screen, but it doesn't tell you how to interpret what's happening.

Paul Middlebrooks
Oh my gosh. You're speaking to a-- I make some really pretty, naturalistic behavior neural U map graphs right now, and gosh, they're pretty, but they're not the solution. I'm not done.

Joe Monaco
When this ties into the whole, interpretability explainability and mechanism, discussion, how do we get at what the important factors are that are driving that high dimensional neural activity?

Paul Middlebrooks
Wait, now you're just jumping the gun and going right into maybe that's the way that we should do it is talk about some of the topics in the paper and then bounce back and forth. I don't want to also not come back at least to-- Well, actually let's hold off on going-- I know that they're all related--

Joe Monaco
You're saying that you need both, right?

Paul Middlebrooks
Yes, okay. Then I'm not sure what we've missed about the workshop, but I wanted to get your general reflections on how it went.

Joe Monaco
Yes. From the point of view of the workshop, we were both incredibly thrilled and pleased at the discussions that we had and thank you Paul for stepping in as a discussant on the first day and for helping with the wrap up on day 2. We've been reviewing the recordings. When you're in the middle of it, you don't get to really listen and experience it. It's really great conversations and discussions and questions that we had.

Paul Middlebrooks
At one point in my first discussant panel thing, I got to yell at Terry Sanofsky, and I thought, "Oh, why am I saying something negative to Terry Sanofsky about what he said?" As I was saying it, I was having this moment like, oh, you're not in a position to even talk about this with him. He was, it was fun. It was fun. Because he's a hero, intellectual hero to many people, including myself. You're in that situation and this is maybe very meta, so I apologize, but you're talking to your heroes sometimes, and you realize either are they a colleague or are they a hero, and it's, not surreal, but an interesting feeling?

Joe Monaco
This is why we brought everyone together. We want the leaders of the field who have been around and just driving things forward for as long as anyone can remember, sorry, and his colleagues had won the Nobel prize, just a few weeks before the workshop. If you talk to, people around that and you go back and read those papers from '84, '86, Terry is a coauthor in all of that.

Paul Middlebrooks
I was wondering how he felt about that. I'm not sure if it's the right place to discuss that, but I imagine a lot of people have wondered, does he feel like he was missing from that?

Grace Hwang
The day the Nobel Prize was announced, the tell you right neuromorphic community overtly wondered why Terry was left off.

Paul Middlebrooks
Yes. No, I've seen, I saw a few things like that. Yes.

Grace Hwang
That was shared upon all of us who attended Telluride this year. It was good to be able to acknowledge Terry's contribution at the workshop through multiple talks.

Paul Middlebrooks
Yes.

Grace Hwang
I think, I think you were fine, Paul. You were correct.

Paul Middlebrooks
Oh no, I, yes. He corrected, he corrected me, which was, which was wonderful because he correctly corrected me. Then I was like, all right, I'm not going to get into a back and forth with this person.

Grace Hwang
I have to say every time I felt something didn't go well, I went back and listened to the recording and I was like, oh, it went way better than I thought. I think the workshop, was really a great success. We, I learned things that I didn't expect to learn.

Paul Middlebrooks
I came away thinking that it felt like a win, a great success. I'm not sure if you guys want to reflect more on how you felt about how it went and maybe even what may have been missing that will happen next time or, how reflecting on what happened here affects how you think about moving forward.

Grace Hwang
I feel the neuroethical conversation was a really important one because NeuroAI is going to bring about a lot of new challenges and that Karen Rommelfanger's talk was really insightful. I think, if we were to have another one of these workshops in the future, I felt like we didn't give her a chance to ask her a question because we ran out of time. Actually, to be honest, though, I don't think he could do NeuroAI ethics justice in a one-hour session. There ought to be more conversations about ethics and as well as regulatory questions.

Joe Monaco
No, we can't speak directly about the future, but, clearly, brain put on this workshop is interested in the space and we agree. We think it was a great success as far as our goals of hearing from everyone.

Grace Hwang
Paul, do you think there was a scientific community that was not represented at the workshop that should have been?

