NUMENTA
ON INTELLIGENCE

Episode 12: Conversation with Jeff Hawkins – On Defining Intelligence

Matt:	00:00	Today on the Numenta On Intelligence podcast.
Jeff:	00:04	You see something to believe the earth is flat and you give them all these facts that don't fit that model. Now, the Earth is flat predicts certain things. Right. And, and, and they will do their damnedest to try to fit these new facts into their model, no matter how, that's what they're going to, otherwise you have to throw the whole thing away and start over again, which is a really disconcerting thing for them.
Matt:	00:26	It is. It is uncomfortable, to have to throw away a whole frame of reference.
Jeff:	00:30	Basically, the models you build in the world are your reality.
Matt:	00:34	It's your belief system.
Jeff:	00:34	It's your belief system. It's your reality. This is what you know, that's the, oh, that's what it boils down to. There is no other reality in terms of your head. This is it. Instead of, if you say, well, you know what? Everything you believed these references, and you have to start over again. Um, then it's kind of like, oh, you're going back to square one.
Matt:	00:49	Jeff Hawkins on defining intelligence. What is it? What makes us intelligent? How was our intelligence different from a mouse'? We're narrowing down the answers to some pretty existential questions. I'm Numenta community manager, Matt Taylor.

Thanks for listening to Numenta On intelligence. All right so I'm here with Jeff and we're going to talk about intelligence in general and about how we might define it. And so I like to think about this from the question of evolution in a way. Like what are our brains do for us now? How, what did they evolve over time to do for us as humans or as this life formed animals?

Jeff: [01:31] Yeah. Yeah.

Matt: [01:32] It's a good place to start.

Jeff: [01:33] That's a good place to start. Well, sure. Evolution is a good place to start for almost anything related to living. Um, yeah, I mean I kind can viewed it that way too. I start off by saying, uh, I tell a an evolutionary story. And the evolutionary story says, well, life formed on earth, you know, billions of years ago and it was pretty dumb for a long time. Bacteria basically for billion or so years. And, um, and then when things started to move, um, movement requires, there's two ways you can move. You can move blindly, you know, which is pretty stupid. No point in moving if you don't really, you know. If you're going to move with a purpose to achieve something, you have to have some sort of knowledge about the world. Um, and that knowledge can be encoded in genes, the genetic material, but it's some sort of model. Even a bacterium has, uh, a motor plan, which is, you know, if I can move and I'll move, I'll turn towards gradients that, you know, that's my, that's my strategy. That's my model of the world. But any, any animals that start moving around a lot, um, then they needed, they need to have some sort of concept of where they are and where things are. And so all of a sudden, if you're going to have mobility and then you're going to have to have some sort of way of learning the structure of your environment and knowing where you've been, how do you get back to some place that if you're going to a home and you get away from it and you have to come back, you have to have some way of doing that, right? And so nature evolved a series of ways of sort of building maps of your environments and different animals do the different ways. And uh, and mammals have a certain way. And that's, uh, that was first evolved in the hippocampal complex, the grid cells and the place cells, which is a really a map of the world and learning where you've been and knowing where you, what you did this morning, things like that. And then, um, and then finally, the theory that I propose here is, uh, that released from Numenta is that, um, that sort of mechanism that was highly evolved for in the, in the, in the hippocampal complex grid cells in place, cells and things. That method of, of mapping your environment, um, became, um, uh - evolution discovered they could make us sort of a generic

		version of that and make lots of copies of it. And this is, this is what we think the neocortex is. So this is the whole Thousand Brains Theory of Intelligence.
Matt:	03:57	So you are saying that the mechanism that the hippocampal entorhinal complex has done for an organism to navigate through its own environment is copied thousands of times.
Jeff:	04:06	Well, it's been refined and copied. There's not a direct copy.
Matt:	04:10	But it's running in parallel.
Jeff:	04:11	Yes, so you know our theory that came out in the paper in December on the framework paper, and a couple of papers before then basically is arguing that there are grid cell like cells in the, throughout the neocortex and that the same mechanisms that were basically used to figure out where you are - typically a rat is in some environment or we are in some environment. We have a sense of where we are. That's in the grid cells in the entorhinal cortex. Um, that, that same mechanism was sort of nature say, hey, I can make a, a sort of slim down generic copy of this mechanism, which is a mapping mechanism. It's basically learning maps of the world if you can think of that way. And um, and now it's copied to throughout the neocortex and now that, so intelligence has taken that sort of mechanism which was evolved from specifically one thing, navigation and remembering where you've been.
Matt:	05:05	Egocentric navigation.
Jeff:	05:06	Well yeah, I don't want to use those terms. It's navigating in some rooms. I have to learn like, uh, where's, where's my environment? Where whether it's a woods or my house or the rat's nest under your house or whatever it is. Um, and learning that, where things are and how to get around and navigating whether it's, and so this is complex. That system is very complex to do this. The right there is egocentric and allocentric models there in the entorhinal cortex. So that, but again, that's been the evolutionary pressure for a very long period of time. Of course, the neocortex, which we think of as the organ of intelligence is fairly new. It hasn't been around very long at all. And it got big really rapidly. And so the idea that got big really rapidly by taking a single thing and making many copies of it. And that single thing we believe is sort of a, an essence of what you see in the old brain and the hippocampal complex, the entorhinal cortex. So now we are, um, so this is a very long answer to your question, but, um, instead of building just maps of, uh, uh, you know, rooms and, and, and where you are in the

forest and things like that, we now build maps for objects and instead of moving just my body moving through the world, it's my finger moving relative to objects or my eyes, moving relative to objects. So now we started, we have started having the ability to build, uh, uh, models of other things besides environments. Sothe bottom line of all this is, is that intelligence is tied to the quality of the model that you build about the world. Right. And to say we're smart and we say, Oh, you're smart. Cause maybe we know about, you know, the planets we know about, you know, black holes. That just means we have a model in our head for those things. And,

