Episode 3: New to Numenta? Top 5 Things You Need to Know

August 28, 2018

Christy:	00:00	Hi, this is Christy Maver.
Matt:	00:02	And I'm Matt Taylor and you're listening to the Numenta On Intelligence podcast.
Christy:	00:10	Today's episode is designed for people who might be unfamiliar with Numenta and our work. Hopefully this episode will pique your interest and entice you to follow along with what we're doing. If you want to stay up to date, the best way to do that is to sign up for our newsletter. You'll find a link to subscribe on our website at numenta.com. We typically send a newsletter every one to two months. Basically, anytime we have news to share about upcoming events, new papers we've published, and partner updates, just to name a few. Subscribe today to make sure you don't miss any Numenta news.
Matt:	00:45	Thanks for joining us. Hi Christy.
Christy:	00:46	Hi Matt. How are you?
Matt:	00:47	I'm doing great. I was thinking about the last episode, which was episode one we did with Jeff
Christy:	00:53	Jeff Hawkins, our co-founder. Two parts.
Jeff:	00:55	We're really bad at numbering episodes, so we have episode zero, episode one, part one, and episode one, part two, and now this is episode two.
Christy:	01:05	This is episode two slash three.

Matt:	01:07	Makes total sense. Anyway, that last episode went really deep and we just want to let everybody know that it's not always going to be like that.
Christy:	01:15	Right. So Matt and I were thinking after listening to the last episode,
Matt:	01:19	Let's do an episode that's like the opposite, not a huge deep dive, just a shallow skim on the surface of what it is.
Christy:	01:26	So maybe if you're new to Numenta or a beginner in brains or just mildly curious, this is a great episode to start on for you.
Matt:	01:36	Especially if the last one just went way over your head. Don't give up on the podcast.
Christy:	01:41	We're going to mix it up.
Matt:	01:42	We're going to have some interesting interviews with neuroscientists and interns and ceos coming up.
Christy	01:47	So what we want to do in this episode is we want to break down the top five things you need to know about Numenta,
Matt:	01:54	The first one being what is our mission? Why do we exist, what are we here for?
Christy:	02:00	What are we here for Matt?
Matt:	02:02	Well, I'll tell you what we're here for. When I first heard about Numenta is right after I read the book On Intelligence by Jeff and that really affected the way I thought. So I looked it up. I looked at what Jeff Hawkins was doing and I saw that he had founded this company called Numenta and their mission was to understand intelligence and I thought that was just an amazing mission for a company in Silicon Valley to have and that is still our mission even now over 10 years later, is to understand how intelligence works in the brain.
Christy:	02:32	Yeah. And you know, we often talk about how we actually have this dual mission of number one to understand how the brain works, understand intelligence, and that is very much a scientific mission. Brain theory driven, deep neuroscience details.
Matt:	02:50	That is number one by Jeff because he wants to know, like honestly because he's curious, right?

Christy:	03:00	I think more than just Jeff wants to know.
Matt:	03:01	Yeah, I mean it's, there's a lot of passion behind our mission because another thing I like about Numenta.
Christy:	03:06	So the first part of the mission is a scientific one. There is a secondary mission, which is to apply those principles of intelligence to software to enable intelligent machines. So there is a machine intelligence, AI component.
Matt:	03:23	There is, which gets us, I think grouped in with a lot of AI companies. Uh, but I, I would never have considered us to be a real AI company. I don't like the term AI, artificial intelligence. I think a better term would be machine intelligence because there's nothing artificial about intelligence. It's either intelligent or it's not.
Christy:	03:42	And in order to enable machine intelligence, we first have to have an understanding of what intelligence is. And that's what leads us to the brain because that's really the best, if not the only example that we have.
Matt:	03:54	We run in a lot of neuroscience circles. So it's sort of obvious that if you want to, to understand intelligence, you know, the only thing in the universe we know of that's intelligent is the brain. So let's try and understand that. But I do think there's, there's a lot of computer scientists and mathematicians who have been working from the other direction trying to crack that problem of intelligence in lots of different, other interesting ways. So it's just, that has never been our mission. It's not to build intelligent machines. It's to understand intelligence.
Christy:	04:23	Yes. So yeah, so that, that's really been our mission from the very beginning. And, and we've been at this for more than 10 years. So let's talk a little bit about what we've learned so far along the way about how the brain, how the brain works.
Matt:	04:39	I think we've had a couple of big achievements and one of them is about sequence memory and in 2013 when we released all of our, our core code as open source. That, that was the big discovery that we had was about sequence memory. And it's really about how your brain memorizes spatial patterns over time. And the time thing is really core in that it's a, it's how cells, how they, how they connect together to, to create these sequences that can be replayed, like a melody,

