



Episode 1: Research Update with Jeff Hawkins – Part 1

Released: July 18, 2018

Matt:

[00:09](#)

Welcome to Numenta Intelligence, a monthly podcast about how intelligence works in the brain and how to implement it in non biological systems. I'm Matt Taylor. Today I'll be talking with Jeff Hawkins about the latest research Numenta has been doing with regards to grid cells and hierarchical temporal memory or HTM. While this is the first episode of the Numenta on Intelligence Podcast, Jeff and I will be going very deeply into the theory very quickly. Therefore, I prepared in the show notes a bunch of educational resources so you can learn at your own pace about HTM in grid cells. I also broke this episode up into two parts, so it's easier to digest. If you liked this conversation and you don't know how HTM sequence memory works, I suggest you watch through the HTM School videos on YouTube. One of the main ways we share the research Jeff talks about in this episode is by giving talks at various academic events and workshops and we've had a busy summer doing just that. Our research team spent some time at the CNS 2018, which stands for computational neurosciences. The event was held at the Allen Institute and the University of Washington. Our VP of Research Subutai Ahmad, gave workshop presentations and a couple members of our team joined him to present two posters as well. Like we do with all events, we make the presentations and posters available afterwards. So visit numenta.com if you'd like to check those out. So without further ado, I'm Matt Taylor

and here with Jeff Hawkins. We're at the office. Jeff is the founder of Numenta and he's my boss.

Jeff: [01:44](#)

We work together. Hi Matt. How you doing?

Matt: [01:47](#)

Uh, so I wanted to talk to you specifically about some of the newer stuff that's going on in, in HTM research here at Numenta. Last time we talked we did the HTM chat video, that was right around when we released the layers and columns paper. Uh, so it would be great if I could get you to elaborate a bit on to that whole location idea in the brain and what have we learned over the past year or so about location in the brain.

Jeff: [02:11](#)

Okay. Well that's very exciting. So let's just go back a little bit to the time you talked about a year ago, so we uh, well it was last October, released what we called the Columns Paper, right? And that introduced a really big idea and that is that everywhere in the neocortex is a location signal. And so when we think about how the brain processes information, if you think about just how the sensory input comes in, that's only half of it, right? And the paper we published in October talked about some of the consequences of that. And we also said we didn't know where this location signal was coming from or how it was generated. We knew where it was, but we didn't know how it was generated because it's a kind of an odd idea. But we knew that there was another part of the brain called the entorhinal cortex which has a location signal. And we said, hey, maybe it's using the same mechanisms. So we put that in that paper. We said we think that it may be the same mechanism used by the neocortex, things called grid cells, which many listeners may have heard about.

Matt: [03:08](#)

Right.

Jeff: [03:09](#)

So since that time we've been exploring that idea. It's, um, it's certainly true. First of all, there's been other experimental evidence suggesting there are grid cells in the cortex. So that's something we would have predicted.

Matt: [03:22](#)

There's a lot of research in the field right now. Yeah, well grid cells are a very hot topic. It's been a, it's one of the few places in the brain when people have made a lot of progress. We started thinking about, okay, how do they play in the cortex and what kinds of things can they do? And we've discovered several major new ideas related to grid cells, which no one knew about. Um, and uh, we think they apply probably in the old parts of the brain, but they certainly apply in the neocortex. So it's, it's actually quite exciting. We're showing all of a sudden that a

whole bunch of things are coming together that, uh, we kinda knew what has to happen somehow and now when we study grid cells in the cortex, it all sort makes sense. Um, so we can go through some of exactly what some of those things are if you want.

Matt: [04:07](#) Well, the interesting thing to me is all of the grid cell literature from the Nobel prize days and everything is all about egocentric location of an organism within an environment.

Jeff: [04:17](#) Yeah.

Matt: [04:18](#) So what's the difference in how we're trying to apply that in the neocortex?

Jeff: [04:21](#) What we think is - the big idea in the paper I'm working on right now is the following: the old part of the brain where grid cells are found in the entorhinal cortex, they basically, as you said, a say where an animal lives in like if you're in an environment or in a room, something. As you move around, these grid cells modify their activity that reflect where you are and this is how an animal and you, you too, you have an innate sense of where you are, even in the dark, if you move, you know where you are and these cells are updating as you do that. And the whole point of the old part of the brain is sort of to map out your world and know where you are and how to get back to someplace. It's like, oh, how do I get back to the kitchen now that I had been wandering down this hallway.

