

UCL Brain Stories Episode 18 - Brain Stories Live

📅 Wed, Nov 08, 2023 10:21AM ⏱ 1:35:47

SUMMARY KEYWORDS

ai, work, neuroscience, brain, study, models, problem, people, fact, machine learning, data, learn, understand, ucl, neuro, question, caswell, metacognition, architecture, decision



00:02

Welcome to brain stories live. I'm Castleberry This is Lena ray.



00:07

So this is new. This is new to us is to you who's seen or heard our podcast?



00:15

reasonable numbers. Okay. So in the podcast and tonight, we've got pretty much the same format. First of all, we talk about the science or people are doing what I guess are doing, what they're researching. And then we talk about what makes them tick the journey that brought them to where they are, and we're going to be doing the same thing tonight.



00:31

By way of audience research, who we got here we've got anyone any undergraduates here, show of hands, no, undergrad, I saw you there. No.



00:41

Any any postgraduates? few of those? Anyone who's not at the University at all, feel it. Okay. Good. Good mix, right. Fantastic. Now, now we know who we're aiming for. So Celina for a moment. Well, welcome, everybody. And thank you for coming to join us for this experiment. And so if Caswell hinted at this is there's a series of firsts this evening. It's our first ever event of recording where we've had more than one guests. So we will welcome a panel up to the stage in a moment. And it's also our first ever live recording, hopefully not our last, please be nice to us, so that we come back and do this again. And the reason we decided to do a live recording is because we wanted to meet some of our listeners, but also because this year is the 50th anniversary of the UCL neuroscience domain. And so the UCL neuroscience domain is a network across UCL that tries to connect scientists working in different disciplines bring us

011' Together so that we can work more effectively, but also to do outreach events so that we can showcase our work more broadly. And so you're all very welcome tonight, and we're really looking forward to sharing some of our amazing scientists with you. And we will, the way the evening will work is we will do this in two parts. The in a second, we'll welcome our panel up. And in the first part, we will discuss their research where they think their field is going and what they're excited about, will then take a short interval in about 15 minutes time. And after the break, we'll get a little bit more into the start the scientists story. So why did they become interested in their research areas? What did they study? What key career decisions did they make that brought them here? We really want this to be interactive. So you're very welcome to ask questions throughout, just give one of us a wave and we will come to you. And if people have questions, but they don't really want to stick their hand up, I've got as you can see, in my hand, I'm clutching a few pens, and I've got a little bit of paper and I'll leave them down at the front. So you can write things down in the interval, and then we'll ask them at the break. So I'll hand back to Caswell, we have an organising theme today. So the topic, the thing that unites three guests, is neuro AI. And so what I'm going to do is attempt to define what that is now, I'm sure when I guess Come on, they'll tell me I'm wrong. And that it's nothing like that. So here's my working definition of neuro AI. So it's a portmanteau of neuroscience and AI. The point being that these two fields sort of share a common lineage. So there's sort of a natural connection between them. And what that means is that various points, there's been sort of exchange of information. So some of the ideas that are sort of driving AI development now been borrowed from neuroscience, the ideas of neural networks, the ideas of reinforcement, learning the way animals learn, things like convolutional networks, which is how we think some of the visual system works. Has that's been happening for a while. But what's increasingly been happening over the last sort of five or 10 years is, neuroscientists have got kind of wise to this and have been taking lessons from AI. And so neuroscientists are increasingly using machine learning tools to deal with their data, maybe to diagnose diseases, but also using these sorts of neural machine learning models as models of the brain telling us, you know, if there's a problem, how ought this be solved? And then we go and look in the brain to see if the brains respond like the machine learning models. So we're going to find out whether the guests agree, you've heard enough from me already. So I'm going to say the magic word. Let's bring on the guests. And hopefully three eager researchers will bust through the curtain.



04:28

Fantastic.