Matt:	06:42	and the, the meat of that is in the new part of the NEOCORTEX. It's interesting that, that the neocortex spread so quickly over, evolutionarily that means that it was doing something that was really useful.
Jeff:	06:54	Yeah. The only way you can get something big like that, quickly in evolutionary terms, it's just make more copies of it.
Matt:	07:01	Yeah. Which essentially what it is
Jeff:	07:03	Yeah. So, so that was Vernon Mountcastle's big proposal in the 1970s was that the, you know, the NEOCORTEX got big rapidly. And by you look at the structure, the structure is very similar everywhere. So he says this is just, it's the same algorithm being operated on 100,000 times. You know, we have a hundred thousand or 200,000 copies of the same circuit,
Matt:	07:24	But each one, they're working together to provide a rich model of reality.
Jeff:	07:28	Well, this is the Thousand Brains Theory, which of course is in the December paper. The Thousand Brains Theory of Intelligence says you've got all these parallel models. So basically there's tweaks but they're basically doing the same thing. But they model different things. They're not all the same. All of these models are different. They're not like, it's not redundant. They're parallel. And so there are thousands of models that are modeling the visual space and there's thousands of models of modeling auditory space and there's thousand of models that are being used in language and so on. But each column, um, is essentially modeling its inputs doing of sensory motor modeling of it inputs. It doesn't know what - each column doesn't know what it's doing. But depending on what its inputs are and what motor behaviors it could control, that's what it builds a model of.

Jeff:	08:15	And some of them are going to be building models with input from your eyes and your fingers. Some of them are going to be building models of models, models of models and they're getting input from other parts of the neocortex. But they're all doing basically the same thing. And then they vote. Again, this is all in these papers we've written. Um, they vote by, um, uh, because often they're, they're observing the same object or they're thinking about the same thing. And then, and so even though they're slightly different, um, each, like each part of my somatosensory cortex represents different part of my fingers, my skin. And when I touch this cup, they're all measuring a different part of the cup. Yeah. But they're all sensing the cup. And so they can vote and say, yeah, I only have this limited input. I have this limited input, I have this limited input, but let's vote together. And together we can agree that this is a, we all know about cups, but, uh, we all just agree that, you know, we all can only agree on right now that this has to be, this cup. It can't be anything else.
Matt:	09:11	So how can we take what we know about how that works and help it to define what intelligence is?
Jeff:	09:18	Yeah. Yeah. So I think this is a, one of the, um, I think today we're in this sort of a weird state about intelligence, especially in machine intelligence, which is dominated by the idea that intelligence is a capability. Can I do something? Can I play G better than the best human player? Can I drive a car better than, you know, a human, or can I, you know, whatever, pick your favorite task, analyze a medical scan and detect cancer. And so these are, this is how we measure today, this is how the world of AI measures their progress.
Matt:	09:54	The value of the application.
Jeff:	09:56	Well, just, you know, against human level behavior, right? Take a task that a human does. Say, can I get a machine to do that task? Um, and this goes back to Turing with the Turing test, right? Uh, he proposed this the imitation game, which we now call the Turing test, right? But, which is all about matching human level performance. But that's a really poor way of measuring intelligence. Um, there's lots of intelligent things out there today, animals, who couldn't possibly pass the Turing test, they can't speak to us, you know,
Matt:	10:26	But are vastly more intelligent than
Jeff:	10:28	Yeah, a dog is pretty intelligent. Right? And some dogs more than others. And, so, you know, the Turing test doesn't capture

		that at all or a human who speaks, you know, Portuguese, is asked to take, you know, do a Turing test with someone who speaks English. It's meaningless. It's like, you know, they can't do anything.
Matt:	10:47	It's like we're framing intelligence in the wrong way.
Jeff:	10:48	yeah, we're framing it in the wrong way. So we're, we're using this sort of ends method to do it. Like, well, if you can do x... The problem with doing it that way, is that you can always engineer a solution to do x better than, than anything. Right?
Matt:	11:01	That's what we have today.
Jeff:	11:01	So we end up with, we have basically these AI systems that are really good at one thing, but they can't do anything else, nothing else. And then you could argue, well, I could take this convolutional neural network and I can train it to play Go. And then a moment later I can train it to play, you know, Pacman. Okay, fine. And then a moment later I can train it to drive a car.
Matt:	11:22	It doesn't understand anything it's doing.
Jeff:	11:25	Well, yeah. So there's a real question, let's go back. There's a couple of things. First of all, we have to ask what kind of model of the world does it have.
Matt:	11:30	Or does it have a model?
Jeff:	11:32	Yeah. Well, anything that acts has a model, it can be a very poor model. I mean, a model, basically
Matt:	11:37	It could be a hardcoded model.
Jeff:	11:38	It could be hardcoded model. The point is, you know, you, if you have some input and you're gonna act on it, then you need a model to decide how to act. Could be really stupid model.
Matt:	11:46	Could be a lookup table.
Jeff:	11:47	It could be, but it'd say, here's instructions on how to do this right? Or it could be learned. But the point is in our brains, we have this a fairly general purpose way of learning models the world. My brain can learn mathematics, it can learn Portuguese if I so desire to do so, I can learn to play a guitar. You've been playing guitar recently. Um, we can learn all kinds of things. I can learn art, I can learn engineering.

