NUMENTA
ON INTELLIGENCE

Matt:	00:00	Today on the Numenta On Intelligence podcast.
Subutai:	00:04	If you're a computer scientist, you might know about multiplexers where you set a particular switching function. Depending on a state, you, you set the system to switch input in a certain way and when input comes in and gets moved around and switched around, so the Ballon boundless could do that through by control of the via control of the neocortex.
Matt:	00:24	That was Subutai Ahmad on the Thalamus. I'm Matt Taylor, Numenta community manager, and you're listening to the Numenta On Intelligence podcast. I'm here with, Subutai Ahmad our VP of engineering at Numenta and he recently attended the COSYNE conference and I wanted to talk to you about it. High. Subutai.
Subutai:	00:45	Hi Matt. Happy to be here.
Matt:	00:47	Awesome. So, um, tell me about the COSYNE conference. What, who attends and what is it all about?
Subutai:	00:52	Yeah, so the COSYNE conference is pretty much the main conference in computational systems, neuroscience. That's what COSYNE stands for, every year. And it's probably one of the main conferences that we attend every year. Uh, it originally split out from NIPS, um, way back in the 90s. I think, um, which is now called [inaudible], it's called Neurips as of this year. But back then it was Nips. Um, and basically it's, it's people who really wanted to focus more on the neuroscience rather than the math and pure math and machine learning. And so it's a great conference because it's, um, it's grown a lot and it's a really nice mix of experimental neuroscientists and computational neuroscientists getting together in one place and

discussing a lot of different topics. So it's one of the only conferences where you'll see almost like a 50, 50 mix between experimental papers and computational papers.

Matt: 01:45 So how well does Numenta fit into that mix? Where do we fit in there?

Subutai: 01:51 I think we're like the prototypical lab in some sense because we, we interact a ton with experimental neuroscientists. A lot of the stuff we do really draws on, on their work and their research. Um, and we have computational models that take into account a lot of biological details, so we always get really great reception at, uh, at COSYNE whenever we present stuff.

Matt: 02:13 Great. What did you find interesting there since you attended this year?

Subutai: 02:17 Yeah, this year was, um, uh, so this year it was in Lisbon, Portugal, which is different. Uh, mostly it's been in Salt Lake City or Denver, uh, you know, in the rocky mountains, basically.

Matt: 02:29 Not very exotic.

Subutai: 02:30 Not very exotic and it's been hard for Europeans and international people to get there. This year they made a point of mentioning this. It was a lot more international and uh, that was great. And you could see the mix of ideas. It was a little bit different. Um, so that was pretty cool. It's also, it wasn't nice to be in Lisbon, um, uh, in, in terms of kind of what I found, interesting there. There's a lot of stuff. Um, uh, you know, Edvard Moser was there and he talked about their work on grid cells. Um, and that was, we are familiar with a lot of that work, but it was just great to see him present and talk about it. So that was really cool. Um, there was, uh, after the main conference there's workshops and this year there were two workshops that were extremely relevant to our, our work. So there was a full day workshop on continuous learning.

Matt: 03:21 A full day workshop, full day workshop mean you come bring your computer and do actual work or you just sit and listen.

Subutai: 03:26 Yeah. So the, these workshops are like mini conferences, so they're split up. So the COSYNE as a whole, has about a thousand attendees roughly. And each workshop had anywhere between 20 to a hundred people attend.

Matt: 03:39 That's pretty big.