Paul Middlebrooks
Oh geez, putting me on the spot here. I was re reading your position paper. Maybe this is a segue into that because, yes, there's a lot of embodiment, and I'm trying to reflect now myself because as I'm reading through this paper, the title of your position piece, "Neurodynamical computing at the information boundaries of intelligent systems." I'm reading through this thing again, and it's so rich and dense and makes the case for embodiment and the importance of environment, body, brain, continuous cycle interactions. I'm reading, I'm like, "Oh my God," I click on every reference and half the time I'm like, "Oh good, that's good. I've read that." Then the other half, I'm like, "I got to add it to my reading list. The irony is my original job at the end was to synthesize the workshop. To be honest, I had an idea of how I was going to do it, but I didn't have a set plan. It ended up being more of a moderated discussion, which was great. Then a lot of interaction from the audience as well. The reason why I said synthesize, and I think Joe might've mentioned that term earlier is because your position piece synthesizes a ton of stuff, with the goal of using so much historical perspective and what's maybe missing these days in AI, to synthesize what you call a base layer of computation. You're going to correct me on this. I don't have the exact quote, but a base layer of neural computation.

Joe Monaco
I think that's what we called it.

Paul Middlebrooks
Okay. I know it's called a base layer. You asked me what I thought might have been missing. It might have been that bringing together of the historical contexts and why these things are important. Then an interesting thing happened, that two people, Blake Richards and Zach Pitko, Zach at the very end said a great goal for us would be to record the connection strength of every synapse. That's such a reductionistic approach that is in line with modern reductionist neuroscience. It just flies in the face of what you guys argue for, in this position piece a little bit.

I thought it was odd that there was still this reductionist assumption underlying all these things to measure more. Here's the level that we need to measure at, and modern dynamical systems theory, manifolds, looking at larger populations that lower dimensional structure is somewhat antithetical to that story. If anything, I thought maybe the whole even manifold and more talk about levels and going across levels, what's the right level of abstraction, why it's the right level of abstraction. That's more on the theoretical side. If anything, I thought there could have been more of that, I think. Man, that was long-winded. I apologize.

Joe Monaco
It was more of a comment than a question. [chuckles] No, I totally agree. Here I have to be very careful to differentiate my perspective, opinion on the field as Grace and I wrote this paper a couple of years ago now. This did come out of a wide range of frustrations, which is why I went deep historically and I brought in ideas from philosophy of mind, philosophy of computation.

Let's think about what is computation? What does computation mean in the brain? What does information even mean in the brain? Because one of these other things that pervades all these fields, and I do give a historical capsule at the top of the paper, cognitive science, neuroscience, AI, is this--they have the same conception of what information is, and it's Shannon information, which, as we know from the '48 paper and the later one, it's about communication, where you have a transmitter and a receiver and you have a shared alphabet.

That may not be the right metaphor or framework for understanding information in the brain, especially how brains construct semantically meaningful structures and processes and dynamics.

Paul Middlebrooks
Shannon was even aware of this conflation of information with meaning. I could put in the show notes. He actually wrote up a very short piece saying, "Look, people, this might not apply to your field," because everyone was applying Shannon information to their fields. He was warning against his own work being applied too broadly and misconstrued.

Joe Monaco
I think I've seen that. It's reminds me-- but also Tony Zador in the opening keynote brought up the concept of the hardware lottery. In a sense, Shannon's information theory is a theory lottery. It provided everyone a readily accessible tool for like, "This is a really important concept. How do we measure it and grab hold of it?" It's like, "Okay, well, here you go. You just run these sums and averages over this kind of distribution." It's a purely a statistical process now. It's purely syntactic. I think it was Shannon, maybe in that piece, who said, "This is purely syntactic, this does not get at semantics of what this actually means."

If we think about this paper, "Neurodynamical computing and information boundaries", it's because there's different types of information and they are transformed across the boundary of an organism.

Paul Middlebrooks
Now you're not using information in the Shannon sense. You're using it in a semantic-

Joe Monaco
Right. What is different about information in a biological cognitive organism? It's that organisms construct a boundary. It's our skin, but it's also our exteroceptive senses. We have ways of taking in information. You have the whole universe of sensory input coming in, but then you also have the internally generated universe of goals and drives. Organisms are in constant conversation with the environment. We depend on the environment for energy. That's why foraging is such a fundamental problem for. Animals, fundamentally, what defines them is that they move. They move in order to forage, find food, safety-

Paul Middlebrooks
They can move.

Joe Monaco
-and shelter.

Paul Middlebrooks
They can find more food so they can move more.

Joe Monaco
It turns out ecologically there's a lot of niches that open up if you can move through the environment, otherwise you're a coral or you're a sea squirt attached to a rock somewhere and you're a filter feeder.