Christy:	05:14	Like a melody. Yes. And that's the example that I love, that works best for me because when you think about how many, how many melodies you recognize, right? Hundreds, at least, right? If not, if not more. And when you're listening to a song, you are constantly predicting what, what the next part is, because you know the song, you know what the next note will be, and if it's different, you will detect that as something different. You will detect that as an anomaly.
Matt:	05:44	Right. And you know, you, you could even invoke sequences in people's brains listening to the podcast right now. So get this. (Matt sings) Everybody who's ever heard that song that classic, I'm assuming it's Beethoven, I hope it's Beethoven, I'm going to assume, I'm just going to call it Beethoven. But I'm not into classical music reading neuroscience here. So it doesn't matter if I'm wrong about about classical music anyway. Everyone knows that song. And, and you're probably hearing it in your head right now. We've invoked it just by those four notes, just by playing those four notes in that order, at any instrument, in any voice, you can sing those notes and invoke that sequence of how you experienced that song in your life in everybody's brain. I think that's pretty interesting. But that idea, you know, that sequence memory that ties those notes together, the way your brain has memorized that over time, that's at the core of that first discovery, right? We're talking about in the neuron paper. We'll link to it in the show notes if you want to read the details.
Christy:	06:44	Yeah, so as Matt said, that that was really one of the first big discoveries. I think the second area of discovery builds on that sequence memory and is about how we learn objects over time through movement. So when you think about listening to a song, you can be sitting absolutely still and take in the song and you're still predicting what note will be next, but most of how we experience the world is as a result of our own movements. So that's what our second kind of area for discovery starts to address.
Matt:	07:18	Yeah, and so I think about, I like to think about this new part of the theory about object representation and about movement in sort of three different ways. There's a learning an object and the best metaphor for this is if you pretend to reach into a dark box, you can't see anything in it, but you just put your finger in and you touch something in the box. I did this experiment with Jeff in the last episode, but at a basic level, if you're feeling something that you've never felt before, you've got to touch it all over with your hand. You've got to traverse that object. You, you've got to move across it to like build out a representation of it in your brain and your imagination. And that idea of

movement over time is core to how we build up object representations. It's through sequences, through sequence memory, and that's why we need the first part of the theory that we just talked about.

Matt: 08:08 So it's crucial to have the first part for this object recognition to build on top of it. Um, so you learn objects through sequences and also you have to be able to infer objects. If you were to touch something in a box and say you've touched a lot of things in your life, you could imagine what it might be. If it's furry, it could only be a certain amount of things. So that's inference, you know, that's, that's like, I'm going to guess what this is based on interest based on what I know so far. Um, and then the third part is prediction. Prediction also involves sequence memory because it, it gives you the ability to predict an object state knowing that it's been, you've seen it in this sequence of states over time you get an indication of how likely a state of that object will be in the future. Um, so I think those three parts are really core to this, uh, this idea of, of object representations with sequence memory. So that's object recognition in a nutshell, I guess.

Christy: 09:06 So if you're interested in learning more about these discoveries, we have papers and we'll link to them in the show notes. The paper is pretty detailed, but we have some additional supporting resources as well. So I think a natural next question would be, okay, so that's what you've learned. What about applications? What have you built?

