Brain Science Podcast #38*- Interview with Jeff Hawkins, author of *On Intelligence

This episode first aired on May 30, 2008

Note: All times are approximate. If you are listening to a download that contains inserted advertising, the times will become less accurate later in the recording.

0:00:00
This is the ***Brain Science Podcast***. The podcast for everybody who has a brain. And I'm your host, Dr. Ginger Campbell. On the ***Brain Science Podcast***, we explore how recent discoveries in neuroscience are unraveling the mysteries of how our brains make us who we are. For more information, including show notes, links to previous episodes, and information about how to subscribe, please go to the website: **www.brainsciencepodcast.com**. We also have a discussion forum at **www.brainscienceforum.com**, and you can send me e-mail at **docartemis@gmail.com**.

0:01:00
Welcome back to the ***Brain Science Podcas***t. This is episode #38. Today I have a very special guest, Jeff Hawkins, author of *On Intelligence*. When I originally reviewed Jeff's book back in the summer of 2006 for the sci-fi show, the experience inspired me to start the ***Brain Science Podcast***. So I'm really excited to have Jeff on today. There's one thing I'd like to mention before I get into the interview. I want to let you know that there is a new episode of my other podcast, ***Books and Ideas***, out. And it is an episode about meditation. I didn't put it in the ***Brain Science Podcast*** feed because it's not that much about meditation and the brain, although the subject does come up. So especially for those of you who are waiting for an episode about meditation, I think you will enjoy the episode. It is an interview with forensic psychologist Delany Dean. And you can find that at **booksandideas.com** and in iTunes™. Before we get into Jeff's interview there is one thing that I have to warn you about. Jeff was on a cell phone while we were talking and so his voice is not as clear as it could be. I want to apologize for that and promise you that in the future I'll be more insistent with my guests to try to get them to avoid using cell phones. However, despite the sound quality, Jeff's enthusiasm for brain science definitely comes through and I know you're going to enjoy the interview. Please stay tuned at the end for some closing announcements.

0:02:50
Ginger Campbell: Ever since the book *On Intelligence* was featured back in episode 2 of the ***Brain Science Podcast*** I've wanted to interview the book's author Jeff Hawkins, so today I'm really proud to be able to welcome Jeff onto the podcast. Jeff, I really appreciate your taking the time to talk with us today.

Jeff Hawkins: No problem Ginger, I'm happy to be here.

Ginger Campbell: For the sake of my listeners who may not have read *On Intelligence* and may not be familiar with your work, could you start by just telling us a little bit about your background?

Jeff Hawkins: Yeah. So, I had this sort of Dr. Jekyll/ Mr. Hyde type of lifestyle. I've done sort of two different careers. I'm 51 years old and I've worked for many years in the computer world. So many people know me from that. I was the founder of a couple of companies, Palm computing and

Handspring. I was the designer of the Palm Pilot and the Treo Smartphone. But throughout all that time, going back to 1979, actually my main interest has been neuroscience and brains. So I've been sort of living in two worlds at once, and as part of that I did a stint as a graduate student at Berkeley. From 2001 to 2004 I started and ran a neuroscience institute that was focused specifically on the theory of neocortex, and that was actually very successful, and that still continues on now at UC Berkeley. But my interest today has been understanding how the neocortex works and some other related brain areas like the thalamus. Understanding it from a theoretical point of view, and a biological point of view, and then there are machines that work on those principles and we've made a lot of progress in this area. My book when it came out- I guess in 2004- we made a fair amount of progress understanding how the neocortex functions and forms memory, and the book described that, and since then we've figured out more about that and we've figured out how to make machines work on the same principles.

0:04:48
Ginger Campbell: I want to ask you a question about... I think we're just about the same age, I think I might be a year older than you, but I also came out of engineering school around the same time as you did but went to grad school and ended up in medical school. I was interested in the fact that your first job was at Intel because wasn't that in the very early beginnings of microprocessors?

Jeff Hawkins: Yeah, well. I was right out of school so I really had nothing to do with any of that but yeah, it was just when the microprocessors were starting to become popular. I think the 8086 had just been introduced or something like that. And they had only been around for a few years before then.

Ginger Campbell: My husband was starting his job as electrical engineering professor in 1977 and we were using the Motorola chip. We had something called the D2 kit that we had to program in hexadecimal- I can still remember that!