Paul Middlebrooks
People will take umbrage to the idea that those don't move, by the way, some people will, even plants, but yes, I get the point.

Joe Monaco
There are sessile animals. I did not mean to offend the sessile organism community, but movement is fundamental to all of this. It's that dynamical coupling at those informational boundaries that allow the goals to basically stream against the incoming sensory inputs, along that kind of hierarchy of both perception, ascending, and the hierarchy of drives and movement and behavior control, descending.

That was our perspective that you can think of. It's related to the predictive processing framework like Karl Friston's unifying theory, where he sees the brain as a distributed internally generated feedback model and you're canceling out prediction errors as they ascend with the top-down expectancies, and then there's trade-offs that are governed by his conception of free energy.

That unifying theoretical framework has had trouble gaining traction, making direct predictions about what people should be doing in neuroscience. What type of experiment should you design to figure out, "Oh, this particular function, it operates this way within the predictive processing framework in a way that's distinguishable from some other framework"?

I'm just trying to step back and not put a name on things, but fundamentally organisms construct meaning through managing basically the entropy at this interface.

Maybe that's prediction error, maybe it's something in some other quantity, but you need to manage entropy. Fundamentally, that's what you're doing. Thermodynamically we're far from equilibrium. That's the whole game.

Paul Middlebrooks
That's the control theory aspect of it? Is that where that comes in as well? What I was going to say originally is, injecting meaning back into

neuroscience is not the default. The default in neuroscience is this reductionist, "Brain is a computer," and then we can go from sensation back to-- we can disregard goals and meaning and purpose. That's been the default position of most neuroscience for a long time, although in the paper and elsewhere it's pointed out that early cyberneticist movement, was more about control. I guess that's why I'm asking, is that where the control theory aspect comes in?

Joe Monaco
In this paper, which again was our perspective, we found the most simpatico framework out there to be what's called Perceptual Control Theory. This is an alternative branch of cybernetics essentially from the '50s and '60s that was brought or initiated by Bill Powers in the '60s and '70s.

Paul Middlebrooks
Henry Yin, whom you cite in the paper and he's been on this podcast like many of the references in your paper, which also-

Joe Monaco
You influenced us too, Paul, your podcast.

Paul Middlebrooks
Just, maybe by having people through the podcast. Thanks. That's awesome. It's awesome to see so many references in a paper. It's like, "Oh, that person has been on the podcast. That's the person's [unintelligible 01:07:35]" I just mentioned Henry Yin, and he was saying one of the problems with early cyberneticist research and control theory in general actually is that, the reference signal of a machine is external to the machine, whereas we have internal reference signals that we're trying to control to match for those reference signals. That's a fundamental difference and that's what neuroscience is missing. I don't even remember my question, but that is what you speak to in the paper as well.

Joe Monaco
Right. Henry's written a couple book chapters, with this perspective and they're bomb-throwing chapters, in a way. I think it's helpful to have strong opinions out there, because it really makes you think, "Okay, this sounds interesting. It's provocative, but where does it go wrong? Does it go wrong?" That's reading that. I'm not a motor systems person, but I have passing familiarity with the motor control paradigms, those theories, coming out of the '90s and 2000s, optimal feedback control conceptions of motor commands, the theories and models about efference copy and corollary discharge systems.

All the traditional motor control systems or frameworks were based on building up more and more detailed and refined internal models, basically, forward models, to predict the consequences of action and movement, and then using that to evaluate different commands and behaviors, and then putting that in this much larger, much more complicated control loop.

PCT or perceptual control theory is appealing because in a way it reverses it. It says, "No, it only matters is that, you're only making comparisons at each level in this hierarchy because you're making direct perceptual comparisons." If you have a direct perceptual goal at the highest level, then those reference points come down or compared to the ascending perceptual input. Then that descending reference to the next level gives you what you need.

Paul Middlebrooks
Eventually it filters out through your muscle actuators and your movements in the world.

Joe Monaco
It's not eventual because it's all simultaneous.

Paul Middlebrooks
It's all simultaneous, right.

Joe Monaco
It's a staged flow of control signals, the sensory and control signals across the hierarchy.

Paul Middlebrooks
I'm individual in terms of time, where one signal flow starts, it takes time to propagate, not that there's a central organizer that says, "Go," from nothingness, because yes, part of what you push for also is this consideration that perception and action cycles are continuous flows.

