Episode 7: Conversation with Numenta CEO Donna Dubinsky

Christy:	00:00	You are listening to Numenta On Intelligence, a monthly podcast about how intelligence works in the brain and how to implement it in nonbiological systems. I'm Christy Maver. For the past couple episodes, my co-host Matt Taylor has continued his Interview with a Neuroscientist series and gone pretty deep into various neuroscience topics. For our final episode of 2018, we're taking it back to the business side. I interviewed our CEO Donna Dubinsky actually a few months ago, but we saved this episode to be the final one of the year. During our chat, I asked her about the business evolution of Numenta, the challenges of navigating both a scientific and a machine intelligence mission and why she believes it's so important, even for non-neuroscientists, to understand how the brain works. As always, you can keep up with our latest progress by signing up for our newsletter and visiting numenta.com. All right, here's my interview with Numenta CEO Donna Dubinsky.
Christy:	01:07	Hi, I'm Christy Maver and I'm sitting here today with Donna Dubinsky, CEO and Co-Founder of Numenta. Hi, Donna.
Donna:	01:14	Hello, Christy.
Christy:	01:16	So, thank you for being on Numenta On Intelligence.
Donna:	01:19	My pleasure.

Christy:	01:19	I'm specifically excited to talk to you because as you know, we do some pretty detailed neuroscience research as a company here. And like me, you do not come from a scientific background, so I'm really excited to talk to you more about Numenta, our work, what it means and particularly why you are drawn to it. So the first question I want to ask you is one that I get often which is, Numenta is such a unique company. How do you explain to people what it is that we do?
Donna:	01:53	Well, Numenta is a unique company. It's not so much that it's a unique effort, but I think the form of it is really unique to have created a scientific entity that also has a commercial mission is a unique structure. There certainly are scientists working on hard problems throughout the world and academia and private research institutes, but other than maybe big companies that have research components you don't certainly see it very often in a startup. In fact, I think it's interesting because so many of the venture capitalists around look at it and say, "Oh, nobody's doing big thinking in big work anymore", but we are. But I don't think the venture capitalists are really interested in that kind of big work. They're really interested in incremental work. So we're weird because we're doing science and big ideas in a commercial form.
Christy:	02:45	Right, right. And you mentioned this dual mission, right? So there's the scientific mission which has to do with reverse engineering the brain, and then there's a business mission of applying that to machine intelligence. So, how difficult has it been to steer the ship and navigate between those two missions?
Donna:	03:05	That's pretty difficult, and I think it's not clear yet what the ultimate resolution of that will be. We don't really know what the best way is to commercialize it, but let me just address why is it we want to commercialize it. It's fine to have a scientific mission. Why do we need to commercialize it? And the reason is not necessarily a monetary reason; the real reason we want to commercialize it is for impact. Both Jeff and I (Jeff Hawkins and I) have had extensive experience with developer communities over the years, where we've seen clever, amazing, fascinating work done by people who build on a platform we've created, such as the Palm Pilot or the Treo, and our idea when we created this company as a for-profit was to enable a whole generation of developers to build value-added products on top of this technology. We felt that a commercial opportunity is what would draw those developers much more than a nonprofit structure, but we're still actually kind of early in that process, so we're not sure how it's all going to work out.

Christy:	04:14	Right, right. And there's actually a Harvard Business Case about this.
Donna:	04:18	Yes, there is. I recommend reading it. It covers a lot of our different business models.
Christy:	04:24	Right, right. So let me ask you, why do you think it's important to understand how the brain works? It's such a big problem, but what does it mean to you? Why is it important?
Donna:	04:38	Once you start thinking about this, you start looking around your world and realize everything manmade that we see is created by our brains, our collective brains. I look at a giant skyscraper, and I have to stop and marvel. How could we– little, puny, weak humans– build such an incredible structure? How could we create these incredible medicines and surgeries? How do we do all of these amazing things? Well, it's an accumulation of a bunch of individuals' brains that have solved a million individual problems to make us stronger and faster and healthier and a lot of positive attributes. So as I think about that, I marvel at the brain and its capability to do this. And the question is, can we figure out how that works and figure out a way to take some of that power to have machines help us do those amazing activities and be even stronger and healthier and more powerful?
Christy:	05:40	And there must be certain examples that come up from time to time that really kind of, you know, those "Aha" moments of, oh wow, that's what the brain is doing. There are so many things that we don't really think about, and one of the examples that always resonated with me I think is something that you had said, which is you observed that with any other part of your body, like you could have a heart transplant and be fine, right? But the world only exists in your brain, and we are our brains, right?
Donna:	06:12	You are your brain. If you have a heart transplant, you're still Christy. If you had a brain transplant and put my brain in your body, you'd certainly be different. But I think that it's really true, but it's more for me the output of that brain. It's the idea that so many people have solved so many amazing problems through human history, and it made our lives better. You know, I was a history major in college. I look back at a century, two centuries, a millennium ago, and you think about the number of people who died of disease and natural disasters and poor buildings and all of the things that created human heartache and trauma in the world, and so many of those have been addressed over many years through incredible engineering and

science, and there's many more yet to do. So, can we be a part of that human progression and making the world a better place? That's the question.