04:33

Well, thank you for joining us this evening. And we're very excited to have the discussion with

Sainsbury Welcome Centre. I sort of said between systems and computational neuroscience, and yeah.



05:19

Benedetta Hello. Hi, I'm Benedetto de Martino. And I'm a cognitive neuroscientist working at the ICN. That is the author of cognitive neuroscience, not much fun to see you in.



05:32

So maybe we can start just by giving the floor to each of you for a few minutes to expand, maybe introduce your research in in kind of general terms. What are you working on? What sort of techniques are you using? And why is this important? And so if we go in reverse order this time, Benedetto would you like to start? Sure, I do. Mostly my work is in human cognitive neuroscience. And my hi all is strange. I mean, I come from molecular biology, but I haven't seen like a little mouse for a long time besides, this is for the occasional one.



06:11

So we do use neuroimaging methods like fMRI, E, G, M, eg, but mostly my lab is also interested in computational modelling, and things that are like crossroad between neuroscience, economics and machine learning.



06:31

I've just completed a PhD and during my PhD I was studying how different brain areas especially like corticostriatal circuitry, coordinate the acrohy to dothy,,



07:48

that's might be one for us to pick.



07:52

So, I work in this kind of subfield called Computational psychiatry. And there's a mixture of things that goes on there, I've done bits of all of them. So sometimes using computational models of cognition, to work out how my brain processes proceed, just like these guys do, but then also to think how they might go wrong. So how perception might become a hallucination, for example. And then also biophysical models of imaging data to try and infer neurobiology and neurobiological properties of the brain,



08:34

which is, so to link to things like drug targets, and that kind of stuff. And then lastly, kind of more machine learning based methods to analyse very large datasets, for example, comparing patients and controls and this kind of thing. So I just wanted to pick up on some of the things you were saying, Rick, so would you say it's true that we're already at a point where machine learning models are useful in the clinic? Or is that is that something that's yet to come? I'd say it's very much depends on which clinic so so in some areas, they're really kind of ready to go, it looks like from how well they work. So ophthalmology, there was a really landmark collaboration between DeepMind and the Institute of Ophthalmology down the road, where they had a million labelled retinal image scans. And this network learned how to diagnose diabetic eye disease, hypertensive disease, and tell the age and sex of the person in a way that they don't know how it can do.



09:38

But when the problem is



09:40

brain data, like structural MRI data or functional MRI data is such.



09:47

So many more data points in there. There's like 100,000 data points instead of the number of pixels in a retinal image.



09:56

And we don't have a million labelled scans we have like maybe



10:00

1



12:28

other type of engineering stick approach in which you as an engineer, you



12:33

come into engineering what the things has to do er, your engineering that the learning architecture, so your practical your engineering, not towards, in the case that



12:46

it was talking about, they didn't engineering something specifically for the AI is engineering, some architecture that it doesn't care if it's an AI, images, whatever, like a scan from an airport. And then through these like massive training, normal is a training through labelling. So this was the reason why it was saying, the problem is we do not have enough labels, can the machine start to understand this association, this relation? Now this is something quite a bit different from the way you learn and small children learn. So in a way, the big challenge Hi there is getting this massive data or understanding something's more about human learning habits, small children, I mean, they're not so small anymore. For adult they're also small, but when they learned were horses, they didn't need the million horses a billion horses to detect well horses, you show them for horses, and since then, they can pretty much detect where horses so is a learning architecture. Now neuro AI specifically



14:02

is clearly as we joke not super clear what it is that everybody sees in one way, but you can imagine both way like is the use of AI. In neuroscience. You know, Caswell, for example, has done himself some work he should be talking about that rather in interview, in which he has like, he trained and that machine to just detect the like movement of mice and things and, and doing like some work that was very tedious and very long for human to do and not very well. And then the most ambitious part of neuro high that unfortunately, is one as under delivered, in my opinion, is the contribution of neuroscience to do exactly what I told you to have a computer ^{ere} t v rse o nherince thververold

a quess que~~s~~sess f à o



17:52

training set data from all over the world from all kinds of different scanners if it's going to be exported, because obviously, the hope is we can use this technology in places that a resource poor, but obviously, the absolute worst thing is gonna happen if it's if it doesn't work in those in those places. So.