Jeff Hawkins: Yeah, in fact it was a great time to get into the mobile computing or the computing industry, but my story is in this fall of 1979. There was an issue of Scientific American and the September issue is always dedicated to a specific topic. And that year the topic was the brain. And many neuroscientists remember this issue because it was influential for quite a few people. And it had all these stories about the brain, about the neuron, about diseases of the brain and so on, and the very last article was written by Francis Crick, who was one of the co-discoverers of the DNA molecule structure. And Francis had turned his interest to neuroscience and he wrote a piece- I think it was called "Thinking about Thinking" but I'm not sure about that title. But anyway, essentially he said, "look," he just read all these facts about brains and people act as if we know how it works, but in reality we really have no idea how it works. We have all these data but no overarching theory. And the words he used was, it was conspicuously lacking the framework to understanding these different ideas. And that struck me as correct and a little challenge. And it got me excited about brains. I said, "You know what, I'm gonna work on that." I had just started a job as an electrical engineer and I said "I'm gonna become a brain scientist, and I'm gonna dedicate my life to understanding how brains work.” And that's what got me going in this field.

0:07:09
Ginger Campbell: The only reason that I took that aside into talking about your work at Intel is that I think as we get into our conversation, I want my listeners to understand that you are a person uniquely

qualified to comment on the difference between brains and computers, because you really understand how they work at the very guts level, which a lot of us don't. So that's really why I wanted to bring that piece out of your background there.

Jeff Hawkins: OK, well there's probably some truth to that, yes. And in fact, its an interesting commentary because when I first started studying brains, it was intuitively obvious to me that brains do not work at all like computers. And I understand computers pretty well. You know, I studied quantum mechanical devices in college and I knew about operating systems, so I understand computers pretty well, but it was just obvious to me that brains didn't work on the same principles. But it wasn't obvious to a lot of other people.

Ginger Campbell: [laughs]

Jeff Hawkins: Since the beginning of the computer era back starting when Alan Turing wrote his famous paper about computing machinery and intelligence in 1950, since that time people thought of brains like a computer in that they thought "We can just program computers to be and they would do things like a brain does." And most people who were trying to build intelligent machines thought there was nothing to be learned by looking at brains. They just thought it was a waste of time and that there was no reason to look at brains and I went through a lot of roadblocks, because my push from the very beginning was "hey, if we're going to build intelligent machines we have to understand how brains work, and they work on some different principles. We need to understand those principles before we can go about trying to build a machine that works like that. But that was not an accepted idea back then.

0:08:57
Ginger Campbell: Can we talk a little bit about your book? I know it's been a few years since you wrote it, but I'm sure the ideas are still pretty fresh in your mind. In your prologue, you said you wanted to propose a new theory of intelligence. Why do you think we need a new theory? Or did we have an old theory?

Jeff Hawkins: No. We don't have an old theory. Well, it depends on who you speak to. The whole field of brains and intelligence and intelligent machines... there are thousands and thousands of people who are interested in this, working in various fields. There's linguistics experts, there's computer scientists, there's neuroscientists. We have tens of thousands of neuroscientists in this country. So this is a big area and there are a lot of ideas. So on the one hand, I generally like to just say, "you know, we don't need new ideas. We need to understand what we already have." But when you're writing a book, you really need to say, "look, there's a synthesis here that I don't think the reader will have ever heard before. And I'm presenting it, and I called it "new" at that time. And you know, there are some very new things about it, but in general, there hasn't been a unified or any sort of coherent, overarching theory about brains and brain function. That's what I said earlier about Francis Crick. He said that we're missing that theoretical framework and you know while I was trying to do my book, and I think I succeeded to some extent with providing a theoretical framework for understanding a lot of disparate pieces of data about brain science and behavior.

0:10:28

Ginger Campbell: Did you ever get to talk to Crick in person before he died?

Jeff Hawkins: Yes I did, but I got to know Crick and James Watson fairly well. I had a nice social tea with Francis Crick down in San Diego, and then I gave a talk at the Salk Institute where he was hanging out at the time, because he had spent the latter part of his life interested in neuroscience.

0:10:50
Ginger Campbell: I've got to interview Kristof Koch, but I obviously missed Crick because he's been gone for a few years. So one of the things that stuck in my mind about your book was that you criticize the approaches taken by artificial intelligence and neural networks, especially for their failure to consider how brains really work. What do you think are the most important features that we really should be taking into account?

Jeff Hawkins: Well, just for starters, we have one example of intelligence, and that's the biological brain. People have tried to replicate brain functions in computers for many, many years and you generally fail. I don't want to say completely because there's all kinds of things computers do that are useful. But the basic ideas of how we- the most simple things that humans do- perceiving the world, vision, language, motor control. These are things computers can't do. And so I would start off by saying, "Let's study brains. You know, that's the only way to get to this point. Now at this point I have a lot of understanding about the difference between computers and brains, but in the beginning of my book I actually tried to make the case that you have to stop thinking about the brain like a computer, and most people just don't realize how poor an understanding we have about brains. And you gotta understand how poorly we've done at replicating brains in computers. My intent was less to criticize- it's just to point out the state of recent endeavors. Brains are still a mystery- I think i called it "the last great terrestrial frontier of science." You know it's like one of the last things that is so everyday important to us that we really have very little idea on how it works. And I also just wanted to point out that some people have the impression like "oh, don't we have intelligent computers? Can't they do these things? And the answer is "no, they can't." I couldn't even make a case before I said "look- we need to rethink our whole approach to this." You have to realize how much we struggle with this over the decade.