Matt: [05:01](#) A useful skill.

Jeff: [05:01](#) Yes, it is. Um, that same mechanism which was evolved a long time ago because the animals needed to know where they were, is now been applied to two different tasks in the neocortex. One of them is to map out objects in the world. So imagine this: you're an animal. You are and I am - humans or rats, whatever. And as we move around, we have a location in the environment and we sense things in that location. And we build up a map of the world by moving around, knowing where we are and what we sense.

Matt: [05:31](#) You mean our body has a location.

Jeff: [05:31](#) Our body is moving. And, and as you move, the grid cells says, Oh, I moved over here and I moved over here. So it has this interesting way that it's represented, but we'll leave that aside for the moment.

Matt: [05:40](#) Sure, if you want more information there's links in the show notes to details.

Jeff: [05:43](#) Okay, that's great. And so the same basic approach can be applied to figuring out what some object is like a computer or a Coffee Cup or telephone. Imagine now that I have my finger and I'm moving around and as I move it's mapping out my location and it's discovering the structure of an object the same way it discovers a structure of a room. Now we had this in the paper in October, but we didn't know how it was grid cells. Uh, so let me tell you a couple of things that have really come out of this since. Uh, so if you're following me so far -

Matt: [06:14](#) Well let me, but I am sort of. Well, let me just point out, you're talking about the finger at being a sensor as if it were sort of like the organism in a room. The entorhinal cortex - flipping it a little.

Jeff: [06:25](#) Yeah. Imagine research is mostly studying rats and rats running around some maze or environment. And as that moves, it knows where it is. We're saying right now that every part of your body, let's just talk about your somatic body or your fingers and your hands and your skin and so on. Every part of your body, actually, when I touch an object, those parts are touching objects, are like little rats. They're all individually knowing where they are scurrying around this Coffee Cup. I'm holding the Numenta Coffee Cup in my hand right now.

Matt: [06:52](#) Interesting visual. And my fingers are like little rats scurrying around all simultaneous and it's like I could have five rats in the room at once. Sure. Yeah, yeah. Um, and they all work together figuring out, hey, what is this thing we've got here? And your eyes do the same thing, which is a little hard to imagine, but the parts of your retina are looking at the Coffee Cup as well and saying, Hey, I know I'm seeing this feature at this location. I'm seeing this feature at this location and so on. So now we have a very concrete mechanism. The grid cell mechanism for how that signal is generated, how the location signal is generated. But now let's talk about where we went from there.

Matt: [07:23](#) Okay.

Jeff: [07:24](#) One big insight we had, um, and this came out of Marcus who was one of our team members here.

Matt: [07:30](#) Marcus Lewis.

Jeff: [07:31](#) Um, he, so that you could do, you could take two locations. This is a little tricky now. Imagine I'm - and this is the thing we did in the paper I'm writing now. I have the Numenta Coffee Cup in front of me right now. And you can imagine this generic White Coffee Cup with a Numenta logo on it.

Matt: [07:49](#) Yes.

Jeff: [07:49](#) Now the Coffee Cup is something I know. And then the Numenta logo is something I know. How is it that I learned that the Numenta logo is on this cup? What does that mean? How do I learn a new cup that has a logo and a cup? And what we showed is that each object, the Coffee Cup and the Numenta logo and every other object in the world has its own sort of - it's like its own environment and it's its own sort of space. This is one thing that we learned from grid cells that uh, every environment has its own space, its own physical mapping.

Matt: [08:19](#) A dedicated -

Jeff: [08:21](#) A dedicated set of locations that say this, it's not like, I mean location x and y one and two. It's like I'm in this location in this room and that's unique to that - another room and every other location.

Matt: [08:33](#) Can we call it to like a reference frame, a coordinate frame, something?

Jeff: [08:37](#) It's sort of like that. Although it's more like, it's just unique.

Matt: [08:42](#) I don't like using the term coordinates, but it's some type of unique frame of reference.

Jeff: [08:47](#) Yeah. It is, it is. It's just every. Yes, but it's different than we learned in high school about -

Matt: [08:53](#) You wouldn't represent two objects in one of these things, right?

Jeff: [08:56](#) Each object has its own space. There's like every point in space is unique to a particular object.