Matt:	12:12	Anything
Jeff:	12:12	Well anything we know of, but not anything. But pretty much all of the things we think about of human endeavors across all time are done using the same basic algorithm. And not only that, we do them all simultaneously. It's like, it's not like I'm not dedicated to one thing. So I have a, I have this, you've got to picture that in your head, in your head right now. You've got this model of the world. I know you're interested in this, you know, astronomy and space and physics and time and space. And then you also have a model of your family and your car and how to fix the things in your house and all these things. So we have this very, very rich model that we've learned and the single organ, the neocortex learns, can learn this very, very rich model of the world. And then if we want to ask how intelligent something is we really need to ask what is its model like and how rich it is and, and by what methods does it work? Because some models of the world are very, very specific and they in some way, let's say this way, so many things learned. It can be very, very specific. If I set up a Go playing, uh, let's say chess. If I set up a chess playing computer, I want to build the world's best chess playing computer. Um, the model may be built around chess. You may be like, oh, there's a, the the only moves this thing can make are chessboard movements. And the only thing you can know are chess board positions. And so we structure the information in this computer, in chess board coordinates. And so that's a great a framework for learning chess.
Matt:	13:37	It's the only frame of reference the system knows.
Jeff:	13:39	That's right. So we talked about, you know, grid cells, like a frame, a reference frame, right? And, but it's, it's a very general purpose reference frame. I can apply grid cells to all these different things. It's more like Xyz Cartesian coordinates. It's very general purpose. You can apply x, Y, Z coordinates to lots of different things.
Matt:	13:54	So you could say grid cells are like a general purpose tool our brain has decided to use to map things in their own reference frame?
Jeff:	14:02	Well, yeah, it's a general purpose. So, um, we talk a lot about reference frames around here. So um, a reference frame is just a way of locating something and assigning knowledge to someplace. So, um,
Matt:	14:14	in relation to other things.

Jeff:	14:15	Yeah, you need to know where things are, in relationship to other things. And that's what you need to know. Basically that's how you assemble information to something useful. Like, you know, what makes a computer a computer is because these are components that are in relation to each other and how they move relative to each other and so on. So you need a reference frame for storing structure about anything. And there's different types of reference frames. You know, a good general purpose reference frame is the one I just mentioned the X, Y, z one.
Matt:	14:40	Cartesian coordinates.
Jeff:	14:41	Cartesian coordinates. Um, that's, that's pretty general purpose. You can apply to any three dimensional or two dimensional, one dimensional structure and you can add more dimensions if you want. Grid cells are another general purpose one. There similar in that regard, but they work differently than Cartesian coordinates in a clever way. And, and uh, just a little aside here, the reason they're really clever is because there is no origin to them, but they, the locations are tied together by movement. And so it's all, it's, it's a, it's, it's like a Cartesian coordinate frame and in some sense, in that it's general purpose, but it's, um, movement decides how you get between things. So, um, but you can think of it as a general purpose reference frame. So I can, I can use a general purpose reference frame to learn chess. I can say, oh well there's a set of a board here and maybe I'll use my sort of x, y, Cartesian coordinates for that. And that won't be as good as a chess coordinate frame because I could build a specific one, a reference frame just for chess. And it'd be probably better.
Matt:	15:38	You could encode the movements of the pieces.
Jeff:	15:40	Yeah. So basically the only movement that exists here are the chessboard movements and, and so it's that system could think in chess very, very well. But if I asked it to think about coffee cops, to learn to structure the Coffee Cup, it would totally fail.
Matt:	15:53	So you're saying today's AI or weak systems are sort of hard coding their models to specific reference frames of reference?
Jeff:	16:00	Well, I wouldn't say that. I would say some might be. Um, it definitely kind of case by case basis. Um, if I, if someone built, you know, I don't know how the team that did Alphago did their thing, but if they encoded knowledge about chess boards into it specifically

Matt:	16:16	Go boards
Jeff:	16:16	Oh excuse me, Go boards, then, then it might be spiffy.
Matt:	16:20	Well they'd have to,
Jeff:	16:21	Well, no, actually it turns out most of what's going on and convolution neural networks these days, there's no, there's no actual...
Matt:	16:28	Oh, it's just mimicking or...
Jeff:	16:30	Well, there's no encode-, there's no assigned reference frames, there's no even knowledge. It's essentially has to be learned and no one really understands how it's learned. And so you end up with this sort of weird sort of mapping between inputs and outputs that no one really understands very well. Um, so the fact that there is a mapping between inputs, outputs tells me there has to be some sort of model. Whether that model has a reference frame is hard to tell. If all I'm doing is sort of saying, here's an image and here's the label, Here's an image and here's a label. I actually don't need a reference frame. It could just be this big complicated lookup table in some sense. Um, but if the system is going to um, to um, you know, make predictions and move and make, you know, and, and, and sort of say, if I do this, what's going to happen here? And so on it then you think a reference frame is required. But yeah, but we know in brains, brains have reference frames and if you don't have a generic or general purpose reference frame, you cannot be smart. You cannot be generally smart.
Matt:	17:29	Right. And you're talking about, so there's different types of reference frames. Like I like to think, there's the easy separation of, I have an egocentric reference frame. I am in the center of it and my environment is the frame of my, where my organism exists. And then there's the reference frame of every object that I could imagine has its own reference frame and then I can somehow combine them. So I can imagine an object in my hand, it's not there.
Jeff:	17:53	Let me be clear there. When you say there's different types of reference frames, there's two ways of thinking about that. One is what is the physical of the reference frame? Like chess boards versus Cartesian coordinates, versus grid cells versus latitude and longitude. Those are different types of reference frames.
Matt:	18:07	True. I'm not talking about that.