Joe Monaco
Right. That phrasing ties into this conception of everything as being linear, input, output. If you have a cycle, it's like, first you're at this step and then you're at the next step. You're at sensation, then you're at cognition, then you're at motor commands, and then you're at behavior. Then that changes the pose, the orientation of the organism with respect to the environment, which changes the sensory inputs, and now you're back at the beginning of the cycle.

If you have a very complex computational forward model in that control loop, then you have to imagine that the delay of computation is now a delay in your control loop. From a control theory perspective, from a control engineering perspective, the more delays you have, the weaker your control is. One of the, I think, the most dispositive or fascinating properties of animal behavior is that it's really good, it's really resilient. Animals can accomplish their goals.

Paul Middlebrooks
Not always, but they're very good at it.

Joe Monaco
Right. Compared to the variability of the observed behavior. The robustness of accomplishing goals far outstrips the actual movement. It seems like, wait a second, how is any of this possible? How is it possible for a rat to make its way through this very complicated burrow with basically no light and only a small set of sparse cues, but it can navigate that burrow really well. It comes down to, maybe you don't have all this complex computation going on in the loop.

This is a conversation that I've tried to have over the years with people, working in motor systems and they think, either I have the wrong idea or Henry has the wrong idea or actually, all of their theories already encompass this idea, "Don't worry about it." It seems to me, from my personal scientific opinion, this is an open question of forward versus inverse models, ascending prediction error comparisons versus perceptual reference point comparisons up and down the hierarchy.

That's some of the discussion that I wanted to open up at the workshop. From BRAIN's perspective, I'm not going to impose my views on this, but I see that that's an important conversation. I think that will open up a lot of potential opportunities for driving theory forward. Data is a part of it. The reductionism is, obviously molecularly characterized cell type atlases of whole brains, very fine grain Connectomics datasets, the FlyWire dataset that was just released, which was brain supported on a number of grants, and there's more to come.

We launched the Brain Connects program last year, and we'll start seeing data from those projects in the next few years. Lots of exciting stuff to come, but that's obviously from a very reductive approach to neuroscience, break things down and just so we can see everything that's there. I think that does need to be in the loop with this more holistic way of thinking. I think that's where there was a lot of talk about digital twins, multi-scale biophysical modeling, and then thinking about different ways of putting this in the loop with behavioral neuroscience and different ways of understanding that.

Grace Hwang
That was definitely a surprise for me at the NeuroAI workshop.

Paul Middlebrooks
What's that?

Grace Hwang
That there was so many talks about digital twins. Even in sessions 3 and 4, it seems like the community is really ready and really want digital twinning in their respective research areas.

Paul Middlebrooks
What's a digital twin? Why do we want one?

Grace Hwang
That's another definition question.

Paul Middlebrooks
You don't have to define it, but roughly.

Grace Hwang
There actually is a definition that The National Academies of Sciences, Engineering, and Medicine put out. I have it in front of me. I'll read it to you because a digital twin is a set of virtual information constructs that mimics the structure, context, and behavior of a natural engineered or social system or system of systems, is dynamically updated with data from its physical twin, has a predictive capability and forms decision that realize value. The bidirectional interaction between the virtual and the physical is central to the digital twin. That's it. That's the official definition. I think people have their own definitions.

Paul Middlebrooks
I'm sorry. I asked for the definition. Just kidding.

Grace Hwang
They were using their own definition, which is a subset of this official definition. This actually was a question that came in in the email. I was surprised to hear so much digital twin talk. I think that's potentially a new, exciting area that the NeuralAI workshop participants can continue to engage in.

Joe Monaco
I think there's an important continuum there as well. We heard some of that discussion in the first session of the workshop between neural foundation models on one side and digital twins on the other. These are both large scale ways of using large scale neural and behavioral data, but they have different goals. Neural foundation model is like foundation models in AI, where you want to have a base model from which you can generalize to downstream tasks and application-specific domains or to answer particular questions.

Digital twins, parsing the definition that Grace just gave, it is really more focused on using lots and lots of data to make very clear predictions about a particular individual system. It's individualized, or in the health context, it can be personalized.

Paul Middlebrooks
You can test hypotheses about the natural system using the digital twin.