Christy: 07:09 So let me ask you– There was a blog post that you co-wrote with Jeff Hawkins a couple of years ago called, "What is Machine Intelligence vs. Machine Learning vs. Deep Learning?", right? And we often get questions because those terms are thrown around so often that they've almost become meaningless. Everything now is defined as AI, Siri is AI, you know. But that post from two years ago, the two of you really took a look at the landscape and broke it down at a high level to say, "Here's how we're different." So can you touch on that perspective of what it is that makes us different for people who might hear, "Oh, machine intelligence. Got it, so you're building robots."

Donna: 07:56 It certainly does all sound the same. I think for outsiders and even for me, I read websites of people working in this field and I think to myself, "It sounds just like what we're doing. How is it different?" But when you spend a couple of weeks here or even a couple of days and you listen to the meetings, it's pretty clear pretty quickly how different it is. The team here is delving into incredible detail on the actual neuroanatomy, what is happening with one neuron and you know, dendrites and synapses and talking to each other, whereas all the people I think in everything else has referred to as AI are working on much less of a realistic neurological model. It sounds the same. They say it's neural nets. They say it's got different weights for cells, and they call it the same kind of terminology, but it is nowhere near as in-depth of a true biological constraint, if you will, as it is that we do here so our approach is much more biological in nature.

Christy: 09:08 And do you see that biological approach being embraced, or do you see it as more of a fringe approach?

Donna: 09:17 It's not yet, but I think there's a growing realization that it is an interesting place to look. I think most of the AI world is realizing that they're hitting a wall is a bit strong of a term. I mean, they've accomplished a lot. They have a lot left they could do, but what they're accomplishing is really with old kinds of algorithms and adding more power and more data. That's really what has created this resurgence of "AI", and the smart people who work in the field have looked at that and said, you know, we're not really doing this in any fundamentally smart way. We're doing it by brute force, so how could we do some of this in a smarter way? Well, maybe we should look to the one

instantiation we know of in the world that is actually intelligent, which is the brain. So we see people coming around at that point of view, but they don't really know how to go after it. If you think about it, it makes sense because there's two wholly separate groups of people. There's neuroscientists who delve down deeply into issues of the brain relative to disease or learning or other concerns, and then there's computer scientists who really know nothing about neuroscience, and they don't really intersect. And this is, I think, one of the few places in the world where they intersect at a world class level. I mean, we have world class neuroscience and world class computer science together, sitting, working these things out, and in most places it just simply isn't that.

Christy: 10:49 So this touches on something I've heard you talk about which is related to Numenta, but also kind of related to your history and career about having a front seat at these evolutions of computing. Can you explain that for listeners?

Donna: 11:03 Oh, I've had an incredible career and very lucky and very happy where I ended up over the course of my career. Not that it's over yet, but it's definitely been a long time. And that career has put me at the front row seat for four major generations of computing. The first one was personal computing when I joined Apple very early in its inception and we had the notion of putting a computer on every person's desk, and that seemed outrageous at the time. Of course, now everybody looks at it and says, "Well, of course!", But it wasn't so obvious then, and we helped make that happen and Apple. And then the second major revolution was at Palm where we said, these devices need to come off the desktop and go into your pocket. Personal computing is not the computer sitting on your desk. It's not very personal. The one that's going to be the one you carry in your pocket or your purse.. that's your personal computer of choice, and this was really Jeff [Hawkins]'s vision that I bought into and we made that happen at Palm in a significant way. We really led the handheld computing revolution and then at Handspring, we led the revolution in smartphones and really saying that these handheld devices need to be connected to each other and connected to the world. That was a front row seat at that revolution, and now my fourth revolution is a front row seat at intelligent computing, and so it's been an extraordinary privilege for my career to be able to be a part of these amazing revolutions, particularly since I'm not the technical person– I'm not the engineer, I'm not the scientist. I never could have done these things as a protagonist, if you will, but I have been able to be a key part of the team that builds the infrastructure around it, the marketing, the sales, all of the things that these ideas

need to come into the world. I've been able to be a part of that team.

Christy: 13:03 And what is your vision for Numenta? I mean ultimately the way you're looking back on those first few evolutions, I assume, will be the way that you and others look back on Numenta. But, where do you see Numenta going?