Jeff:	18:09	A reference frame can then be applied to different things. So reference frames are always, are always in some sense anchored to something.
Matt:	18:15	Yes.
Jeff:	18:16	Right. That's what you were referring to.
Matt:	18:17	That's what I was referring to, in space.
Jeff:	18:21	Even if it's not -
Matt:	18:22	- or to each other.
Jeff:	18:24	So like you can anchor a reference frame to my body and that becomes an egocentric reference frame. So that, you can think of it as coordinates related to my body.
Jeff:	18:30	And that's what we think is going on in the way of regions in the brain. Then you can have reference frames that are anchored to physical objects, meaning if the object moves, the reference frames move with it. That's what we think is going on in what regions in the cortex. Um, then, uh, there are, um, yeah, those are, those are the two big ones.
Matt:	18:49	And those working together allow us to model all the objects, sort of in our space.
Jeff:	18:56	Well, I can model objects using like what reference frames, like, you know, a reference frame around that object.
Matt:	19:01	Yeah but for me to know there's a chair there, there's a table there...
Jeff:	19:04	Yeah, for me to know where it is relative to you and how to get to it. So, you know, the, there's, excuse me, there's these well known, um, when people first discovered the what and where regions in the neocortex, these, these parallel sensory paths for vision, hearing and touch and they first discovered them in vision and um, uh, it was very odd because if some, if a human has a damaged, um, What pathway in vision, they, they can look out into the world and I can't recognize things. They can look at, you know, the computer and this cup and they will say, I don't know what that is. I can't tell you. But they know something's there and they can reach for it.
Matt:	19:44	They can actually grab it.

Jeff:	19:45	They can grab it. They don't, they're surprised when they grab it. It's kind of like, ah, and then by the way, then they know what it is.
Matt:	19:50	Yeah, because they touched it.
Jeff:	19:51	Because they touched it. So they're like, oh, that was a coffee cup, of course. Um, but the point is with, with, uh, with only the egocentric reference frame, you can know how to reach things and you can know where things are relative to you, but you cannot recognize what this thing is because you need, you need an object-centric reference frame to do this. You flip it around, the people who have a damaged Where pathway but a functioning What pathway, they say, oh, there's a coffee cup. And they say, and you say, well, why don't you pick it up? And I said, I don't know how to do that.
Matt:	20:21	They can't relate it to their
Jeff:	20:23	They don't know where... They see it. They know it's out there, but they can't, they don't know how to move their hand to that thing because they've lost this, this visual egocentric reference.
Matt:	20:31	I don't know what's worse.
Jeff:	20:35	So when we, when we say we model the world, we need to be careful here, I model the Coffee Cup with object centric reference frames. Um, and uh, I model the solar system that way and I model the, you know, the universe, you know, cause I can't actually go out and touch those things, move around. But, um, but um, but to be a functioning animal with moving limbs, you have to have egocentric reference frames. And, uh, and so, you know, so you know, you take a more primitive animal, I always like to talk about like a crocodile and you know, they see some food and they may not have a clear image of what that food is or how to exactly shape and... but they know how to reach it. Then know how to move their face towards it and bite the thing. So it's really important to be able to like move your body to capture some prey or something like that. Um, so there's these limited abilities to do these things in other animals. So these, these ideas aren't just in the neocortex that are in the old parts of the brain too.
Matt:	21:41	The interesting thing you've mentioned about the crocodile, you think about a crocodile versus a mouse and how they eat. A mouse will have little bitty fingers and it will, it will manipulate the food. Really carefully.

Jeff:	21:52	I saw this video, they're showing this just in the last year for the first time. And I didn't realize this. You think a mouse paw is not very interesting. But they, they pick it up like, you know, like the Shakespearean actor looking at the skull, you know, it's like, ah, what's this little thing? I mean, I'm going to eat it this way. And it's like a mouse does this and it's really surprising.
Matt:	22:10	A crocodile couldn't even
Jeff:	22:14	Yeah, yeah, a mouse looks at the piece of cheese and decides which way to hold it. It's pretty impressive. They do it really quickly. So if you don't slow down the video, you miss it.
Matt:	22:23	But it must mean they've got a rich model.
Jeff:	22:25	They have a model. It's not nearly as rich as ours of course. But the idea that they can see the structure of that object. What that tells you is that they see the structure of the thing they're about to eat. They can then they know, they recognize each orientation and what features are on it and then they can move their hand and grab it in the right way to bring it to their mouth in the right way to eat it. And so it's not like it's just some, you know, piece of food and they don't know what it is or what its shape is or you know, they're not just stuffing in their face. They're picking at it, like you and I might pick at a piece of fruit that's got some bad spots on it, you know.
Matt:	22:58	So while we're talking about, um, different types of animals, let's talk about the intelligence spectrum cause cause there's, we're more intelligent than mice obviously. And you would just probably say mice are more intelligent than crocodiles.
Jeff:	23:09	There's so many things are confused here. So, first of all, we can ask, um, the mechanisms by which the animal makes a model of the world. Are those general purpose mechanisms or are they specific mechanisms? Then we can ask what is the capacity of that model, because clearly a mouse's neocortex is quite small and, um, you know, it's the size of a small postage stamp at best, a rat's, not that big. So, a mouse is probably even smaller. Um, so it doesn't, it's not going to have a very big model of the world, but it learns that model the same way you and I learned that models. So it's a general purpose model, but very limited capacity, right? Whereas, you know, a, an ape, it's got a much bigger model like, you know, a monkey. And, um, and then there's humans. We have this really big NEOCORTEX, so we all learn the same way. We all use the same mechanism. We're all general purpose learners. We all learn, we learn rapidly, continuously and using sensory motor inference. Um, um, and