Subutai:	[03:40]	It's still pretty big. Um, but, uh, what you have is a set of talks and it tends to be more discussion focused so people can sometimes interrupt during a talk. Um, and there's a lot more room for a discussion and you're pretty much looking at one focus topic for the whole day. So everyone there is kind of interested in that topic and it's a nice way to meet people.
Matt:	[04:01]	Yeah. Working on the same stuff that you're interested in.
Subutai:	[04:04]	Yeah, exactly. Yeah. And so it was nice to see a whole topic on continuous learning. Um, so again, they had some experimental work showing how people can learn or when they forget and things like that. And then some work in machine learning that has tried to address the problem. So continuous learning,
Matt:	[04:21]	Is it the same as thing as online learning? Would you say...
Subutai:	[04:23]	yeah, it's, it's, it's similar. Yeah. And some people call it lifelong learning...
Matt:	[04:28]	Oh, I haven't heard that.
Subutai:	[04:30]	Um, so that was that. I think I learned a lot from there. So that was really cool. And then the next day we had a full day workshop on coordinate transformations and reference frames.
Matt:	[04:40]	Oh Wow.
Subutai:	[04:41]	So that's of course extremely relevant.
Matt:	[04:43]	Grid cell heavy?
Subutai:	[04:44]	It was very grid cell and hippocampus heavy I'll say. So it was a lot about hippocampal stuff and a little bit on the on Cortex. Um, but there was a lot of really interesting kind of ideas and, and it was just interesting to see how people think about, uh, these, these terms. And so, so,
Matt:	[05:01]	so, uh, did you attend just to see the, the workshops and the presentations or did you present anything?
Subutai:	[05:07]	So, uh, we had a poster there so it was not just attending but also presenting some of our work and I can talk about that and it's, uh, and you know, that way you, there's really a two way interaction between the, the other conference attendees and our stuff. And um, so that's good.

Matt:	05:24	So that's where there's like a hall of posters and what do you have like six feet by eight feet or something like that to shove as much information and graphics as you can.
Subutai:	05:33	Um, I might have tweeted the picture of the poster session, I forget, but it's, yeah, you have a large hall and maybe 50 posters or something or, or a hundred posters sometimes in, in a hall. And you're separated by about five feet or six feet. And each poster is about, you know, six by three years, something like that. And then you stand in front of your poster and anyone who's interested can come by and talk to you about it. Usually there's an open bar at that time as well. So that kind of lubricates the discussion sometimes. Um, and it's a great way to engage with people one on one or one on two or, or so on and going into depth and, and, and the topics. And you also get to walk around and meet other people. And I, in many ways, I like it better than the oral, the oral sessions are also great, but this is a completely different type of interaction.
Matt:	06:21	You get a lot of one on one.
Subutai:	06:22	You got a lot of one on one.
Matt:	06:23	Meet people, press the flesh as they say.
Subutai:	06:26	Exactly. Yeah.
Matt:	06:27	So what was your poster about?
Subutai:	06:28	Uh, so this year we had something slightly different. We had a poster on the Thalamus.
Matt:	06:35	It's a hot topic.
Subutai:	06:38	It's a hot topic. I did a blog post on it, on why more people should pay attention to the Thalamus. And uh, the, so this was a kind of a novel proposal on how, uh, the Thalamus, a dendritic mechanisms and the talents could be used to make the Thalamus a very dynamic kind of routing and multiplexing system. Um, and so we walked through it in a fair amount of detail, how exactly that might happen in the, in the Thalamus.
Matt:	07:04	And you worked with a collaborator?
Subutai:	07:06	Yeah, so this was a collaboration with Carmen Varella who was at MIT when we started discussing and now she's a professor at Florida Atlantic University. So she, and I had been working on

this since this was, this actually came out of our guests scientists program. So she was at MIT and then she visited us here for a few days and we discussed all things thalamus. And then, uh, this is one of those things where she was talking about some specific mechanisms and I was thinking about from a computational standpoint and we both kind of realized, hey, we could, it could be used to fulfill this computational need that we think we have. And it could be kind of the perfect mechanism for that because I think it's a really good example of experimental theory collaboration.

Matt:	07:48	Yeah. I think that's neat too. To see a neuroscientist come and visit and then we collaborate with them and go present a concept that's very, that's perfect.
Subutai:	07:55	Yes. And this was not my main job in, this was not her main job. This sort of came out suddenly out of this collaboration. And so we both been kind of working on it on the side, trying to develop the idea.
Matt:	08:06	and if you want to find out more about your Carmen Varela and her work on thalamus, there's an interview that I did with her, uh, last year that I'll link in the podcast show notes. Um, so let's talk about some of the details of this paper if you don't mind. Uh, and our last podcast I talked to Jeff about, about the Thalamus, and I'm going to talk to you about some different things about the thalamus. So in this poster, you talk about the difference between the modes that the thalamus might be in, one's called the burst and tonic. Can you describe what those modes are?
Subutai:	08:34	Yeah, I can, uh, maybe before I jump into that and maybe I can just give like a high level picture of the thalamus, we can dive into that specific question as well. So the Thalamus has, um, uh, you know, many different possible roles. And Jeff has, uh, ideas about matrix cells and how they can be involved in timing and also controlling of oscillations and stuff. And that's one type of cells. Um, but in general, the Thalamus is right in between all cortical regions. So whenever a cortical region projects information forward to another cortical region, there's sort of a copy of the signal that gets, uh, or not a copy of that exact same, but there's also a pathway through the Thalamus. And for a long time people thought that the Thalamus was just a gating center. It would just relay information from one part of the cortex or the other.
Matt:	09:25	These cells are called relay cells.