0:12:53
Ginger Campbell: That was something I came to the book sort of knowing, so I guess that point might have gone right by me since I wasn't the one needing that point

Jeff Hawkins: Well, you'd be surprised how many laypeople would just say "oh yeah, we understand how the brain works. There are stories all the time in the paper." Or "I read about that computerized car. We must understand how computers can see or understand language" and things like that. And it's not true. So I did need to sort of set that level of expectation.

0:13:23
Ginger Campbell: I was wanting to sort of get into a couple of these little features that brains do that really are different from computers, or at least the way they do them is very different. Can you talk a little bit about the importance of feedback?

Jeff Hawkins: I'll step a little bit back from that. A brain is not a computing function. It doesn't have software, it's not programmed to do anything. The best way to think about the brain is a form of memory. It's a memory system. When you're born, your brain knows very little about the world. It doesn't know about cars and trucks and buildings and books and podcasts and doctors. It doesn't know any of this stuff. You're born with sort of a brain that knows nothing about the world and you have to learn these things. No one programs it- you have to learn it and you learn it only one way- through exposure to patterns in the world through your senses. And so you have to start thinking about brains as a memory system. And it doesn't work like the kind of memory we have on a computer. It works on different principles, but in some sense it stores information, and it can retrieve it. And that's what you are doing when you are thinking and when you are perceiving the world. You are taking new patterns into your senses and recalling stuff that's previously been stored in your memory. And my work is all about understanding how those memory systems work. How does the brain memory system work and how's it different than other types of memories we might be familiar with.

One of the clues that help me make progress on this problem was a biological fact- that if you look at brains there's a huge amount of feedback going on. That is, the information comes from your senses and it goes to different brain regions and different areas in your cortex for example. But information is also flowing backwards. And many people think of it more like a computer- you have your inputs, you process it and then you have an output. "Brain must be like a computer." But in reality, it looks like as much information is going backwards as is going forward. And this was a big clue, because most people's models of thinking and brain function ignore the feedback. This is just one of the things that they ignore. And another one is the importance of time. But feedback was one of those big clues. The anatomists told us "hey brains have all this feedback." It's there. And yet you look at your patients and you say like "well, I don't think it's doing anything, because you know, I don't understand why I need the feedback." And then they would literally write papers saying, "I don't think that feedback is important." But it's a dominant feature of all brain regions, and clearly in my mind it was important, so that was just one of the clues that we used to say "look, you know, we can't ignore this." There are just several things we can't ignore and that was one of them. And I can talk about its role in brain function if you want, in the cortex, and the importance of it.

0:16:07
Ginger Campbell: Yeah, let's come back to that in just a minute. Your work kind of has focused on the cortex. Would you like to kind of just explain why you choose to sort of isolate down to mainly the cortex?

Jeff Hawkins: Yeah, sure. If you look at a human brain, if you open up the skull and you see that wrinkly surface on the top, that's the neocortex. And it's actually a full sheet of cells- it's about three or two millimeters thick, so double business cards type of thing. And there's a big sheet- it's like the size of a dinner napkin. And it gets wrinkled up to get in your head. Now that neocortex is volumetrically about 60-70% of the volume of your brain. So it's a dominant feature of a human brain in terms of its size. Besides it's size, what's most interesting about it is that the cortex is where we store the vast majority of learned knowledge in the world. Cortex is in charge of all high level language. All the meanings of words and how you use language- it's all in your cortex- understanding language and creating language. Neocortex is almost completely responsible for vision. There are a few other areas but mostly it's the organization. So everything we know about how the world works and expectations about visual things is all in the neocortex. And it's also- the neocortex is in charge of all high level motor control. So planning and moving and creating my speech right now. So it is a place which is really all high level- almost all high level thought processes in the brain and if you're interested in understanding language or building machines that can see, for example, and the way that brains do, it's the cortex you want to start with. You probably can't ignore the other parts of the brain, but it's a good place to start. And another thing that's great about it, is that all mammals have a neocortex, and non-mammals don't. The neocortex came onto the scene in evolutionary time fairly recently. And it got large very rapidly in evolutionary time. And so it's a fairly uniform structure. It's doesn't have a lot of... even though it does all these different things, it's actually very very similar throughout. And so it gives you the promise that you can understand a lot about what the brain is doing and about what thought processes are, without studying a lot of different things. In fact, this is one of the premises of the theory that I'm working on, and it's well-known in the neuroscience world that the way the brain actually does language and hearing and motor control and so on, is actually the same- this has been known now for almost thirty years. So, it gives the promise of understanding a whole bunch of high level thought processes all at once, if you can understand how the neocortex works.

0:18:46
Ginger Campbell: Would you explain why the work of Vernon Montcastle is so important?