Matt: 09:29 Yeah. Putting aside the whole object recognition thing, the, the, the older stuff about sequence memory, and that was in the "neuron paper," that's been around since 2013 in the public. Um, it has some interesting applications for anomaly detection. We really think this whole sequence memory was a breakthrough and I think there's a lot of still untapped potential there. There are some commercial applications being built around streaming anomaly detection. I think there's other opportunities for people to work in that area, especially what I'm still excited about is geospatial, objects moving through time and space. That's something that current HTM techniques can take advantage of right now. So you can track things moving through space over time and, uh, encode how anomalous that object's behavior is. So, uh, if you're talking about logistics or routing for, for planes or trains or automobiles or, or whatever. There's a lot of potential applications there in anomaly detection.

Christy:	10:31	Yeah, it's really anything, anything that has data streaming from something, right, whether it's a sensor or a GPS signal, or medical device or I mean anything that has a sequence of data where you can make predictions, you know, which is a lot of things, especially in the, the, you know, Internet of things era.
Matt:	10:56	Yeah. And a lot more things nowadays have locations that move around too.
Christy:	11:01	So we have a number of example applications that are available for anyone to look at it, to play with.
Matt:	11:08	Oh, HTM Studio, Talk about HTM Studio.
Christy:	11:12	Yes. So HTM studio is near and dear to my heart because it doesn't require any technical skills. So it's a tool that we released a couple of years ago that's specifically designed to let you test our technology on your data to see if it finds interesting anomalies. So anytime people come to us and they think they have a use case or a potential application and they want to know if they can use HTM, that's the first place we send them because it's, you know, it's a minimal investment, it's easy. It's a good way to test it.
Matt:	11:46	All you have to do is get your data in a CSV format. That's pretty simple. And uh, and it takes care of the rest, honestly. A lot of people will come into the technical forum on, you know, the hackers forum in our, in our community and ask, "Is this HTM going to work well on this data and I'll just tell them, go download HTM Studio. It's just, uh, it, it works on Mac and windows, so you can install it easily and then just upload a CSV into it. It's got sample applications. It'll run NuPIC in the background. It runs HTM in the background and it'll tell you where the anomalies are and that streaming data. So if you have any of that data that Christy was just talking about, you know, from devices or whatever, give it a shot because if it's, if it has patterns over time, hourly patterns, daily patterns or even weekly and monthly patterns, it'll find it given enough data.
Christy:	12:33	That's really what makes this approach to anomaly detection unique is that it's not just about finding a spike or using a threshold or you know, kind of those more traditional techniques. It's really able to find the more subtle,
Matt:	12:47	the nuanced stuff.

Christy:	12:48	Yes. So all of our example applications are available for you to see, the code is available in open source, which actually brings us to our next thing you need to know about Numenta.
Matt:	12:59	Yeah. And that's that we are very transparent in everything that we do. We believe in open science and so if you follow Subutai, our VP of research on Twitter for example, you'll know he's always posting about open science issues. We only publish in open journals that are free and open access. So we, we really encourage that. We put as much as we can online, you know, a lot of times I'll record meetings and put them on YouTube, or if someone on a forum is asking about something, I'll just grab someone and, and we'll answer the question and throw it out there.
Christy:	13:33	Yeah, the daily research code, I mean it's, it's all out there, right? It's all accessible. There's nothing really hidden.
Matt:	13:39	No, all of our research papers are online and free, like I said. We're trying to get more of the research paper code in a standardized format so that it's really easy to run for people and to replicate all of our experiments. Christy has a great events page on numenta.com that's got all of our upcoming talks and conferences, including previous talks and conferences and she always put slides and videos in there. So there's a lot of media that you can consume from our previous events and our events page that's really nice.
Christy:	14:13	We also have an educational youtube series called HTM School, which Matt hosts and you want to talk about HTM School?
Matt:	14:19	It's a from the ground up explanation of HTM theory. So we really start from the beginning, assuming you know nothing about computer science, about neuroscience and, and really attempt to build up the theory in a way that makes sense from scratch. Uh, and right now it's up to 14 episodes or something from, from binary encodings and semantics all the way up to grid cells and columns and layers and all of that stuff. So it's, uh, it's unfinished. I know there's probably going to be a few more episodes at some point, but it's a great educational resource if you don't feel like reading the papers, you can get a full understanding, I think just by watching these videos. You may have to watch a few times because I packed them full of information, but you can get the concepts of HTM theory down just by watching these videos
Christy:	15:09	and each one is about 15 minutes long. Right? So it's