Matt: [09:00](#) This is a point that I really want to hammer in on because it introduces something that was counterintuitive to me initially when I first understood that. Because I thought one of the things about this theory is that every cortical column is attempting to represent objects on its own. And it has an idea

of what object, uh, is, is the sense that it's getting, what that might be depending on its location in sense. And, and, and the idea that each one of these cortical columns has a completely unique frame of reference for each object was counterintuitive to me because I thought, how do they compare objects? Right. Does that make sense?

- Jeff: [09:36](#) Uh, well, up until the point you said, how do they compare.
- Matt: [09:38](#) Yeah. Well, then correct me then because that's always been my point of confusion. Like I would think that you'd want these cortical columns to be able to compare with their neighbors. What are you sensing? What am I sensing? Why don't they use the same reference frame?
- Jeff: [09:50](#) They, uh, well they do compare to the neighbors but not at that level. It's like each column is saying, uh, I think I'm looking at something here and these are the things I could be looking at. I might be touching a coffee cup, I might be touching a telephone, I might be touching a pen and they communicate at that level.
- Matt: [10:09](#) Oh, not their spaces.
- Jeff: [10:11](#) Not the spaces. So I'm looking at space, you know the location, that might be a location on the cup, but it might be location on a pen, so on, but the representational location is going to be unique for that particular column. The other column, it's a little bit more complicated than this, but the basic idea is they really can't share locations. Each one learned its own model, but they can share what they think the model is like what the output is.
- Matt: [10:34](#) So they're still sharing. They're not sharing at the very low level.
- Jeff: [10:37](#) Yes. We had this, uh, if the listeners read our paper last year, we had this idea in the paper that the columns are sharing in layer three, we still have that. What's unique now when we didn't understand back then is that the locations themselves are unique as well. We didn't have that in that paper. That's something we learned from grid cells. So alright, we're off on a tangent or maybe not.
- Matt: [11:02](#) Yeah, I threw you way off.
- Jeff: [11:02](#) So we start with this idea that every column knows the location and it's sensing something. Now we've, now we've got this idea that every column, the actual locations are unique per object.

It's a detail that's not super important at this level of discussion right now, but it makes a big difference in mechanisms. What we've now learned is how is it I can represent an object as not just a set of features but as a set of other objects?

Matt: [11:26](#)

Compositionality.

Jeff: [11:27](#)

Compositionality. So I don't want to learn - Again, looking at my Coffee Cup, I don't want to have to learn this and the Numenta logo on the cup. I don't want to have to relearn a Numenta logo. I want to be able to say, hey, there's a thing I already know called the Numenta logo and it's here in reference to, it's relative to this thing. Yeah, so what we've come up with is a mechanism and we think this is happening everywhere in the brain. Now. It's a pretty big idea that you can take two objects which have their own spaces. That's an important detail, but the point is you can take these two objects and we can define another, a sparse representation, another cellular representation which says, Hey, I represent the logo on the Coffee Cup at this location. So it's an associative link, a link between two existing objects and it's very efficient. It's extremely efficient

Matt: [12:18](#)

It essentially places an object inside the reference frame of another object, right?

Jeff: [12:22](#)

Basically what it does is it ties them together in a way that you could say these two are at some position. It's not inside of it. It's just like a link between.

Matt: [12:32](#)

It's just a wink, like in programming. It's like a link. So you can follow it. And it's very efficient to follow that.

Jeff: [12:37](#)

Yeah, very efficient. So now objects in the world are built of other objects. Um, and so that's how we build optics. So that's how we've learned structure.

Matt: [12:45](#)

That makes sense because you think about a car, you just think about a car, you don't think about all the components that are composed and - you CAN dive all the way down.

Jeff: [12:52](#)

That's right. But when I learn a car, it has wheels. I don't have to relearn the wheels and the wheels have tires. I don't have to relearn the tires. And the tires have valves. I don't have to relearn the valves. All of these things are sort of connected by linkage through this associative linking mechanism.

Matt: [13:06](#) And you just learned the links?

Jeff: [13:07](#) And you just learn the links. So, and when I'm forming, if I see something new, I say, oh, what are its components, it's got this and this and this, and it tells me where they are arranged relative to each other. Um, and now I have a new object very efficiently represented which has all the knowledge of previous objects. So this, we now understand the neural mechanism for this. And we realized this is the basis of many things. So for example, objects have behaviors.