Jeff Hawkins: Sure- Vernon Montcastle is a neurophysiologist at Johns Hopkins. He studied a lot of different aspects of the brain. But he's really focused on the neocortex. He wrote a wonderful book and many papers about the neocortex. And one of these papers, which he wrote in 1979, which I read for the first time I think in like 1985 or so, is in a little monograph, actually. And this was the first time I read this idea. It think it was the first time it was published. He made the argument that if you look at the different areas of the neocortex- the areas involved in vision and language and motor control and touch, and so on, that the architecture- the detail- of how the cells are arranged and how they connect to one another, is really identical everywhere. This idea that they look similar has been known back since the turn of the previous century- Cajal, back in 1900. But, Vernon Montcastle made the following speculation- he said, not only do these areas look the same, but I believe they're doing the same thing. And a vision area is doing the same thing as a language area, which is doing the same thing as a motor output area. He said, what makes it a vision area is because it's hooked up to the eye. What makes it an auditory area is because it's hooked up to the ears, but the fundamental property that these different areas in neocortex, what they're doing, is the same. And this, to me, is one of the greatest discoveries in science, because it's one of these things you look and nature suprises you, and you go, "oh that is so amazingly beautiful." And it explains so many things we know about trauma and about physiology, and it gives us like, "isn't this beautiful that nature's come up with a single way of doing all these different things." And that, to me, was just a tremendously liberating idea. Now you could search for that sort of uber algorithm, that uber idea, those principles, that span modality, that span behavioral tasks. That really clarified it for me quite a bit. It's just one of those great discoveries, and Vernon Montcastle was the first guy to basically nail it, saying, "not only do these look the same, but they're doing the same thing."

Ginger Campbell: And now there's lots of evidence to support his predictions.

Jeff Hawkins: Yes, there is. In fact, there was a lot of evidence back when he wrote it, but it was such an unexpected idea that not everybody liked it. And over time, I actually- I've held a whole bunch of conferences with neuroscientists, and I usually start by asking how many people believed in that. And in the beginning when I started asking people I would find maybe a little bit less than half who said- half of the neuroscientists believed what Vernon Montcastle proposed. Today, I don't have a number on it but it's well over half- I would say it's more like 70 or 80 or 90% of neuroscientists who say, "yeah, there's truth to that." And the evidence is overwhelming. The only obstacle people have with it is it's kind of hard for them to believe. If you spend your life studying vision and you come up with all these vision algorithms and in speculation about how the brain does vision, and then someone comes along and says, "Well you know what, there isn't a vision algorithm. It's just a cortical algorithm that looks at vision," then you might object to that. You might find it like, "hey, I don't quite believe that." But the evidence is overwhelming. It's true, actually. It's a fact. It's a beautiful fact.

0:22:11
Ginger Campbell: I'll be right back after a brief word from this month's sponsor, US Navy Medicine

Note: Depending on when you download this episode it might have a 30 second ad here which will make the times slightly different from the ones listed.

0:22:20
Ginger Campbell: Neuroplasticity has been one of the themes that we spend quite a bit of time on on this show, and we talked about that experiment where they rewired the ferret's brain so that the- I can't remember- is it the auditory cortex to the eyes, or something like that?

Jeff Hawkins: I don't remember the details either, but I remember- I know the experiment you're talking about- it's famous.

Ginger Campbell: So the animal was able to see with it's auditory cortex, I think was the bottom line.

Jeff Hawkins: Yes.

Ginger Campbell: That's a real attention-getter!

Jeff Hawkins: So you see in the humans if a human is born with trauma with the eyes, for example, then the area of the cortex which is normally used for vision becomes tactile. You see that in trauma, and you see it in birth defects. There's tremendous evidence for it. The brain is very plastic, and essentially cortex is cortex and they'll look for patterns in whatever is available [laughs]. It's a wonderful sort of idea and I always enjoy thinking about it. It's a beautiful idea and there's lots of evidence for it.

0:23:19
Ginger Campbell: Getting back to your book, one of the starting places that you worked from was the contingent that behavior is not an adequate measure of intelligence. I love the quote, and I'm going to say for the sake of my audience it's on p. 33, where you said, "Intelligence is something *happening* in your head. Behavior is an optional ingredient [Jeff Hawkins laughs]. Why have you concluded then that memory and prediction *are* keys to understanding intelligence?