we have the same sort of rich, um, modeling structure. Um, but we clearly, so we have to, we can make, we can divide animals or systems along that line. What is the mechanism you're using to learn and then there's a capacity issue. So often we think about, you know, intelligence is people who know a lot of things, right? Um, well that's a separate issue. That's like, okay, well a mouse I would say has the same sort of mechanism for learning that we do. It's a mammal, so it's going to have exact same neocortical structures, or very similar, uh, but it's just limited. Uh, sometimes I make the analogy with computers. What is a computer? Well, there's a formal definition for a universal turning machine. It's a mathematical thing, whether something's a universal Turing machine or not, and you can build them out of tinker toys. You can build them out of silicone and, um, and all computers that we think of today as computers are universal Turing machines. Not all, you can build an ASIC chip that does something, it's not even a universal Turing machine, but computers are. But computers come in all these different sizes, right? You can get the teeniest little computer, still has to have a CPU Still has programming and maybe only eight bits, uh, has limited amount of memory.

Matt:	25:25	Powers your toaster.
Jeff:	25:26	Powers your toaster. Um, and that's still a universal turning machine. Um, and then you have the room size.
Matt:	25:33	That just means you can program it to do different things.
Jeff:	25:35	It means it works on the same principles of general purpose computing. If you gave it enough memory and enough time, it could compute anything.
Matt:	25:42	General Purpose is keyword there.
Jeff:	25:43	Yeah. Right. It's basically, it has this set of attributes, which in theory could solve any problem if you gave it enough memory and enough time.
Matt:	25:53	But there's this huge spectrum between toaster and supercomputer.
Jeff:	25:56	That's right. But they are all universal Turing machines. And there are other systems that are not like that. There are other systems that solve problems. I could, I could have a toaster, my toaster. It's funny cause I think it's really amazing. But my older toaster wasn't so amazing. It had some mechanical things that

let it figure out how to toast much. And so it did sort of the same functions, but it wasn't as good, but there was no computer inside of it. Right. So, um, but it didn't many of the same things. So, um, uh, so anyway, so you have this one spectrum, which is the mechanisms that are being used and in the computer world, that's universal Turing machines. And then there's assessment section of capacity. And then there's a third thing which is like, okay, what does it been programmed to do? So now we can have, we have three similar metrics on intelligence. We have, what are the basic mechanisms by which the model is being learned. So a mouse and a dog and human, all of the same mechanisms.

Matt:	26:49	And we're talking about reference frames.
Jeff:	26:52	We have a general purpose reference frame. And we, and we, um, the whole mechanism that we use for building a model of the world is what we wrote in the frameworks paper. Yeah. And so we all, we just all use the same one. Um, then there's a capacity issue. So mouse has very limited capacity. We have much more capacity because we just have bigger brains. And then finally the equivalent to what the programming is, is like, well, what have we learned? I could have a big human brain and not have learned very much. Right? I mean, maybe I just didn't get an education.
Matt:	27:20	I didn't think about that very much. That aspect of intelligence. You can have a system that is potentially intelligent, that hasn't learned enough to be classified as intelligent.
Jeff:	27:28	Well, when we think about intelligence, we have to think along these metrics because the same thing with computers. It is those three metrics for computers, the method which operates the capacity of the system and then what has been programmed to do and it tells us we have the method of learning a model which is reference frames and general purpose and sensory motor inference and all those things. Then because they can vary their, so maybe the Go computer doesn't have a general purpose reference frame. Um, then there is the capacity of the system. How much memory does it have? How big a brain, how many columns do I have, how many neurons, how many synapses, things like that. And then there is what is, what is it been trained to do? What has it learned. And so you and I can be very intelligent people and have very different sets of knowledge about the world. So if I had never been, if I was raised in the woods by wolves and you came along and said Jeff, what do you think about, you know, about the Milky Way? And,

and I said, what the hell are they talking about? I don't know what the hell the Milky Way is.