Matt: [13:33](#) Yes.

Jeff: [13:33](#) They move, they change. Think about a car. Car Door opens, it closes, the wheels turn, the, uh, the steering wheel turns, you push a button and the radio turns - on all of these behaviors. An example I'm using in the paper I'm writing is a stapler. Simple object. But a stapler has some look doing it. Oh, I know what that stapler looks like, but it has behaviors. If you push it down, a staple comes out of it. If you pull it up, it opens and swings open. You can put new staples in it. It's still a stapler. So the, what we now understand the neural mechanisms by how the brain represents not only what the stapler is and what it looks like and it feels, but also how it behaves. And so if you're, if you're following here, and I just told you in a moment ago that an object is composed of sub-objects in a certain arrangement. So a stapler, you can think of it as, oh, it has a bottom of the stapler and the top of the stapler and they have a certain arrangement most of the time. It's hinged, and it has some arrangements, at a slight angle. Um, so that's the arrangement between those two parts. Now I have to imagine the staplers just defined as two things. The bottom of and the top. And so I have one, a sparse representation which represents the bottom half of the stapler to the top half of the stapler. Now as the top half of the stapler moves, imagine I'm opening up and it swinging through about, you know, maybe 120 degrees, something like that. As you do that. So the relationship to the bottom, the, the, the position of the top to the bottom changes, you go through a series of sparse activations representing the top relative to the bottom in a sequence. So as I move it, it starts off like, oh, the top has this position relative to the bottom, and as it starts moving up, that position changes and position changes. And so now I can represent the behavior of the stapler, opening up as a sequence of sparse representations. And so I can just learn a sequence using our sequence memory. And I've now, not only do I know what the stapler looks like or feels like, I now can, I know it's behavior and what it's, what it is, as it's opening, I know where it's gonna

go and what's going to happen to it. But I can represent all of that very efficiently in a single column in the cortex. Now, all the columns are doing this simultaneously. We now have a way of using, um, uh, moving sensors like your fingers, your eyeballs to learn the structure of an object. The structure of the object is really composed of other objects. We can do this very efficiently and now if the structure of the object changes because the pieces, the individual pieces move relative to each other, now we have a way of representing that behavior of an object. And the way this works is that it can be what an engineer would say. It would be, um, hierarchical construction or even a reentry construction, meaning I could, I could define an object, like a Coffee Cup as having a logo on it. The logo itself could have a coffee cup on it. So that's like a reentrant coding screen recursive. That's right. The better word, recursive. And um, this is a general idea and so we are going to represent all knowledge in the world this way, including things that are sort of language. It has these recursive structure.

- Matt: [16:39](#) This touches on a core misconception and I always had that hierarchy in the brain was required to recognize hierarchies of objects, structures in the world. Now what you're saying here, flips that a bit.
- Jeff: [16:53](#) Yeah. I wrote that too. That was in the book On Intelligence.
- Jeff: [16:56](#) Yeah. I said that too.
- Matt: [16:57](#) Probably why I thought that, Jeff.
- Jeff: [17:01](#) Uh, you know, life goes on, we learn, right? It's more, but composition, the reality is when you're just talking about associations, even if they're recursive, there's no hierarchy required.
- Speaker 3: [17:12](#) There's no physical, physical hierarchy required. When I look at the car, I can, I can walk through the associative links so I can say car, car door, car -
- Matt: [17:24](#) Forever. it's not memory intensive.
- Jeff: [17:25](#) But I'm not doing all of it at once. So I can't see all the parts of the car at once. So I can follow links and then I see a door handle. I said, oh, the door handle reminds me of, that's part of this thing over here. And so, um, yeah, we got that wrong about the hierarchy. It's not completely wrong.

Matt: [17:40](#) What happened? I mean it makes more sense this way honestly.

Jeff: [17:43](#) Now in hindsight it's completely obvious.

Matt: [17:45](#) Isn't that the way science is?