Joe Monaco
Right. The idea is it evolves with the system you're studying. If you have a digital twin of, let's say a mouse, and then the mouse is in a particular task, you can be running the digital twin model in silico in parallel with the actual experiments. Now you've got, the title of Patrick Mendoza's talk, "Closing the loop with virtual neuroscience." Closed loop neuroscience, we've got an in-silico ghost or simulation, essentially of the actual animal in the experiment. Then you can do very fine grained real-time predictions, modulation. The idea is that should be a very powerful approach.

Paul Middlebrooks
You can imagine all sorts of things like tracking through the lifetime changes and through development and the lifetime, and not necessarily in humans, because that's really longitudinal, but you can do it on a faster timescale with something like--

Grace Hwang
It actually came up during Kai Miller's presentation that he, as a functional neurosurgeon, would like to have digital twinning in the future.

Paul Middlebrooks
Sure. I just mean, the particular idea of tracking over the lifetime, but if he had that in his surgical suite, then he could test things very quickly, and then decide whether or not to implement some surgical technique.

Grace Hwang
Exactly. That was really a conversation I wasn't expecting from session 4, but it was very insightful.

Paul Middlebrooks
Those are the fun things at workshops when something like that surprises you.

Grace Hwang
Yes. Also, the combination of Chris Rozell and Kai Miller on the same discussion, both are very clinically savvy, and taking opposing views for neuromorphic, was exciting session 4. When Joe brought up the concept of a hardware lottery in the concepts of Shannon information theory, it reminded me how neurotech certainly also could suffer from the hardware lottery, given how hard it is to get devices approved. We're often just stuck with what's approved and not what's necessarily the best.

Paul Middlebrooks
I'm aware of our time and I want to make sure that we talk about that. One of the interesting and maybe surprising things, because the paper does so much, is that you end up arguing for a base layer of neural computation, like we talked about before. You don't have to define it, but what is a base layer, roughly, of neural computation? Then, why do we need one? Why do we need to determine what the base layer is?

Joe Monaco
This is not a term that I coined, it's just I referred to it as the base layer in this paper, but that came out of thinking-- I was reading some philosophy, philosophy of computation, what are the different types of computation? How do you do computation in physical systems, in dynamical systems? Is a system of ODEs, differential equations, and you just evolve them forward with Runge-Kutta or whatever your algorithm, is, can that do computation? Can continuous dynamics compute?

There's a lot of really interesting questions around computation, and brains are a particular kind of physical, dynamical, material, chemical system. Again, falling into the default concept of the neurocentric framework for understanding brains, is maybe leading us astray. Obviously, throughout biology, cells are super important. Neurons are super important to brains, but also neurons are not the only cells in brains, as we know, there's all sorts of glia, astrocytes, oligodendrocytes, microglia, which have important roles in structural plasticity.

Looking beyond, digital computing has gotten us into thinking about computational systems as, "Okay, there has to be a transistor." There's some unitary element that's the lowest level thing. If you think about a silicon disk with chips that have been burned into it in your photolithographic process, all it is a material carving in silicon and other types of materials. That's the base layer. The transistors in the CPUs that we use in all of our computers and phones right now, that's the base layer of computation for digital computing on conventional CMOS processes. Is our brains just like that? Are cells like transistors? Is that it?

Paul Middlebrooks
No, cells are binary event action potential generators. I think McCulloch and Pitts were right, and nothing has changed, and so we should still consider them that. Right?

Joe Monaco
No. McCulloch and Pitts, that was fundamentally wrong. It's not a binary signal. It's an event. A spike is an all-or-none event. Some people will say, "Oh, okay, so it's spike timing. What's the timestamp on that spike or versus that spike? Then that'll tell us everything." No, because you don't need an absolute timestamp. You need to know what the role of that one spike is in this dynamical system, because that spike is propagating to downstream neurons, and at some point you hit recurrent connections and it feeds back, and then you go up a layer to the next higher level of cortex or whatnot.

Everything is causal, dynamical, interconnected, so you can't just say, "It's a one or a zero." Obviously that inspired von Neumann, and it's an amazing insight because that's why we have digital computing technology now, but it's not how brains work, as Walter Freeman pointed out.

Grace Hwang
As we also heard from Yiota, there's all this dendritic computing that's happening at the dendrites. There's just, I think, so much more richness that we are now aware of that wasn't available to McCulloch and Pitt.

Paul Middlebrooks
That's true, but now, Grace, you just went down a level physically from the point neuron to dendritic computing, which would make Panayiota very happy, but you guys want to go up a level and talk about the role of oscillations in this dynamical coupling. I also found myself wondering, "I don't know how engineers think about oscillation."