Matt:	28:26	They still talk to me about the spirit in the sky.
Jeff:	28:28	Yeah, whatever. Or I'll talk to you about these plants that, you know, you don't know about these plants? You know.
Matt:	28:32	Don't eat them!
Jeff:	28:32	You don't know that? So, so we can have, then we have to, we have to be really careful. And I think the problem we have in the field of AI is that people aren't talking about it in a structured way like this. They're not talking about these different components of intelligence
Matt:	28:47	How something is intelligent?
Jeff:	28:49	Yeah. These three components: the method by which it works. Is it general purpose or not, do we have a general purpose learning algorithm for learning models of the world or not? Then your paths through the system and then what it's been trained and how it learns. Right. And so right now we're focused basically on, hey, can I beat some humans at something. Which is really the wrong metric completely. If I'm going to, I think the world of AI in the future, if you go far enough, it's going to be dominated by machines that don't do things humans do at all. So why, you know, why try to recreate a human for all these things, you know. Let's have them do things, there's some overlap perhaps, but it's not the goal to pass the Turing test. That's, that's the wrong goal.
Matt:	29:31	Um, along those lines, I want to talk, we just talk about the intelligence spectrum. I would assume that we humans are probably near the top of that, intelligence spectrum on this planet.
Jeff:	29:40	Well on this planet. I would say, I would say, yeah, there's no question about it.
Matt:	29:44	So one of the things that we do that perhaps intelligences down the spectrum, don't maybe is have this ability to have these rich abstract models of things that don't actually physically exist in nature. How does the idea of reference frames apply in that arena?

Jeff: 30:01 Yeah. Um, we touched on this in the frameworks paper and I'm writing a lot more about it right now. Um, uh, first of all, the way I approach this is not saying like, oh, I see that these are references. It's more like, no. The neocortex does these things. And neocortex looks the same everywhere. Therefore these things are based on reference frames. It's like, let's figure out how as supposed to starting with saying, you know what? I think language is based on reference frames. No, it was like, dammit, language must be built on reference frames. So let's think about it a bit. Um, so, boy it's such a big question, Matt. There's lots of ways you can attack it. Um,

Matt: 30:40 We can use examples because I, I'm a programmer and so I think in different ways in different engineering situations and that feels like, for example, if I tackle a programming problem in one language, I will pull a certain reference frame into my brain to execute commands in that language. It has similarities to other languages that I've used in the past. Like all these different areas of expertise.

Jeff: 31:04 Let me get really, really, um, uh, highbrow about this. So it's kind of hard to bring it down to like everyday experiences, but, knowledge. I'm just trying to define what knowledge is. Knowledge is sort of information or facts arranged in a useful way.

Matt: 31:26 Correct. Agreed.

Jeff: 31:26 Okay. So I can have a whole bunch of facts and I can say, yeah, you and I see the same facts, but if I'm knowledgable about it, I look at those facts and say oh, I know those things. And what does it mean to know those things? What, what I believe it means to know those things, is that I take that information and I can assign it to a reference frame. Everything has a location in that reference frame and what you get with a reference frame is you get the ability to navigate through these facts and navigating between the facts is like navigating through a room or moving my finger relative to the Coffee Cup. But in this case, I'm moving my location in this space of, of facts or space of things, um, to achieve certain results. So an example I've often used is a mathematician where you say, okay, some people look at mathematics equations like, Oh my God, it's like Greek. I can't understand a word that's going on. But the mathematician looks at it and all these equations are friends. They're like they're friends. So like, I know that and I know that. And I know that. Even numbers are friends, you know? Oh yeah. You know, 169 that's 13 cubed. So, and what they do is, and so the one thing you think about is like if I'm trying to do is I'm trying to solve a mathematical theorem or something like that. I start with set of uh, points that set of equations or set of things I know are true, in some reference frame, and I'm trying to get to a new point in some space and I'm trying to figure out the right behaviors to get me there. And the behaviors are mathematical operations. So I do a plus transform here, I get this and if I do this kind of transform here, I get that. And if I do a multiply, I get this. And so in this sense where the behaviors are mathematical transforms, that move you through a space of mathematical constructs or concepts.

Matt:	33:04	Like a theoretical space?
Jeff:	33:05	It's literally using grid cells. Um, but the space is, there's two things. Think of it like, grid cells are like a map. Okay. N Dimensional map. So what do you assign to the map locations? You assign mathematical constructs. You can think of like equations. Okay, when you move, what are you moving? We're not moving physically, but you are mentally moving through this space by applying certain operators instead of moving the, instead of the operating being flex this muscle, the operator is multiply. It's like, yeah. And um, and you get the same basic thing. You're moving through this space of ideas. And so that's what you do when you're thinking about something. So this, so this is the general idea, that high level knowledge is the organization of concepts in a reference frame and, and your ability to be really, really smart at something is knowing how to navigate and what kind of behaviors you do to get some place. Another example would be a politician. You know, a politician wants to get some bill enacted. They have something they want and they have all these obstacles to get there, right? So they say, I'm at this point here where this was the state of the world and I want to get to this point here where this thing gets, you know, gets passed. How do I get there? And they have all these potential things they could do. They could have a rally, they could hold a forum, that could do some sort of publicity, they could get some endorsements and they, and then politician's experienced enough to say they know what will happen when they do these different behaviors. Um, and so since your choice of behaviors is your choice of which way you want to move through the space of political things.
Matt:	34:43	So his model of politics as he knows it is, is a reference frame that he's evolved over time due to his experience in that field. Right?
Jeff:	34:53	More likely a woman because they're smarter about these things. But um, um, yeah, what you gained, what'd you, what

the comes an expert is an expert can look at the same facts, but they have through experience, they've learned how, what happens when you do things. So as a programmer, you sit there and you said, ok i want to solve some problem and you say, well, I might solve it using this kind of sorting system. I want to use this kind of structure or this framework I want even this tool. And you as an expert know exactly what would happen to do those things. It's very equivalent to being an expert in the woods. You know, I go out for a walk with you and we say, you know what? We need to find the source of water. We're getting really thirsty. If I spend a lot of in the woods, I'll say, yeah, I understand all this. I bet you I know there's water down there because I learned these kinds of patterns before. Yeah, and so I recognize the scene as, yeah, I understand this. I see these trees and I see that kind of tree and I see these hills over there and I hear these sounds. To me, I understand this stuff. I have a reference frame.