18:10

So there's a, there's a really good point, actually, I want to come back to those sort of questions of generalizability. And sort of how equitable we can use these things. But just want to finish with one finish up on one of the points that Benedetto mentioned, which was this sort of disappointment that the dream of neuro AI to sort of paraphrase was that the information would go both ways. And at moment, neuroscience is doing pretty well. But some people would claim and indeed, when I put on your AI meetings, the lack of machine learning, people would seem to indicate that actually, the information isn't going back the other way, so much. And then, an example of sort of a story that was told me was, you know, it's a bit like birds and planes and flying. If you want to, if you want to build a flying machine, it's great to see birds at the beginning, because they tell you that flying exists. But if you didn't spend 50 years trying to build something with floppy feathery wings, you're not going to get very far you need to give up and try something different. And the implication being that maybe we've taken the inspiration from the brain, but actually, we shouldn't be trying to copy it too closely to achieve what we want. Do I just wondered whether you've got any sort of thoughts or thoughts about that, whether whether you're disappointed or not, basically, I'm more optimistic. I think from a basic science perspective, I think the fact of the kind of techniques which have been developed in AI of late have helped us make better models of like, how neural responses look in the brain. But I think now is the time that we can take that inspiration from neuroscience back to AI models. So AI models tend to be really brittle as you were saying that they don't generalise well, they pick up on random features, whereas brains tend to learn functions which are very smooth and generalise well and are not brittle in the same 'sense' as the 'tar' d' r' d' th



gent i4y So qg

function starts to degrade before others do and then, and then you become much more reliant on your expectations of that model, rather than the information that's coming in, and then you that can generate hallucinations. But you only really think about these properties when you have to build these networks that can recognise speech or recognise facial things. And you realise what is so important about the differences of it also are like one of the major departures in transformer architectures is that earlier, when people were trying to predict speech or like sequences and time, they were using recurrent neural networks, which were probably like, closest thing you can get to what brains look like in like artificial neural networks. And the advance in Transformers is that they got rid of that component and said, We do not need that we can just do with these matrix multiplications. And these attention heads. So it's like a major departure from what looks like brain like, artificial units to like a completely different thing. All the people are trying to map that onto an RNN recurrent neural network, but it's like that mapping is still incomplete and debated.



26:38

Great, thank you. So I want to move on. Now to go back to the points you're mentioning earlier and pick up on those about the sort of

motivation. We have a society in which intention intentionality matter. Because if I say I did, because I wanted to kill that guy is very different from idea to just. So this is an issue that is actually incredibly important. In human brain architecture, we have this module that I happen to study quite a bit, that is called meta cognition, and is the fact that we are able to introspect what we do and verbally report. This module hasn't been very much of interest fB ê ú



30:34

They've often found that in the lab, they outperform the humans, but in the clinic, they're no different. And when they look into it, the reason one of the reasons they're no different is because when they give a recommendation the doctor would not have made, the doctor just overrules it, and ignores it. And so really, the the this reasoning, this big ability to present its reasoning is probably going to be the only way that actually gets people to do the recommendation to follow it. Because it may not actually make sense rationally. But practically, they might not be useful unless they can do that. So my sort of take home, I guess from that little bit of the discussion is we are not in danger yet of having our GPS replaced by a or do we say chat GP?



31:50

I get that wrong.



31:54

Yet.



31:56

But maybe that can lead us into a bit of a discussion of, I guess, overall, what we're seeing is there's areas of huge potential, but it's still a really early field where there are huge challenges to overcome. So I guess from each of you, I'd really like to hear a minute or two about what you think the next 15 years holds for Neuro AI in your particular areas. What are the things that you're most optimistic and excited will happen?