Jeff: [17:47](#) Yeah at least to us, I think it's totally, to most people now it's, this would still be very a foreign concept. Um, so part of our job is to document this stuff and promoted and get papers written about it and make it clear. I know there's a lot of communications on our forums that people think talking about our paper that came out last year and they, they're asking questions about things we're working on right now. So we really we're, um, we need to get this done. It doesn't end there by the way, uh, this now that we know that grid cells are sort of in the cortex doing these things, it explains a whole bunch of other stuff. So it also explains in the brain how it is that you know how to move your limbs from one point to another. This is something that's associated with what are called where pathways in the brain. So it's not just about object modeling. It provides a framework. The grid cells and locations provide a framework for and explaining everything that cortex does.

Matt: [18:40](#) The way I think about it, forgive me if I get this wrong, but when I understood what grid cells really were, it's when you build up a model of grid cells, as you learn as a baby or any organism does, you're building up a model of space itself, right? You're not just learning about the spaces that you're in right now or those environment. You're just learning to map space with your sensors.

Jeff: [19:01](#) Yes. Uh, well, it's just sort of space is inherent and uh, the, the concept of space is in here, but the brain probably has to learn, you know, when I move, what does it mean to move in spaces like a, there was an internal sort of movement command which is just a bunch of neurons firing and then there's what actually happens in the world.

Matt: [19:20](#) Sure.

Jeff: [19:20](#) It has to learn that connection. This is why the whole system can work for, as I mentioned a moment ago, can work for abstract objects. The whole mechanism that you learn to move your fingers over a coffee cup and then what a coffee cup feels like and how it behaves. That same mechanism can be applied to nonphysical things and movement doesn't have to be physical movement. It can be, um, it can be movement of

concepts or movement of I've got some equations on a book on a whiteboard. I can be moving them around. Things like that. Behaviors can be like transforms in mathematics and anyway, the mechanism says I'm going to try to figure out if there's some sort of quote behavior unquote, um, some sort of input and I'm going to figure out the structure of this thing in some location based space. And now the properties of a, um, I talked about a moment about compositionality and hierarchical compositionality and so on all apply to this stuff. It's a very kind of fundamental theory about how we form representations of the world and how knowledge is represented.

Matt: [20:26](#) Yeah. Speaking of forming representations, learning object learning is really fascinating to me. I do this thought experiment. I think you probably started this. We talk about reaching into dark boxes all the time, it seems like, so it fascinates me that the grid cell space is just an enormously huge, so when you reach into a box and sense something you've never sensed before, you're essentially like sort of randomly picking a space in the ether to start an object,

Jeff: [20:54](#) A point

Matt: [20:54](#) A point, some point and start defining an object and as you move sensors over that object, you're defining that object in your brain over time. That's fascinating I think because it means that you have like almost an unlimited potential to learn.

Jeff: [21:10](#) You do well, it's a well unlimited potential to represent.

Matt: [21:13](#) Right to represent things.

Jeff: [21:16](#) This is a basic property of sparse distributed representations. People who've been following your work and our work know what that is. But basically you can take a set of a few thousand neurons and activate them sparsely and say well, how many things can I represent? It's nearly infinite, greater than the number of atoms in the universe type of thing. So we have this huge representational space. Um, the challenge is in the brain, if you want to learn something, then the amount you can learn is limited and um, you know, we can't learn everything, but it is the same wonder you are expressing. I, I, I feel too, it's like, it's almost unlimited number of things you could learn. You can't learn them all at once. There's almost unlimited but things you could learn if you had time and you had enough neurons and enough synapses. But, um, in some ways you could say, well, how many different locations can a set of grid cells represent and for those who were knowledgeable about grid cells will say

some number of good cell modules, so thousands of thousands of cells. Um, that's, the number of locations you can represent is extremely large and it's unlimited in some sense. Then when you pick one location, you can just start randomly and then you can, all the points around all the locations around that you can move to the movement. They're all nearby. Um, but, uh, and it's sort of an isolated island of points in this huge monster space of points of locations. So like we have this almost unlimited. It's like, it's like imagine that you have your grid cells in your cortex and represent the universe and now you land on a planet where you can explore all the property around that planet or you're never going to easily move to the next planet. It doesn't, you know, it's like, it's their theory. You could move there, but not really.