Matt:	35:58	You may not even have to think about it.
Jeff:	35:59	No, you have a reference frame. It's just like seeing. When you see this coffee cup, you don't say, Oh, where's my reference frame? You say that's a coffee cup. Well, a person who spent a lot of time in the woods will look at that same set of data that you look at and understand it, and they'll say, oh yeah, well I know how to do that. I know how to get to the water because water is typically going to be at this kind of thing down there, right there and you'd be like, How did you know that? And uh, and same thing, if you've never been exposed to something, an everyday object like a coffee cup and you were raised in the woods and never saw any kind of physical thing like this, you would might look out and go, ah, I kind of see what it is, but you wouldn't immediately go, that's going to be great for carrying the fuel we need to carry. You know, or whatever. Right? Um, so that's what makes an expert an expert is someone who has spent enough time looking at experimenting with various pieces of facts or data or concepts or whatever they are, that they've learned the correct reference frame or a good reference frame to assign these facts.
Matt:	36:50	One that's close to reality.
Jeff:	36:53	Well yeah, what makes it reality is that the reference frame is accurately predicting what's going to happen.
Matt:	36:58	I guess that's even more, that's more important to accurately predicting ...even better than reality.

Jeff:	37:04	You know, different people could take the same set of facts and organize them in a different reference frame. And we see this a lot. Um, that doesn't happen for everyday objects as much, but we see it for things that are conceptual and so we can, uh, you know, uh, different religious beliefs take the same facts about the world. We all observed the same thing and they come to completely different conclusions about what would happen when they do certain things. And, um, and that's, that's a case where, um, you know, the reference frames are reality, right? They can't all be correct. Most of the, you know, all the different beliefs we have in the world can all be that conflict with one or they cannot all be correct. So one could be correct. Probably most of them are wrong. All, but, um, but the point is, um, uh, you know, that's, that's the test of reality is or test of the accuracy of the model is how well it predicts the future. And, um, and when we don't really have good data about the future, people can get wrong or if they limit the amount of exposure to what they see.
Matt:	38:02	Yeah. You know, or, or just the data being given to the intelligent system is biased in any way.
Jeff:	38:08	Yeah. Yeah. So you see someone who believes the earth is flat and you give them all these facts that don't fit that model there. Earth is flat predicts certain things, right. And, and they will do their damnedest to try to fit these new facts into their model, no matter how - Otherwise they're going to have to throw the whole thing away and start over again, which is a really disconcerting thing.
Matt:	38:31	It is. It is uncomfortable to throw away a whole frame of reference.
Jeff:	38:35	Basically, the models, you build of the world are your reality.
Matt:	38:38	It's your belief system.
Jeff:	38:40	It's your belief system. It's your reality. This is what, you know, that's the, that's what it boils down to. There is no other reality in terms of your head this is is. And so if you say, well, you know what? Everything you believed, these reference frames, and you have to start over again. Um, then it's kind of like, oh, you're going back to square one.
Matt:	38:54	In a way. I mean, you think about how these are represented by neurons in your brain - a reference frame you can think of as a physical structure. I mean, it is caused by physical structure and

		if you have to tear it down and build it back up, it could be painful.
Jeff:	39:07	Well, you wouldn't be taking away the neurons, but you'd have to redo all of them. It wouldn't be physically changing, but you'd have to redo all the connections, the synapses.
Matt:	39:15	Sure. But you'd have to ignore all of the predictions that you're getting that are telling you it's flat. It's flat, right?
Jeff:	39:21	It'd just be like, you know, starting over.
Matt:	39:24	That's hard because you have to drop your beliefs and that's the hard part is dropping what you think, you know,
Jeff:	39:30	I mean nature has designed us to, to want to build models of the world. That's what we do. That's the first number of years in our life.
Matt:	39:37	You feel rewarded for it and when you realize something works a certain way, you feel good,
Jeff:	39:43	And there's a evolutionary survival advantage of that. So like we can learn new models of the world. Each new child that's born gets to learn anew. And um, and there's an advantage to that. That means we can adapt very rapidly and we adapt during our lifetime. But if you build, once you've built your world model, um, and you've been living your life around it, uh, then someone comes along and says, the entire thing is wrong or some portion of it's wrong. It's, uh, that's, that's very upsetting. That'd be all of a sudden you feel vulnerable. It's like I have a model that makes me successful in the world.
Matt:	40:19	And in a way, it's part of you.
Jeff:	40:22	Yeah. So now, now it's like, it's like feeling lost. You know, if you, if you're, if you're out in the woods and all of a sudden there's, you can't anchor yourself in the woods, every tree looks the same and now your reference frame is lost. You literally, if you cannot locate yourself in the reference frame, you feel lost. That's not a good feeling. That's an emotional thing. And that's how people feel, people who don't like math, they look at these math equations, they feel lost. So if I, if I was trying to understand someone speaking Russian, I would be lost. Um, I always feel like oh my god, I can't do anything. Um, so, um, that's an uncomfortable feeling.