Grace Hwang
The other thing is traveling waves that we really didn't talk much about at this workshop.

Paul Middlebrooks
That's true, which I'm surprised Terry didn't bring up because he likes to talk about traveling waves.

Grace Hwang
Yes. He actually sees oscillations and traveling waves to be one in the same for many parts.

Paul Middlebrooks
How would they not be? An oscillation has some spatial--

Joe Monaco
Oscillations repeat. You can have a wave that travels that's not being generated by an oscillating generation process.

Paul Middlebrooks
Okay. That's fine. You can separate them. That's like having one wave in the ocean, which isn't an odd-- Once one wave leaves, the other wave has to-- I don't know. They are intertwined in general, but I could see you could just do one..

Joe Monaco
Right. Space and time are coupled in the brain. There's a mathematical formalism called hierarchy theory, which is basically if you have oscillations at a fast frequency, you can only maintain coherence over a small amount of space. If you have oscillations at a slow frequency, you expand the region of space over which you can maintain coherence with that clock, with that slower oscillation.

At least mammalian brains have a really well-preserved set of neural oscillations in different parts of the brain at different times that interplay with each other in different ways, at base frequencies, which are at these really interesting incommensurate fractions between each other. It's almost like nature needed to find half a dozen different frequency bands that didn't interfere with each other or minimally interfered with each other, because then you can have a theta and you can have a gamma, and you can nest seven of those gamma cycles in one theta cycle, and then that becomes an interesting packet of coordination.

It's not the spike timing. It doesn't matter that, oh, neuron A fired at time T-zero within theta cycle, whatever, X. It's not that absolute index of time that matters. It's the fact that, oh, you've got this packet of activity that's carved out by this sequence of gamma oscillations or gamma cycles within this theta cycle-- net theta cycles, within this larger set of slower rhythms.

It's hard to ignore these laws almost. This relationship between space and time, we have these conserved oscillations and they do govern the timing of spikes, the activity of neurons. There's this feedback loop that goes up a level to a collective behavior, like an oscillation, and then through haptic effects or through just other modulations, they entrain and feed back into causal mechanisms at the cellular level.

Paul Middlebrooks
Then I want to ask what the base layer is, what the proposal is for the base layer. Maybe even before that, I could list off the three requirements that you posit for a base layer to be a useful computing layer. By the way, I should also say that the proposal is a non-reductionist mechanistic account of neural computation, which is an interesting thing itself.

The requirements that you state for a base layer, and I'm reading directly from the paper here, is, one, that they encompass a macro scale hierarchical control structure, so over which it implements comparator error and output functions. That's the control theory aspect, part of it. Two, to adaptively control access to internal and external information flows generated by physical embodiment and situated embedding in a causal environment. That's the almost ecological psychology interaction between these continuous flows.

Then three, support discrete neurodynamical states and adaptive high dimensional state transitions across timescales of neural circuit feedback, and then you list some specific kinds of timescales. I suppose that links into the reason why oscillations are important, these nested structures of oscillations, the spiking information carried within those oscillations, and how they're interacting across different timescales and structures and flows. All right, mouthful.

Joe Monaco
That's an ambitious framework.

Paul Middlebrooks
It's super ambitious. It is super ambitious. This is a 30-year BRAIN initiative. That's the other thing, is reading this paper, it's like, oh, this is like a whole textbook or a whole four special issues in some journal condensed into one thing. It's ambitious.

Joe Monaco
I should make clear, this document has nothing to do with the BRAIN initiative or views, perspectives, priorities, plans, or any of that.

Paul Middlebrooks
You come out thinking, "Oh my God, where do you begin? How do you start?" So much to do.

Joe Monaco
Okay. I wrote this a couple of years ago and I'm not sure I even remember the three criteria that you just listed. I'm looking at the paragraph now, but this all came out of, again, a frustration and just wondering, well, what if the whole neurocentric paradigm is wrong? Rafa Yuste had a review or perspective paper from a number of years ago, saying that a network centric paradigm is what we-

Paul Middlebrooks
He went from the neuron doctrine saying, "That's old, didn't work." Now we're in a population doctrine era he was advocating for, which is where the field is right now, it's all population dynamics manifold.