Jeff Hawkins: Well, by you asking me to recall how I got to certain beliefs and dogmas and it's a little bit difficult for me to do. For example, the prediction we've already mentioned- that prediction had to be important because it's so prevalent in the neuroanatomy of the brain. So that was a starting function. When I first started thinking about feedback, I didn't know what it was for. But I've always had the sense that the brain's about memory, because I just knew that when you're born you didn't know stuff and now when you do life you're learning things and in order to just understand a visual scene, like what's in front of me right now- the buildings and cars and trains, and so on- I have to have all this memory. So clearly memory is an important function. The reason I'm talking about behavior is because so much of what was done in thinking about thinking or intelligent machines or even neuroscience was influenced by people like Alan Turing who proposed the Turing test, which is to say, we know if a machine was intelligent if it has this outward behavior that we couldn't tell if it was different from a human. You see a lot of philosophical argument- people like John Searle, who wrote about the Chinese Room Experiment and so on. This is all about behavior. People are thinking "well, we'll know that something's intelligent if it has the correct behavior." And that's always just felt to me intuitively wrong- that humans can be very smart without having any behavior- we do that all the time. People listening to this podcast don't have to be doing anything to be smart and understand what we're talking about. And behavior- you can have all kinds of behavior that isn't really very smart at all and you can fool people. So in my journey through life, I was trying to find out things and presumptions that we were making that might be wrong. So this whole behavior paradigm- the Turing test- to me seemed like, "that's a red herring. Don't fall for that." And then the prediction thing came about later, and I can talk about that if you want, but I've already spent a lot of time on that last question.

0:25:50
Ginger Campbell: Do you think we need the prediction in order to be able to talk about your memory prediction framework? Do you need to put any background in there?

Jeff Hawkins: Yeah, okay- I'll try it now, sure. In my book I describe an "aha" moment that occurred to me many years ago. I was very confused by what brains actually do, seems like a silly question to say "what does the brain do?" But when you actually try to pin it down- nail it- it can be difficult to do that. And I was in my office at home one day, and I had this experiment I described in the book. It explains where I realized that there's so many things in the world that were just slightly different than normal. I would notice them. And I had this "aha" moment that essentially said, "you know, we are making predictions about the world all the time. About everything. You predict what words you're going to hear. For example, because you understand English, you can predict what word is at the end of this [pauses] sentence. And it just happens to you all the time. And when you touch things you have expectations about what they're going to feel like, and their temperature, and their texture. As you walk you have expectations about where your feet are going to stop, and what it's going to feel like. So it became coded in the brain and its making these massive predictions about things all the time. And you're not generally conscious of it. You only become conscious of it if an expectation or a prediction fails. So all of a sudden if something happens that is unusual and what wasn't predicted, then your attention is drawn to it. But you're making predictions through all your senses all the time through your waking life. And this, to me, was another big clue. So I said, "OK. I have this brain and its memory system of some sort." At the time I didn't know how it worked, but I said, "it has to be making constant predictions about things and it's making constant predictions about *novel* things- things you've never seen before. Exact details you've never seen before." So if I put my hand down on some table, it may be a new table and I maybe had never been in that exact situation, but my brain still makes predictions about it. So I said, "how can the brain make predictions about novel things based on past experience?" So these are sort of clues that I was collecting, which led to the theory that eventually came out in the book and which we've been working on since.

0:28:03
Ginger Campbell: I want to talk about the memory prediction framework, but I'm wondering if we've left out any pieces?

Jeff Hawkins: Yes. There's a couple of pieces that are big clues that are important to me. One was the role of time. Many theories, whether they were brain theories or whether they were computer science-type theories, they ignored time. People who created neural networks, they mostly ignored time. And the logic for this very much related to vision, because when people say, "look- I can look at a picture, and I can flash the picture in front of someone's face. They can recognize what it is and there's no time involved in that." They argue, "well, we can probably not think about time right now." We can try to understand, for example, computer vision or human vision, without thinking about time. When you think about your other senses, such as hearing, and touch, you can't do anything without time. You really can't do any sort of high level recognition or language or processing of auditory sounds without a flowing through time. And you can't do things- touch. I can't really touch something and know what it is unless I move my fingers over it. And most vision itself is time-based. So as I look at the world, things are moving all the time, my head is moving all the time. So that was another clue that to me was, "hey look, any theory about brain function has to intimately be involved with how patterns flow through time." And if you're not doing that, you're not going to capture it. Because brains are working with time-based data. And then finally, one of pieces that came together late for me was knowing that the neocortex itself has a particular structure to it. It has an organization which we call hierarchical. If you look at all different regions of neocortex and how they're connected together, they have an organizational structure that looks like an inverted tree, like a big pyramid shape to it. That is, there's a whole bunch of regions that project to a whole bunch of regions, project to a whole bunch of regions. And there's this hierarchical structure to the neocortex, and if you took the hierarchical structure feedback in time and prediction, then you put them all together and you think about it for a while, you can figure out how the whole thing works.

0:30:13
Ginger Campbell: Do we need to talk anything about how patterns are stored?