Subutai:	09:25	They're called relay cells exactly because of that. But if you look at the anatomy of the Thalamus and the physiology, it looks a lot more complicated than that. And it can't just be a, a relay station. And the other thing is that the connections between the cortex and the Thalamus are highly regular. They're somewhat complicated, but they are very regular. So, just like the cortical column has a regular structure, the cortical column, Thalamic interactions, and the anatomy there has a very regular structure. And this has been written about by Murray Sherman and Ray Guillery and others. There's a very regular kind of pattern.
Matt:	10:03	Like across cortex
Subutai:	10:04	Across cortex.
Matt:	10:06	Across neocortex. I think Jeff said this, said it in a way, like it's almost like it's a part of the cortical column circuit that the thalamus could be another layer.
Subutai:	10:15	And that's that. And that's, yeah, I think that comes from me, but yeah, cause I, I, the way I see it, I just think of the thalamus as layer seven of the cortex. It's so tightly tied to the other layers.
Matt:	10:27	That decomplicates it a bit.
Subutai:	10:29	Yeah. Yeah. I think it makes it easier. Um, and it gives you a filter of how to read through thalamic papers as well when you, when you read it. And so for us, uh, doing models of the cortical column, I think it's critical to understand the thalamus and how it, how it operates.
Matt:	10:42	So what are these modes that it might not?
Subutai:	10:44	Yeah, so thalamic cells, uh, these relay cells or thalama-cortical cells, uh, they project to the cortex. Um, and they, uh, cortical cells are, you know, the fire and they don't fire. Um, there's an but, uh, thalamic cells, they tend to have a much higher regular rate of firing then cortical cells. Um, but what's really important with, uh, with thalamic cells is not just whether they fire or not, but what mode they're in when they're firing. And the two main modes are tonic mode and burst mode. And tonic mode is kind of the regular neuron firing. It Just fires at a regular rate or you know, at different rates, but it's just sort of regularly firing in. Burst mode is this special mode where it's quiet for a bit and then there's a burst of spikes at a very high frequency that, that

		just gets triggered. Um, and some people think this has very different downstream implications. So if you think about a cortical neuron and your thalamic input is sort of, you know, regularly firing you treat it one way and if it's bursting, you treat it differently. It has a different effect on the, on the post synaptic cell.
Matt:	[11:52]()	so said you said bursting, is that equivalent to the idea of bursting in cortex as well?
Subutai:	[11:58]()	So there's bursts. Um, there's burst type activity in cortex as well. And we've used the term in a couple of different ways. So, um, so if you're familiar with our temporal memory, we talk about columns bursting.
Matt:	[12:10]()	Minicolumns bursting.
Subutai:	[12:10]()	Mini columns bursting. Thank you. Uh, and that's a slightly different thing. That's more like, oh, there's a burst of activity across many cells. So this is a single cell that fires at a very fast rate for just like 50 milliseconds or so. Okay. And that also happens in the neocortex. So layer five cells for example, are known to do bursting in, in certain situations. And I think they've seen it in other cells as well.
Matt:	[12:37]()	So when we talk about the modes of the Thalamus, tonic versus burst, is this, are we talking about individual cells and, and being in different modes in their surrounding cells or are whole sections in one mode or the other?
Subutai:	[12:47]()	Uh, so we think it can be both, but the property is a property of a single cell. Okay. Um, but we think what would we presented our poster is that the cortex can control large parts of Thalamus and, uh, all of thalamus until this, this part needs to be in tonic mode and this part needs to be in burst mountain and you can kind of think of burst mode is like, uh, uh, pay attention to me mode like this, it has a stronger impact or, you know, it's, it's a different signal than, than the tonic mode.
Matt:	[13:19]()	Yeah. Um, so how does the cortex have any control over those cells and what mode they're in?
Subutai:	[13:26]()	Yeah. So thalamocortical cells, or these relay cells, they project to cortex. But what's interesting is there's a feedback projection from cortex as well.
Matt:	[13:35]()	There's a loop there.