Joe Monaco
John Hopfield just got the Nobel prize. There has been movement in that direction, as we've been-- Fundamentally there's a lot of technological determinism here. The better our tools get, and BRAIN Initiative was certainly behind a lot of that, the more neurons we are going to record at better fidelity, at more throughput, the more you can see. We're getting beyond single neurons because we can record millions of them now.

We need to understand what's happening still. Now the focus is on low dimensional representations. What's a low dimensional representation? Essentially, it's an attractor. It's a small subset of a high dimensional space that's been carved out. You can say, "Effectively, these million neurons I'm looking at, the state of the system is somewhere on this two, three, four dimensional-- maybe four, but a low dimensional system," and you can basically understand it.

If you can map the axes, the dimensions of this low dimensional manifold onto task requirements and constraints, then it's like, "Oh, that's explainable. I know what's happening in this neural system." Then there's questions about, well, the whole brain isn't a single attractor. It's not one giant hot field network.

Paul Middlebrooks
I don't know. Some people might disagree. Go ahead.

Joe Monaco
I'm saying, I think there's attractor-like dynamics everywhere, but it's complex and heterogeneous, modular to a degree, and governed transiently dynamically by lots of things at different levels of organization, including things like oscillations, and things like traveling waves. You have coherent organization at time, coherent organization in space. A lot of it's quasi-hierarchically organized.

People will think about natural intelligence and its agility, its flexibility. Your animals are optimizing multiple objectives simultaneously. How do you do that? You do that by flicking on and off different subnetworks at different scales adaptively in the right way. It's like, I'm trying to do three things at once. I want to get food in an hour. I'm trying to wrap up the current sentence I'm speaking, et cetera. You've got these multiple goals in mind. How do you do that? You need to activate different attractors in different ways in a complimentary pattern to achieve the goals of the organism.

That's where the spatial temporal organization comes in, but it's governing this heterarchy, maybe, of attractors or quasi attractors. That's where I went to in this paper is thinking about something-- I regret that terminology, but I call them tokens or causal tokens. I was just trying to think of, how do we think of an attractor or not as this like, "Oh my gosh," this huge task manifold? The animal goes into a maze--

Paul Middlebrooks
This manifold needs to come up and then it needs to go along the manifold.

Joe Monaco
Then all the activity projects onto that manifold. You can figure it all out. If there's a go-no-go, then the selection vector rotates through it, and then boom, the behavior happens and you have Dave Cecillo's work and going back to that 2013 paper, which I think is a great idea. Something is there, but the interpretation of what that means. What is a selection vector versus the command vector or whatever the other space was there? How do we think about communication subspaces? How do they come on and off adaptively in service of goals?

My idea was, causal tokens are these little quasi attractors and they can exist at different scales. Quasi attractor, because you don't get stuck there. The system doesn't get stuck there. You need destabilization, if you fall into an equilibrium state that you can't get out of you're dead. As we know, cognition keeps moving, it's always moving. You're always finding little attractors and then being bumped out of them. There's a competitive process, maybe, but it's just trying to think of, what is the base unit of computation if that's what's happening?

Paul Middlebrooks
What is the base unit of computation that you're advocating for, the base layer?

Joe Monaco
I went back to Hebb, Hebb and Karl Lashley. I should say, there's a great review and perspective put together by Drew Mara and Lynn Nadel from a few years ago. It's cited in the paper, where they reconceptualize what Hebb was talking about. Hebb and his student, Karl Lashley really thought deeply for many decades about what it meant for networks of neurons to be connected to each other, "What are they doing?" Looking at persistent activity as well through self-reinforcing patterns of activity.

Basically that's the causal base layer that I was interested in, if you have a super neuronal group, a group of multiple neurons. The thing is there's a loop of reverberant activity going through those synapses within that interconnected loop at some level, and it's self-stabilizing. You can ignite it essentially. Then the connections are such that the non-linearities all line up and you get self-reinforcing, self-supporting, self-sustaining activity.

That's the basis. That's the base computational unit that I was speculatively putting forward in this paper. The nice thing about that is you can sprout loops off of that, There's always another connection. You can always reconsolidate those connections in a different way. Maybe structural plasticity is-- maybe there's a side loop that's sub-threshold, but then one thing happens experientially for the animal, and then that connection gets twisted up a little bit.