Matt:	15:12	give or take the bloopers and outtakes.
Christy:	15:16	The outtakes, my favorite.
Matt:	15:16	I'm also putting together more visual documentation. So we're continuing to work on docs to make the theory even more approachable.
Christy:	15:24	So that really brings us to our last point, which is how you can get involved. You know, we'd love for as many people to get involved as possible and we also understand that there are different types of people out there, right? Not everyone's a neuroscientist, not everyone's going to read every peer reviewed journal paper, that we release, and we understand that. So if you're looking for more of the, you know, 30,000 foot view, maybe you're interested in AI and you're interested in brain related approaches. We have a few pieces and we'll put these in the show notes. We'll put links in the show notes and a few pieces that Jeff has written about this topic. One called The Secret to Strong AI, one called "The Thousand Brains Model of Intelligence" and an article he wrote for a special edition of IEEE Spectrum that talks about what intelligent machines need to learn from the neocortex. And all three of these articles really highlight our most recent advances in the research and talk about why we think they're so important for and the implications for machine intelligence.
Matt:	16:29	Yeah, they, they try and explain what we're doing and why it's important sort of in different ways to different audiences.
Christy:	16:35	Right. And then of course, if you do want to get down into the scientific details and then hopefully you've already listened to the first couple episodes of this podcast, but we have all of our papers available on the website as well. And for each of our peer reviewed papers, we have an FAQ section that talks about kind of the highlights and what you need to know about each paper.
Matt:	16:59	And if passive learning is not enough for you, you should join our online community or we've got a really robust forum called HTM forum. There's not only a hacker community of, of people that are trying to build things with HTMs or build their own HTMs and code, uh, but there's a, there's a thriving theory community as well that are interested in the repercussions of HTM theory extensions to HTM theory, linking it to other parts of neuroscience or psychology. There's lots of people with lots of interesting ideas on our public forums. So join there if you

		want to learn more and really interact with our community. Yeah.
Christy:	17:34	And Matt manages that community, so he sees all the posts that come in and makes sure things get answered.
Matt:	17:40	I do, yes. You'll have a warm welcome.
Christy:	17:43	It is a friendly community. I think that's fair to say. And it's also a place that, that almost becomes a bit of a, at times a technical support forum where people who have questions about HTM or Numenta can, can go to the forum and some of the community members can answer.
Matt:	17:59	You know, our, our community has gotten mature enough now that I've got a small group of people that I can depend upon to answer questions when new people come in and they, and they know the theory well enough that they can answer the questions and that's really nice to get the community to that maturity level.
Christy:	18:16	And you can also join without posting.
Matt:	18:19	Oh absolutely.
Christy:	18:19	If you just want to read the threads
Matt:	18:22	You can lurk. You don't even have to join, you know, you can just go through and read all the messages. But it's more fun if you, if you make an account because then you can like posts and interact with people.
Christy:	18:33	So hopefully we'll see you there.
Matt:	18:35	I hope so. So that's a broad summary of Numenta. We hope you enjoyed that, you know, five things that you should know. What were they, Christy? We should review them.
Christy:	18:44	Yes. So the first was (1) our mission, our scientific mission to understand the brain,
Matt:	18:50	(2) what have we learned so far, and (3) some things about our anomaly detection applications,
Christy:	18:54	(4) Our commitment to open access and research transparency,

Matt:	18:59	(5) and all the ways that you can get more involved in our community or learn more about what we do.
Christy:	19:07	Yes. So there you have it. Hopefully you found that helpful and hopefully you feel like you know us a little bit more.
Matt:	19:14	It's been a pleasure sharing all this information with you, podcast audience. Please keep listening, subscribe to our podcast Numenta On Intelligence.
Christy:	19:23	Thanks.
Matt:	19:26	Thanks for listening to Numenta On Intelligence. To learn more about Numenta and the progress we're making on understanding how the brain works, go to numenta.com.