Subutai:	13:36	There's, uh, it's, it's different cells that project back.
Matt:	13:39	Oh, okay.
Subutai:	13:39	So, uh, it's so typically Thalmus might project to Layer 4 and some parts of Layer 5. Um, and then Layer 6A in cortex projects back to the Thalamus. Okay. Um, and these are the cortical thalamic cells. The terminology is a little, uh, uh, sometimes confusing. SoLayer 6A projects back to the Thalamus. And the projection pattern back is kind of interesting and intricate. And what we proposed in this poster is that Layer 6A cells when they project back to the Thalamus, they can very precisely control which cells are in tonic mode or burst mode, uh, in a, in a very, very precise way.
Matt:	14:19	Based on its representation of whatever is being modeled. within the column?
Subutai:	14:22	Yeah, yeah, exactly. So the other parts, other layers of cortex could be modeling something. And, um, you know, and then Layer 6A could very precisely control what's going on in how the Thalamus, uh, behaves basically. And what's interesting is that the way Layer 6A projects to the Thalamus. Again, this is very regular pattern, um, and it projects to two different cell types in the thalamus. So one is this, um, what's called the thalamicreticular nucleus, which is a layer of inhibitory cells that surrounds the Thalamus. And then it also directly projects to the relay cells to the thalama-cortical cells. And um, so Layer 6A projects to both of those. And then the inhibitory cells, these trn cells, they also projected the dendrites of these relay cells.
Matt:	15:17	So there's two stops.
Subutai:	15:18	There's two stops there. And so it's kind of a interesting pattern. It's a regular pattern. And so what we realized is that the way the dendrites function, um, on, on these relay cells, and there's this thing called these t calcium channels, which I can describe, I don't know how much detail to go into here, but uh, but um, the, the way these dendrites function, um, Layer 6A could be sending very specific codes down. These trn sales could be recognizing these codes and then causing the relay cells to get into burst mode versus tonic mode through these, through this dendritic mechanism.
Matt:	15:57	And that's not even the direct information they're sending to the relay cells. That's like a contextual signal?

Subutai:	16:04	Yeah. So imagine like the trn cell is detecting context of different types and, um, recognizes that context and then sends these inhibitory signals, uh, to the dendrites of the relay cells, right? And why inhibitory signals? Well, these, uh, the dendrites of the relay cells, they're sprinkled with these, what's called these t type calcium channels. These are voltage gated channels and they have a very, very particular behavior. Um, they have sort of three stages. Uh, they can be in a normal inactivated stage. And if these channels are in an inactivated stage, then whenever the relay cell fires that just fires, it's like tonic mode, right? Um, but if you inhibit the dendrite, um, and these T type calcium channels, if you give it hyperpolarizing input for a period of time, they get into this de-inactivated state. Don't ask me who made up these terms. That's the term, de-inactivated state. And what, what happens is that in that state, if the cell fires, if it gets input and fires, it's going to then trigger a burst. What happens is these calcium channels open up and the, uh, short flood of calcium comes into the cell and it triggers multiple spikes.
Matt:	17:19	And so it's a little chain reaction going gets a little Twitter, Twitter, Twitter.
Subutai:	17:22	Exactly. So this is very particular, you know, three stage thing. It's in tonic mode or it's ready, it's de-inactivated, which I think of as ready to burst and then it's bursting. Um, and so this trn cells can control this. It can switch the dendrites, um, you know, between these different modes, um, to cause the cell to be burst ready or tonic mode.
Matt:	17:45	So the thing that since there's, there's, um, routes going from the Cortex to the Thalamus and from the Thalamus to the Cortex, like what's controlling what? Like, uh, the cortex is obviously affecting the thalamus and what mode it's going to be. Yeah. And then what is the Thalamus then sending back up the cortex?
Subutai:	18:02	Okay. So the Thalamus is sending whatever input it's getting from some other source, and then this Layer 6A, uh, signal is sort of modulating or controlling the nature, the way it sends that input. Okay. Uh, so think the easiest way to think about it is if you look at the visual Cortex, the first stage, um, the thalamic center is called Lgn or lateral geniculate nucleus, it gets input through the retina. And the, um, one of the uses of this mechanism might be as an attentional mechanism. So the layer six AI or cortex would be saying, hey, I really want to pay interest, uh, uh, to this part of the retina. And what it'll do is it'll set those cells to be in burst ready mode through this, um,