- Matt: [23:01](#) Well, you said you have associative links to other things.
- Jeff: [23:03](#) You could do this.
- Matt: [23:04](#) You can compose.
- Jeff: [23:05](#) I can say. Yeah, yes, but I was just saying that actual movement, you can't really, you can move locally and map out your local environment, but I'm, I'm not. The point is as I, as I update my location on earth, I'm not accidentally going to think of all the sudden I'm on some location on planet Xenon, you know, it's not going to happen. Those are the locations on planet Xenon are going to be all unique from the ones on earth.
- Matt: [23:28](#) So, uh, what, how does this idea of these grid cells, how does that interplay with HTM systems? How does that work with mini columns? The idea of groups of mini columns. And the different layers.
- Jeff: [23:39](#) Well you're asking some great, that's a great question. It's a very detailed question. Um, uh, I'm going to have to assume that the listener knows something about our temporal memory.
- Matt: [23:49](#) I have warned them upfront, and provided resources in the show notes.
- Jeff: [23:54](#) Because there's other things we could talk about that are less tactical than that. But let's dive in. Okay. So what we came up with the temporal memory algorithm eight years ago, maybe nine years ago, something like that.
- Matt: [24:04](#) It was before I was here

Jeff:

[24:06](#)

And it has quite a few innovations in it that I think are actual things are happening in the brain and one of them is the use of mini columns and what, here's one way to look at it. You want to represent something, In this case, in the temporal memory, wanting to represent some sensory input. Right? And we've run it through an encoder which you've been talking about. So you have some, you have a representation of that input. In our temporal memory system, that representation is actually the mini columns themselves. It's not the individual cells in the mini columns it's which minicolumns are active. That's the output of the spatial pooler and you say, okay, I have some input. I'm gonna represent it by some active, sparse activation. Every bit in that, in that output of the of the spatial pooler, every bit in that representation is associated with a group of cells. We'll call them minicolumns and then what we're gonna do is we're going to say I can pick one cell in each column to be active at any point in time, so if I have an active minicolumn, I'm going to pick one cell to be active. What that allows me to do is allows me to represent the input in very large number of contexts, so depending on how many cells are in each area, but it doesn't take many. If I had 10 cells per mini column I'd be good forever. It does not take a lot because we're not - It's not like oh, a one cell per minicolumn I can represent one thing. The cells in minicolumns can be used in many different contexts, so I could have - Imagine if I was representing a note in a melody, and so here's a note coming in and I, uh, and I want to represent it, learn that note in many different melodies. The same note could be in many, many, many different melodies, many locations in the same melody, and it's almost unlimited the mathematics work out, like we talked about earlier, even if I have just 10 cells from many columns, I have almost an unlimited number of ways of representing the same thing in different contexts.

Matt:

[25:56](#)

Yes.

Jeff:

[25:56](#)

That's what the minicolumns get you. Now, let's translate it into grid cells. I'm representing a location, right? Well, I'm going to do the same thing. I'm gonna say, well, I have a location, but I might want to represent that location in many different contexts. Now, why would I want to do that? If I have a location on this coffee cup? Why? What are the different contexts? Well, let's go back to our stapler. Okay? Okay. The stapler changed its shape. Okay. Now it's the same stapler, so it's the same space of points at one point in time, one context, a point is occupied by the tip of the top of the stapler and another form of time, the tip of the top of the stapler is someplace else in that space. So there's a location at one moment in time had the physical

stapler and then another time it does not have a physical stapler, but it's the same point, it's the same location, so what I'm wanting to be able to do is say that location is the same location in the stapler, but at one moment in time it's occupied by this component and other moments, not. There's different states of the stapler.

Matt: [27:01](#) Yeah, within the reference frame of the stapler.

Jeff: [27:04](#) Imagine looking at your mobile phone and you look at the screen. There's a location on the screen. Well, different things appear in that location all the time. A menu comes up, then a graph, then a number, then something else. That same space of the cell phone, the same locations associate with that cell phone, have many different features appear at different points in time, right? I need to know the state of the cell phone to know what's going to appear at that location. The same location can have different things occurring there under different contexts.

Matt: [27:35](#) It depends on behavior.