32:26

Starting with me? Yeah. Okay.



32:28

So I think I'm most excited about neuroscience has worked with very like, sensory and motor systems. And that's where a lot of the research has been sort of confined, and we use simple models explain that. And that was all well and good. But now we're moving into this like, era in which we're trying to understand more complex decision making and behaviour in general. And I think currently, AI is the only field which has models for that. And I'm very excited about marrying those ideas with like, what we know happens in the brain and what we know what kind of representations exist in the brain and how brains do those things, and what kind of architectures exist in the brain. And I'm kind of excited to see that kind of, sort of,



33:08

... of things coming together so that we can make sense of how complex behaviour is

sort of things coming together so that we can make sense of how complex behaviour is produced by brains. Yeah. Okay. So it already has impacted in a positive way, my life persona, that before I told my wife isn't native English speaker and I always said was, when I had to write an important document to proofread for me, now I have GDP.



tonight, the decision has made you have you made a good the right decision is not going to hands up with a score of how you how many points you made, now means that you constructed the value of reality around the view. Now the question is machine are really far away from that. And if you want to make them be more like flex, while acumen there might neB therlik



38:32

it just is a much richer, deeper understanding of what they're about than any kind of very biological reductionist type of understanding because it because when, when, when you view people as having a model of the world,



38:47

you can't avoid all of their background, all of their childhood or their environment, all of their social influences, and, and the biology that forms it. And it's not just the, you know, it's not just the serotonin and whatever, you know, it's the whole thing. And so, understanding that is, is super interesting, but I think in the next 15 years, it'll take quite some time to get that I think, in the immediate future.



39:15

The most progress is going to come from these kind of black box machine learning type approaches that just say, give this person this antidepressant or this person, this antidepressant, and it's not totally clear why but it seems to work.



39:28

But even that will be great for now.



39:32

I think there's like I should have told you, they're gonna we're gonna get the audience to rate you afterwards, there is going to



39:39

so it seems like actually, what you've all done to a certain extent is draw a distinction between applying applying these machine learning models as a tool. So you know, diagnosing a thing or deciding whether that's a label of a cat or a snake. And I think probably you'd all agree that's already being useful in neurosciences.



40:00

Don't be useful in the clinic. My lab use things like deep lab coat, which is a Bayesian automatic way of labelling animals and boxes and things like that. And I think that's unambiguous. What's, what's kind of interesting is the way you've talked about using using machine learning models to learn things about the brain, not just to like sort out your data. But you know, like, can I? Can I look at how the machine learning models solve this problem? Or how it compresses data in RNN or a transformer, and then make some inference about how the brain does that. And I

in tn

Welcome back, everyone, first of all, and welcome back to our panel. And we will start with some questions audience questions thanks for for filling these out. And I unless it says otherwise, I will read out the question and invite each of the panel members to comment with their answer. So will we



50:00

Get a AGI defined as an AI that can do everything Caswell does Wait, someone wants to make you redundant. before.



50:11

Before we get close



51:23

process called Bayesian inference, where you use some priors use some experience that you already have to interpret new data coming in. And this is much more efficient than trying to figure out just purely from new data coming in what's going on all the time. And this hierarchical predictive coding is one kind of network that can do this, essentially.



51:45

And so the idea was that this network, if it was imbalanced in some way, it might go wrong.
And



51:54

it might lead to the symptom, the symptoms that we see in schizophrenia, and also in autism.



52:02

So I e



53:32

So yes, I do. People are trying to do this. It's definitely not my area of expertise. But I've seen a talk somewhere, I can't even remember who or where it was, but but people are trying to do exactly this in order to.