Matt:	40:56	And so you're always grasping for anchors like this looks somewhat similar.
Jeff:	41:00	Yeah, that's right. So we have these reference frames. You have all these models of the world. We want to stick new data into our models of the world. Right? And by and by, if you have to abandon one of your basic models of how the reality is, then you feel lost. And that's an uncomfortable feeling and you don't know what to do and you don't know how to act. And it's just like being lost in the woods.
Matt:	41:20	So that's great. Let's come up with those three things that you're mentioning before that sort of that we're saying define intelligence. So Reference frames.
Jeff:	41:27	Well I'd say the method by which the system learns.
Matt:	41:30	The method.
Jeff:	41:30	Yeah. Which is going to, to be intelligent, you have to have a general purpose reference frame, right? So grid cells are general purpose reference frame, right? There are other reference frames that are not so general purpose like latitude and longitude or more specific or chessboards and things like that. So you have to have a general purpose reference frame and you have to, you have to of course build a model of the world around using that method. So there's a lot of things involved in that - sensory motor movement and continuous learning. And so on.
Matt:	41:54	But by building a model, using the general purpose reference frame, you were saying we're going to have a general purpose model.
Jeff:	41:59	It's like saying that that's square one. To be truly intelligent, you have to at least have the mechanisms. It's like the universal turning machine. You have to have the basic substrate, which is can learn different things. Yes. Do Movement. Yes. Okay. And the next thing is the capacity of the system. That's a simple one. Um, you know, you can learn more and build bigger models.
Matt:	42:24	That's what a lot of the work is happening today in Ai is still in capacity or ...
Jeff:	42:31	But if we're talking about intelligent systems, um, then, uh, then of course you say, why is it, uh, the dog smarter than a mouse and why am I smarter than a dog? It's mostly to do with the

		capacity of the system. Not completely, but mostly. Um, and then, then finally is what does the system actually been trained on? You know, again, you could take the world's largest supercomputer and have to play tic TAC toe.
Matt:	42:55	Won't be very smart.
Jeff:	42:56	Well, it's, that's fine, but you know, but I can't take, you know, the world's largest weather simulation and run it on the computer on my toaster. It's, I can't do it in any kind of real time that'd be useful. So, um, so those are sort of the three metrics and we just confuse them all the time. Um, and, and I think even worse than that, we confused them because they only look at most AI looks at what humans do, right. Only what humans do. And then even then, like passing the Turing test brings us into a whole other domain, which is like, well, now you're trying to emulate the emotional capabilities of a human. And my definition of intelligence we just talked about it does not include emotions. It doesn't mean humanlike. It doesn't mean, you know, it's more spotlight then, then, you know, you know, I get that.
Matt:	43:40	Right. I think he was like smart without emotion. Something like that. Yeah. Okay. So think of Spotlight, uh, you know, there was this, you may be, you want to build a machine that mimics a human and it shows emotions, but that to me is not intelligence. That's a separate problem. It's a separate issue.
Matt:	43:59	But something that could be additional
Jeff:	44:02	That's a flavor, you know, you can add to tack on the... if you want to tell them to be a caretaker, a robot, maybe you want that thing to have some emotional states
Matt:	44:13	or be empathetic,
Jeff:	44:15	But I don't think that's intelligence. That's part of the human experience. But if you think about the neocortex itself, then the emotional centers are not in the NEOCORTEX.
Matt:	44:24	They contribute to all of our models of things in different ways.
Jeff:	44:27	Yeah, they contribute to how we learn. Yeah. It's a longer topic, but um.
Matt:	44:32	We'll do another talk about that.

Jeff:	44:34	But basically think of it like we had this super and sort of Spock like thing on our heads, the necortex. It sits on top of a very emotional older system. They interact. Um, but you can think about the three causes of our brain that's just the NEOCORTEX. That's really, even though it's influenced by what we learned. It's influenced by emotions and so on. It is actually an emotionless model.
Matt:	44:54	Sure. It's driven by its input.
Jeff:	44:56	Yeah. It's basically what is the structure of the world and can I discover that structure of the world, whether it's physical objects or, or mathematics or you know, politics, whatever. I'm going to try to learn the structure of the world in an actionable form so that I can now see new things, place them in this model and know how to act related to them and so on. Um, but those, those are the three basic things and we don't, we just don't have that conversation enough in AI. AI is dominated by how well did I do on this particular task. And if I, if I was trying to do an image classifier AI system, all that matters is you got the best performance. That's it. It doesn't matter how you did it, no one, it's important, but that doesn't, you get credit for being the best. And so people are just, you, you're not being, we're not sort of moving in the direction of saying well how general purpose is it? And uh, what else can it learn? And how did it learn it? And um, and all those things. And so we're not asking these questions. No one's sitting around saying,people do point out that the chess playing computer doesn't seem to know anything else. But they don't ask WHY. Why doesn't it know anything else? Cause it's got a chess reference frame and that's all that's ever been trained on and it probably can't learn anything else. Anyway. That's, it's not a simple one binary answer that those three things that all have their separate dimensions. I think that computer analogy might help because that has the same dimensions.
Matt:	46:21	Well I think that was a good discussion.
Jeff:	46:23	That was fun, Matt.
Matt:	46:26	Okay. Thanks Jeff for sitting down with me for the podcast.
Jeff:	46:26	It's great. I hope, I hope people find it interesting. I don't know.
Matt:	46:30	I'm sure they will.
Jeff:	46:32	All right, thanks.

Matt:	46:40	Thanks again for listening to Numenta On Intelligence. This was a conversation on defining intelligence with Jeff Hawkins. I'm community manager, Matt Taylor. Tune in next time.