		through this top down code. And now, whenever, whatever input comes in to the retina, now this, the information that's within the focus of attention will burst and may have a bigger impact on learning, on just inference. Um, both actually and the stuff outside your focus of attention might have a lower number.
Matt:	19:07	So, so, so, so the Thalamus then has a, has a way to sort of focus attention on some aspect of the sensory field and put and have more emphasis on that part of it.
Subutai:	19:18	Exactly. Exactly. Yeah. And because it's dendritic, it can be, it's down at the level of individual cells. Um, and it can, so it can be a very, very precise, uh, control of it. Uh, another possible thing could do is be filtering the content. Um, so if, you know, we know that the retina, uh, is not just like pixels in a camera. It's actually detecting features like motion and you know, its color and so on. And, uh, cortex could be telling the Thalamus, Hey, I'm only interested in things that are moving away from me or coming towards me. Um, and so would emphasize those signals and de-emphasize others.
Matt:	19:58	Oh so that could be like how the Cortex is affecting what the Thalamus is doing by providing information about what should be attended to.
Subutai:	20:04	Exactly. Yeah. Yeah. So it gives a tremendous amount of power into, you know, how you kind of filter the stuff coming in. It's not just a relay cell. You can really contextually filter it. Um, and again, because it's dendritic, it can actually have a very powerful, uh, it can do lots of transformations of the input. Um, if you're a computer scientist, you might know about multiplexers where you set a particular switching function. Depending on a state, you set the system to switch input in a certain way. And when input comes in and gets moved around and switched around, so the thalamus could do that via control of the cortex.
Matt:	20:47	That's one of the hypotheses.
Subutai:	20:48	That's one of the hypotheses. So it's actually a very powerful way of kind of transforming the inputs, um, uh, that are coming in.
Matt:	20:56	What could the brain be doing with that?

Subutai:	21:00	Yeah, that's a great question. Um, so, uh, you know, by talking about filtering and attention, uh, Jeff has talked about ways of doing variance in there, so it could actually be implementing a little bit of a translation and variants or scale invariants. Um, it could be shifting inputs left or right, and there's some theories on how specific types of attention might be implemented with shifting, uh, with what's called shifter circuits. Um, so there's a lot of different possibilities. And in this poster we're just saying, hey, this is a really powerful mechanism of doing all sorts of transformations and then we can figure out exactly what the cortex might want to do with this. And what's cool is that it's not just LGN to V1. This is between every single cortical region. The thalamus is a center routing information or transforming information between all cortical regions in some sense.
Matt:	21:57	Yeah. Talking to Jeff in the last podcast made me realize that, I mean, you can't exist without your Thalamus, the neocortex can't do anything without it and vice versa. The thalamus doesn't, can't do anything without the model the neocortex provides, too.
Subutai:	22:11	Exactly. They're intricately tied together.
Matt:	22:13	Yeah. That was his point is I think of it as the same thing at this point. And that makes sense to me. So it's really important that we figure out what it's doing.
Subutai:	22:20	Yeah. It's, it's, it's part of figuring out what the cortical column is doing. It's no different.
Matt:	22:25	Well, that's what we're working on here at Numenta. And, uh, thanks Subutai for sharing, uh, your, your COSYNE experience with the podcast viewers.
Subutai:	22:33	Thank you. This was fun.
Matt:	22:34	It's been a pleasure.
Matt:	22:36	Thanks for listening to the Numenta On Intelligence podcast. Next episode, I'll talk again to Jeff Hawkins. We'll be talking about the definition of intelligence. Until next time.