58:42



1:03:48

and some of these, some of these biases we're aware of, some of them were not. And so some of them, you could try and tune out in this way that Magneto just described, but there may be many that you might not. And so, yeah, it is a it's a very real problem. But to reassure you on this count, I mean, on this particular account, I think we're very unlikely to be using these kinds of methods to diagnose psychiatric problems anytime soon.



1:04:22

Because there's so many other problems, this is just one of the millions of problems.



1:04:27

But you know, there are, there are scenarios where things think that you know, legal scenarios where potentially legal processes can be potentially replaced by these AI



1:04:43

machines where you definitely don't want these biases to cre

would you go into solving that? I'm going to repeat that quickly for for the record. So the point is, basically, because we've already got this sort of resource disparity

Non machine learning methods is something that I think is really exciting.



1:10:05

So yeah, not quite my area, but that that they're the two kind of themes where I see the most potential if you like, Rick, I don't know, if you want to correct. I'm happy to be corrected.



1:10:17

No, no. I mean,



1:10:19

the interesting thing that might makes me think of is,



1:10:24

you know, it's



1:10:27

would you want to know, if you were gonna get outside was in 20 years, but you could it's an engineering challenge4è 0:19



1:13:42

compared with people like Steve, they studied in their own sake, and from because, you know, in economics, there is always you have to convince economists, they are very smart and very stubborn. Why you want to study metacognition? And, and, and this, I was always trying to find the reason why metacognition is for and metacognition can be for maybe correct because you know, it comes after you made a decision, that's the strange things about metacognition is almost too late to correct whatever you're done. But the fact is, we don't do things only once in like, we repeat them. And now the thing is, is the fact that the complex system as a is almost like, what's the point of IB consciousness, you know, philosophers have debated this thing. It's not like evolution likes to make conscious thinks there might be there is a bar there is a point in which to have this sophisticated level of control. You need a second order system, that is what metacognition is, and then consciousness is going to become a byproduct of that. And some people that work in AI, the I know they it all, you know if consciousness is important is gonna erase on the way to it



1:15:00

I feel that there are very good reason why, you know if evolution evolution has endowed us with consciousness because seeing Right, right, this is an interesting experience. So it might have been very strong pressure for it. And we're going to face the same pressure in those artificial agent and I think we are already facing it because one of the point I've made in some article that the reason why this human can learn with little data you know, the things I've been keeping telling tonight, learning with little data wanderings really helps to learn where little data is having the second order system.



1:15:42

And eventually one system alone with little data. So eventually, we're going to face the fact that we need to this second order system, and then is going to just, you know, generate protocol, viciousness. Maybe, maybe wait when steps come back, and they will finish to answer the talking about Steve. CV is almost like



1:16:06

my best friend I just like a matter, Steve, if you're listening me now on the podcast. We don't want Steve to thi

Question in the back there? Hi, it has been thrown up quite a few times today that





1:18:03

Rather than under generalisation, one, there are



1:18:07

there are things that the hippocampus does, that definitely are not there on or severely impaired in schizophrenia and psychosis. So for example, I don't know about general, I don't know if anyone has ever tested



1:18:25

generalisation itself properly, that would be an interesting experiment to do is to broader concepts, testing one experiment, but but they've they've tried coupling, different inferences together. So learning that A and B go together, and B and C go together. And if you learn those things, then your your hippocampus also learns the A and B and C but but

But it's



1:27:51

Yeah, so



1:27:53

yeah. So a school. Yeah, I had no idea. I had no idea what I wanted to do. I didn't even know if I wanted to do science, sciences or, or kind of English and stuff. I chose sciences. @w h



1:33:58

about the brain. It's also probably the most embarrassing things in neuroscience that we spend, have every day, a little chunk of our life, creating imaginary world, fantastic word in our head, when we sleep, and the fact that today was being saved in order to convince somebody you need the narrative, the fact that we are such narrative animal, and we are not just hallucinating with images we create amazing narrative will happen tonight, to all of us multiple time. And we know if nothing about it. We don't study it, because we feel embarrass y