Jeff Hawkins: Well, we can get into the theory that I proposed in the book and that we've been developing since then, if you want. It turns out that in order for any kind of memory system to make a prediction of any sort, it has to store how patterns change through time. And so it has to store sequences of things. You could take an easy example of a pattern that exists in time like a melody. You can't recognize a melody by hearing one note. You have to hear it through time. You can only make a prediction like "what's the next note I'm going to hear in the melody" if you've been exposed and remembered how patterns move through time- how the notes move through time. So one of the key elements of this theory- the memory-predictions framework- is that the brain is storing in sequences of patterns. You can think of them like little songs, like little melodies, but we do it not just for auditory.

We do it for vision and we do it for tactile. So now the sensory inputs as well. There's a theory about how memories are formed by storing sequences of patterns through time in a hierarchical memory structure.

0:31:23
Ginger Campbell: So do you want to tell my listeners your elevator description of the theory? [Jeff Hawkins laughs] And then we can expand?

Jeff Hawkins: If there *is* an elevator description of the theory, yeah we've touched on most of the bases of it so far. Essentially, the neocortex- and we're talking about the neocortex, not the other parts of the brain- it is this massive memory structure. As I've said earlier, it stores almost everything you know about the world, and you know a lot about the world. All the details. Everything. It's all there. And essentially what happens in this theory is you have this memory structure- it's shaped like a hierarchy- there's memory regions that are close to the sensory input and other memory regions that are one step away and other memory regions one step away from that. And the basic operation- the simplest way I can describe it- is each section in the memory is trying to store sequences of patterns. The ones at the bottom of the hierarchy- the ones closest to your eyes and your ears and your senses- they store little sequences in time of a little part of space . So in the auditory world you might imagine little sounds that are blips going up and down and so on. In the visual world it would be little lines moving left and right in different parts of your visual field. These have all been documented in neurophysiological studies. And then what happens is that these little sequences that are learned at one level, they get passed a level up. But what gets passed up is like the name of the sequence, it's not the details. You can figure it like a melody. If I'm a little section of cortex and I see a sequence of patterns, I say, "Oh, I recognize this sequence of patterns." What I pass to my parent- the next region up in the hierarchy, is the name of it. It's a constant representation. The parent region is looking and sequences of those sequences. And then the next region up looks at sequences of sequences of sequences, and so on. So the end result is when you have a memory system like this, and it self-trains, when you have a very fast changing pattern at the bottom, like my speech or like what's entering your eyes when you scan around the world. As information flows up the cortical hierarchy, it becomes slower and slower, and at the top you find the cells responding to slower concepts in some sense. And this system- when you build a system like this and you train the way we train a human- you expose it to patterns- it ends up building a model of the world. The end result is that you have a model of the world- all the things you knows about the world are stored in this memory model, and it's a hierarchical model, for the model consists of things that are low-level details of the world, and then higher-level details, and then higher-level details still. That's really the best I can give for an elevator pitch. It's a hierarchical memory model. You expose it to patterns, and it learns the structure of the world by basically learning what patterns flow through time together. And it does it in an hierarchical fashion. And believe it or not, that can explain an awful lot about human behavior- about how we learn, about how we think, about how we make predictions, and so on.

0:34:21
Ginger Campbell: So when you say you're having an experience and the information's going up the hierarchy, the higher levels are then feeding back what they expect to happen next?

Jeff Hawkins: That's right. So the first thing is you have a new input coming in on your senses. And the thing about that is in your life you'll never have the same pattern on your retina twice. Even if you looked at your husband or your wife over and over again, it's always different every time. You have these novel patterns coming in, and then it immediately begins to recognize these patterns, by recognizing the low-level details, and then the next-level details, and so on. And in almost all cases, even if I know what's going on- at the top-level hierarchy you say, "OK, that's my husband" or "That's a dog" or whatever. But every region can make predictions about what can happen next, because every region has stored sequences. And so every region is saying, "I know what's likely to happen next." At the high-level, at the low-level, at the middle-level, and so on. And you feed those predictions down the hierarchy. And it basically says, "You know what, you should be expecting to be seeing this at this point in time and you should be expecting to see this at this point in time." Which is useful for a number of reasons. It clarifies confusing data that's coming in. So if you're seeing something ambiguous, your brain doesn't want to see it as ambiguous. It basically says, "You know what? I'm going to interpret it one way or the other." And you tell all the regions below it, "Interpret it this way." So you see that bi-stable patterns like little staircase patterns or Necker cubes, if you know what those are. So prediction is an important part of this whole process. And prediction also is how we think because if I want to just silently think about something, what our brain is doing is it's taking it's current state-its current beliefs. It says, "What would likely happen next?" And you can turn that prediction and ignore the real data from your senses. But take the prediction and feed it back into memory, and say, "Well, what would happen next and what would happen next?" So you can follow a series of basically memory predictions, and that's what thinking is. That's when you're practicing a speech or thinking about what's going to happen tonight when you come home. The brain's just following a series of predictions and it can walk through series of imagined activities.

0:36:28
Ginger Campbell: One thing I like about your theory is that it explains so much of our day-to-day experience, like why when something changes it catches our attention. You enter a room and the one thing you notice is the thing that wasn't there.