Donna: 13:23 It's never obvious when you're in it, where it's going and exactly what the applications will be. I like to say when we invented handheld computing, we never imagined Uber. It just, you know, again, this is where that development community is so important because people will take these core technologies into places you'd never could have imagined, so I believe the same thing will happen with this. We are going to solve the fundamental problems here of the brain and how it works. There will be machines that will be built on those principles. There will be developers that will put those machines to all sorts of incredible uses. And I think if you look down the road in say 10 years, let's make it far enough that you can't hold me into account, but in 10 years I think we will see real impact in the world as a result of this work and in ways that we simply cannot predict. And I know that frustrates people because people want to know today, what's the killer app, and it's just virtually impossible to guess. We could make up a bunch of stuff, but it won't be right. And so we'd rather focus on getting the core right and getting the tools out there for other people to see, how can this apply to a problem that has stymied them? We've seen people come to us for all sorts of problems where they're frustrated, whether it's in text analytics or in farming solutions or people came to us about monitoring beer in kegs. We've had every possible application, and those are the things that are going to be here in prominence in another 10 years.

Christy: 15:01 And do you worry about the, you know, there are so many conversations and some loud voices out there that are kind of the scaremongers and the fear tactics? Do you worry about that either with regards to Numenta or with regards to just the space in general?

Donna: 15:20 I think it's always fair to worry about these things, but I think that any new technology has worries that need to be managed. I like drones as an example. We came along with drones, and now we don't want them being flown near airports where it can disturb an airplane, so we need rules for that. We don't want them flying around and looking into all of our windows. The privacy implications of that are pretty serious. We better regulate that. There's always going to be malicious uses of new

technologies, and there's always going to need to be regulation to try to contain those uses as much as possible, so I don't see this new technology in any way different than that. I think the fears about intelligent machines taking over are just absurd. These machines and these programs.. They will not have agency, they will not have intent, they will not have emotions unless somebody programs that into these machines, and in which case you can un-program it or unplug the programs. They will not develop these things on their own. So, I think that they can be used for malicious purposes. I get that it will need to be regulated. It will need to be controlled in the best way. It can be just as any technology is, and I don't see it as fundamentally different.

Christy: 16:49 So, part of navigating the dual mission over the years has meant that at times we've focused on applications and even one enterprise application in particular, and then other times it's more the research. Do you see Numenta working on applications or is it unknown?

Donna: 17:08 Anything's possible. We might work on applications ourselves. We might partner with people who want to work on applications. Right now, we're taking a partnering attitude. People who come to us and want a the license to the software and the intellectual property, we are very willing to discuss a license with them. We have several licensees, and they're working on all very different problems. And so we're very open to partnering, and we'll just see.

Christy: 17:33 And what do you suggest for people that aren't neuroscientists or maybe aren't even computer scientists, but are interested in the work and want to somehow be a part of this or even wondering, "What should I do? Tell me what steps I need to take. Do I need to act on this now?" If our story resonates with someone, then the obvious next question is, "Okay, what should I do?" What advice do you have for people, whether it's individuals or companies that buy into what we're doing here?

Donna: 18:05 Well, I kind of hear two questions in that. So the first question is, if you're just intellectually engaged and you want to understand, "How does the brain work? It's interesting. You're solving that problem." There's a few things you can do. First, a good place to start is to read Jeff's book, On Intelligence. It's old. It's doesn't at all have any of the new theories in it, but it lays out some of the problems in a very easy-to-understand way that has captured a lot of people. The book has continued to do well other than chapter six, which I give you permission to skip.

Christy:	18:40	And it even says in the book at the beginning of the chapter, skip this.
Donna:	18:44	It's just very biologically detailed, but I always say that because chapter seven and eight are very important. I don't want people to put it down. But I think starting with the book is one thing, and then I think looking at our website and reading some of the more accessible parts of the things we've written, the video that we've put up– Why Brains Matter is a very good video. There are a lot of snippets of information that are there that I think could help that person. Now the second question is, if you're a company and you want to get started on actually applying this work, what would you do? The easiest place to get started today is the problem of anomaly detection on streaming data. Something your brain does is works with streaming data, as opposed to static data. We are making predictions and then finding anomalies constantly in everything that we do. We have implemented software that does that. It's available in our open source community, and that would be the lowest hanging fruit of a way to experiment with the technology. We have a product called HTM Studio. You can easily take your data and flow it into HTM Studio and see if we can find anomalies that you can't otherwise find. We have found latency in networks earlier than other techniques. It's much more effective than threshold systems, which is the way most anomaly detection works. And I think if I were in a company today and I just wanted to start experimenting, that's where I would start, with the idea of getting more deeply in would enable other applications such as prediction or classification. Ultimately, I think the area we might have the greatest impact is sensorimotor integration and robotics, but that's not really ready to be implemented yet. Those theories are still getting the t's crossed and the i's dotted, and we're not quite ready to say, "Hey, all you robotics companies out there. Come jump in today", but starting with anomaly detection is a very doable thing today.
Christy:	20:42	A good place to start. Well, Donna, thanks for joining me today and chatting about Numenta.
Donna:	20:52	Thank you.
Christy:	20:53	And thanks for listening.