Now the non-linearities line up in a different way, and the loop expands. Now you're at a higher scale, causal token or whatever I was calling it, this quasi attractor now means something because it's incorporated this new correlation from the environment. Maybe that's a control signal or control parameter that was updated within the hierarchy or something like that, but it's a self-sustaining bit of activity governed by all the spatial temporal structure that we were talking about with oscillations.

Paul Middlebrooks
It's fun to see you light up like that when you're describing it. You look excited and sounded excited talking about it.

Joe Monaco
I think you reactivated my neurodynamical state when I was writing the paper. [chuckles]

Paul Middlebrooks
All right. We have just a few more minutes. Grace, it looked like you might've wanted to jump in there also, or no.

Grace Hwang
I just wonder if any of this could be measured experimentally.

Paul Middlebrooks
Oh, good God. We don't have another two hours, or was that a plant question?

Grace Hwang
It's a quick question for Joe, because the idea of measuring every synapse came up yesterday.

Paul Middlebrooks
Well, you could do that. That's very straightforward assuming you had the right technology. That's a great question, Grace.

Joe Monaco
You brought it up earlier, Paul, as well. This does tie in directly. I'm not saying my idea here that I just went through is absolutely right. It's just where I went is like, this seems like the most likely useful framework for thinking about it. If this is true, that the actual particular value, the precise synaptic weight of any given synapse is almost immaterial, it doesn't matter.

Paul Middlebrooks
That's why you're bringing that up?

Joe Monaco
Yes.

Paul Middlebrooks
I questioned that too. I actually pushed back in that discussion. What did I say? Something about, just measuring something doesn't give you the theoretical some blah, blah, blah. I can't remember what I said, but then I got pushed back for saying that. I was like, "Oh, I didn't realize that what I was saying was even controversial." Thanks for bringing that up again.

Joe Monaco
Right. In the paper, we say that, if this hypothesis were true, this, what matters is self-sustaining little clumps of neurons that can expand outward and adaptively, then this is completely antithetical to what we see in AI models based on artificial neural networks, where everything we care about is, when you distribute an AI model or a transformer LLM, it's a binary blob full of very precise weights. Then the whole game is to see how far you can quantize those weights down and still preserve the functionality. You can put these things on phones and home computers, and all of that.

Everything that matters is the weights, it's all in the weights and nothing else, the biases too. If this hypothesis is right, then that doesn't matter. To understand the brain, we don't actually need to go around and measure every synapse because they're wildly fluctuating anyway. It's highly volatile.

Paul Middlebrooks
Some people would argue against that because I've received pushback saying that exact same thing that actually they're largely quite stable. I guess we won't know until we measure every single synaptic string.

Joe Monaco
If we do have the self-sustaining quasi attractors, then you're going to have synaptic loops which self-sustain and do maintain strong correlations over time that persist about relative synaptic weights. You would expect that. It's not the weights that matter. It's only lining up the right set of nonlinearities so that that group fires in the way that it does, and interacts with other tokens or other quasi-attractors.

Grace Hwang
Sounds like dynamics was what's largely-

Joe Monaco
Dynamics was switching.

Paul Middlebrooks
In a hierarchical and heterarchical fashion.

Joe Monaco
Just heterarchical. I think it encompasses all of it.

Paul Middlebrooks
Guys, I have to go here in a minute. This is one of those papers that I'm going to revisit and then feel guilty that I'm not reading every reference in that stack that grows ever so larger of what we're supposed to be reading all the time. It is just so rich, and I'm glad to point people to it. Congrats again on running a great workshop, and I think a successful workshop. I really hope that you guys get some rest. I know you have to take in everything now and then reflect, but hopefully that's a little bit more relaxing a process, and then you can take a little vacation.

Grace Hwang
Yes. Thank you so much, Paul, for coming to the various pre-coordination meetings. I was so impressed at how hard everyone worked. We got multiple abstracts, multiple versions of presentations. It was amazing that we had everybody share their files on the NIH box, and you could see how people were changing their presentations in response to each other. Thank you so much to you and everyone else who really made this a great workshop.

Paul Middlebrooks
Oh, it was great.

Joe Monaco
Thanks for having us.

Paul Middlebrooks
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 advanced research. Visit thetransmitter.org to explore the latest neuroscience news and perspectives, written by journalists and scientists. If you value "Brain Inspired", support it through Patreon to access full-length episodes, join our discord community, and even influence who I invite to the podcast. Go to braininspired.co to learn more.

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

[music]

Subscribe to "Brain Inspired" to receive alerts every time a new podcast episode is released.