Jeff Hawkins: Yeah, and that goes back to this idea that you're making these massively parallel predictions about so much in the world, and you're not thinking about it. Our intuition is like, "Oh I look at the world and I see things." You know. There's magic in that. It's very complex. But actually what you see is you're looking at the world and you're comparing it to your memory of the world, and your brain is saying, "Is this right?" And if even the slightest little thing changes, you'll notice it. I have a coffee cup sitting in front of me right now, and it's round. And if that coffee cup was even the slightest bit off round, if it was slightly just a teeny bit oval-shaped, I would immediately notice that. And I'd say, "It's wrong." And that's basically telling me my brain knows what it's supposed to look like in great detail and has memory of this object and it says, "It should look differently than that." And so your brain is doing this all the time where you're not consciously aware of it. It explains a lot about behavior. It explains how we think and how we notice things.

Ginger Campbell: And why when your power goes out it wakes you up?

Jeff Hawkins: Yeah [laughs]. I'm not sure exactly what you're referring to on that question but...

Ginger Campbell: Well haven't you ever had the experience of- you know, you're asleep and your electricity goes out and you wake up and you wonder, "Well, why does the absence of all the..."

Jeff Hawkins: Yes. You know, there's an interesting story, which I don't know if it's true, but I like to repeat it anyway. In New York they used to have elevated trains like they used to have in Chicago. And they no longer have the elevated trains. And I think there was an elevated train that used to run up Third Avenue. And the very first night they stopped running the trains, the police station was getting calls in the middle of the night. Lots of calls. And all the calls came in at certain times. And it came in at the times when the train would have passed people's apartments. And so people would wake up in the middle of the night and they said, "Something's wrong" and "something woke me up." And they called the police stations. Well the only thing that was wrong was there was no sound. There was no train going by. I'd like to believe it's true. I read it- I can't find the original source for it, but it's the same thing. I could believe it.

Ginger Campbell: Yeah, it fits anyway. It sounds believable.

Jeff Hawkins: It's amazing. Our model of the world is really good. And it knows so many things that we're just not consciously aware of. And when things change, even the absence of something, we say, "What's going on here?"

0:39:05
Ginger Campbell: I know your focus is on the cortex, but I have talked a bit about memory in the podcast in the past. Could you just say something about where the hippocampus fits in, even though it's not part of the cortex?

Jeff Hawkins: Yeah. So there are a few parts of the brain that are intimately related to the neocortex. The hippocampus is one, the thalamus is another, the basal ganglia, which is a collection of organs, are others. And we study all these... we have a model for hippocampal interaction with the brain. Now people think of the hippocampus as very important for forming new memories, and that's true. It isn't the ultimate location for most memories. So, for example, if you don't have a hippocampus, you will not form new memories. However, if you lose the hippocampus, the memories you formed in the past are not lost. It's not the ultimate repository for memories, but it seems to be very critical in forming them. I think we understand how this works in the neocortex, and I'll try to describe it as simply as I can. One of the things that the neocortical memory model that we have developed does is that as we are observing the world, it has to throw away details as the information goes up the hierarchy. So I want to recognize my dog, but I don't want to recognize every little hair. I don't want to be conscious of all that stuff. So as the information flows up the hierarchy, you lose data, and you tend to throw it away at different levels as you get to concepts at the top. The hippocampus sits *next* to the neocortex physically and conceptually. And it looks like the hippocampus is very good at forming instantaneous memories. It's like saying, "I can take a flash recognition of something and just remember these bits" or "Remember this data right now." And what the hippocampus can do, is that it can look at the state of the neocortex, and remember the entire state of it. It's like saying, "OK, I want to remember what's going on right now." And it can remember parts of the top and parts of the middle of the neocortex, so it can remember a lot of data that the neocortex itself doesn't remember. It saves it off to the side, in some sense, into another little memory structure. And what the hippocampus can do is it can then reinstate that memory in the neocortex. It's like saying, "Go back to where we were earlier. Reinstate this." And it can basically make you think of details you would have forgotten otherwise. So the way of thinking about this is in the normal way of operation, the neocortex is fairly slow to form memories. And most of the time it doesn't try to remember all the details. The hippocampus can say, "I'm going to remember the details throughout your day. I'm going to remember those details, and I can reinstate them back into your neocortex so you could think about them if you want to." But they only get permanently transferred back to your neocortex if they're important. So let me give you an example. You probably had lunch today- all of us had lunch today or yesterday- and you can probably remember what you had for lunch if you think about it. You say, "Oh yeah I remember I had some rice and a soda." Now if you ask me a week from now, I will have no idea what I had for lunch. And if you ask me a year from now, I have no idea. But that information had to be stored very rapidly someplace. And so, it's stored in the hippocampus. It's like saying, "OK- there are some details. I'm going to store it there." Well if I don't use it again, I'm going to forget it. And if it's not important enough for the neocortex to remember it, it's not part of my world model that I had rice for lunch today. It's not that essential. Now if I had rice everyday for lunch, or something really important happened while I had rice, then I would want to remember it. But the hippocampus gives us a temporary memory that says, "Here's something important. Here's something that happened that probably is not important- I'm going to remember it for a short period of time. And only later, if it really is important, will we transfer it back into the neocortex and reinstate that so the neocortex can learn it over and over again. We've modeled it in our software and it works pretty well, actually. It works just like that.