Jeff: [27:38](#) What app am I running right now, right? It's the same phone, the same location. In this case, the morphology or the shape of the phone hasn't changed, but there's a location on the screen that one moment represents the one app and other moment represents another app, another moment represents a menu icon and so on. So the point is this idea that you represent, you have something like, oh, we started by talking about this sensory input and I want to represent the sensory input in different contexts. I now have a location on an object, but I want to represent that location under different contexts because if I'm going to predict what's going to be at that location, I need to know the context or the state of the cell phone or the state of the stapler. So this basic idea, this you asked me about the, what's the, what's the relationship with minicamps? We just kinda, we think this is happening everywhere that you represent something like a sensory input or like a location or like this display factor I was talking about earlier. Um, and, but I want to be able to represent many different contexts. And um, and so that's where the role of minicolumns come into play. And even when, even animals, they don't know, they don't see that there's many columns, but the same basic principle seems to apply that there's a bunch of cells that have the same sort of receptive field property and they differentiate under context. So it doesn't have to be, some animals like humans and monkeys have physical minicolumns

you can see. Some animals, you don't see them physically, but they're kind of conceptually still there.

Matt: [29:06](#) Sequence memory is happening and at least one layer of cortex, and so there's minicolumn activations happening there driven by sensory input. Elsewhere there's some grid cell stuff happening that somehow synched up with that?

Jeff: [29:19](#) Well, let's getting complicated. I don't think we want to go there now. We're writing a paper, another paper about this right now about how it is that grid cells and sensory input coordinate.

Matt: [29:29](#) Okay.

Jeff: [29:30](#) Right, so imagine I reach my hand into that black box you mentioned earlier and I touched something with one finger. Well, I can't tell you what it is. It could be a lot of things. I might say, well, it feels a little bit the Coffee Cup and maybe it feels a little bit like the stapler. Maybe it feels a little like a pen. I don't really know. What happens is, is that now you move your finger and you sense something else and what the brain does now though, what the columns representing the fingertip do, they say, oh, what object do I now know that's consistent with the first sensation and the second sensation displaced by that movement?

Matt: [30:01](#) Oh, right, right, right.

Jeff: [30:03](#) And so you start eliminating things very quickly. Um, so if you only have like one sensor, like you're touching with your finger or you're looking at the world through a straw, you have to move around and look at and as you move you're sort of saying, oh, I see this feature, this feature, this position, and you have to sort of, it's not just the features, it's the features in relative positions to one another. Right. And, um, so, uh, that's that mechanism by how that works. Uh, we think we understand a good portion of it. It's related in the Cortex to layer four, which is the sensory input layer and layer six, which we think is a grid cell layer and how they interact. So exactly how that works, we don't know exactly, but we have some pretty good ideas of the basic mechanism and so we're in the process of preparing a paper on that concept as well.

Matt: [30:51](#) Great. And Subutai said hopefully when we get to a point where we're ready to try and get it published somewhere and we'll give it an open access.

Jeff: [30:59](#) Oh yeah, totally. I mean, um, I mean, I don't think these papers are in a form right now that would even make sense to share with anybody right now, but my goal, um, is, is, uh, as you know, I have some talks in the fall I'm going to give and I want to have those two papers, that cover all the topics we've been talking about. I want to have them available public at that time. They will unlikely be accepted into a journal at the time, but we will post them on bioRxiv for the preprint servers. So people could read them. Um, and uh, we'll do it as quick as we can.

Matt: [31:36](#) We've always tried to be as transparent as possible and I like that we do this.

Jeff: [31:40](#) Yeah, I don't, I, you know, we haven't ever talked about like posting our, you know, scrap writing as we're going along here. I don't know if that would be a good idea or not, but -

Matt: [31:49](#) I do it all the time. But I wouldn't advise you to do it.

Jeff: [31:53](#) We'll try out language and then we change the language, you know, and it can be very confusing to people if we suddenly start using different language for the same things and no one knows, understands and there's holes and it's quite messy at the moment. Uh, but we're making good progress on both of those papers.

Matt: [32:10](#) This was part one of a two part interview with Jeff Hawkins. The next podcast episode will contain part two. If you like what you hear on the podcast and you want to discuss ideas like this with intelligent, friendly people, be sure to join HTM forum at [Discourse.numenta.org](https://discourse.numenta.org). Our online community was created around the Numenta open source project and continues to thrive on HTM Forum. Hundreds of folks interested in HTM and related theories, share ideas, experiments, and open source code. If you are an HTM theorist, engineer or a programmer, or just a hobbyist, HTM forum is a friendly place to keep up with the latest in HTM technologies. Thanks for listening to Numenta On Intelligence. Be sure to subscribe to our podcast on your favorite podcast service. To learn more about Numenta and the progress we're making on understanding how the brain works, go to numenta.com. You can also follow us on social media at Numenta and sign up for our newsletter.