0:43:09
Note: The times indicated will be wrong if you have downloaded the episode with ads. The second ad would appear here. (This is dynamic and not part of the original recording.)

0:43:17
Ginger Campbell: Since you wrote *On Intelligence*, you've been really focusing on building software that implements the theory. Are we ready to talk a little about that?

Jeff Hawkins: Sure. Yeah. I don't want anyone to get the impression that we've sort of given up on the neuroscience side of this. So since I wrote my book, I've had many many interactions with experimentalists and theoretical neuroscientists about what they liked about the book, and we progressed the theory quite a bit. I haven't given up on that. But I've just decided a number of years ago that one of the best ways to get people to work on these theories, work on the theoretical side of neuroscience, is to turn it into a technology that's useful. And I'm using the Silicon Valley computer world-type of model, and I I can get lots of people working on a technology if it solves some problems. I can get them networking on it much faster than I can through like academic means of, you know, just publishing papers and so on. So I've been approaching this from two sides at once, and we created a company called Numenta, which is taking this theory and implementing it in software, and going and seeing- does that work like the neocortex? We're having some good success with it. A lot of this work came from a colleague of mine, Dileep George who was a Stanford graduate student who did the math, and when I wrote the book, I didn't know how to do it. I didn't know how to build it at that time. But very shortly after the book was published, we figured it out, much of its Dileep's work. We said, "OK, let's build this stuff." And it's exciting, and it's useful, and so I've been spending most of my time, but not all of it, working on this.

0:44:48
Ginger Campbell: Several months ago I interviewed Rolf Pfeifer. He also wrote a book, but his was focused on embodied artificial intelligence, and he said that one of the advantages of actually building things was then you can't cheat the rules. And then you have to see what really works. And from that, you can then go back and have a better understanding of what the brain is really doing.

Jeff Hawkins: Yes, you have to imagine that also [["Carmen Meade"?]] has a great quote I use. I'll have to paraphrase it right now, but he basically says, "You don't really understand something until you can build it." There's a lot of people who have theories about different brain things and so on, but in my mind, if it doesn't match the neuroscience *and* you can't build it, then you don't really understand it. So there's a real advantage in building this stuff and seeing what happens. As we've been building the software and testing it, it helps the theory and it goes back to the biology, because we're seeing things we see in neuroscience that the software is doing that we didn't expect and then all of a sudden we go, "Oh yeah, we find that in the brain too." Theory and experimentation go together pretty well.

0:45:54
Ginger Campbell: Is there anything that you've discovered that's really surprised you?

Jeff Hawkins: There's a couple of things that I won't say they're dramatically important, but you asked what surprised me. And there's a few things that surprised me. One of the things that surprised me is that as we started building these memory models and training them- we basically design a memory system and then we train it on... like, we show it movies, literally, of pictures of things moving and so on. We found that it takes a long time to train them. Much longer than it takes to do the recognition. So it may take me days to train one of our memory systems, but I can recognize things, that is make predictions about things, in just a second or a tenth of a second. But it may take me days or hours to train the system. And this is the fundamental property of these systems and then we talked and we said, "Well, we shouldn't have been surprising at that, because that's what humans are like." You know, I never thought of this. Let me ask a question. Why do we take 18 years to train a child? I used to think, "Well, the brain is growing, the body's not ready." But now the theory says, "No. It takes a long time because it takes a long time." And that's just how the memories are formed and there's nothing you can do about that. Another surprising thing was- as we got into these models, as we're building them, we can tweak the parameters a lot. We can try different things. And one of the things we can do, is we can take these memories and you can make them go between two extremes. One extreme is sort of like what brains do- you don't want to remember all the details but you want to generalize. So I don't remember all the details of everything I've done in my life, but I've learned to model the world and I generalize rather well. But if we tweak it the other way, we can make these memories remember all the details but they don't generalize well.

Ginger Campbell: Just like people.

Jeff Hawkins: Well, it's like savantism.

Ginger Campbell: Right

Jeff Hawkins: People who have savant syndrome, which is also associated with autism, but other diseases as well- they can remember huge amounts of data. There are documented savants who can remember every word in 5000 books. Every word on every page. Now they have a brain like you and I do. It's not some magic thing- they have a brain, but somehow it's slightly different. It struck me as like all of a sudden I have a model for that. I said, "Look, I can make our models do the same thing. It's an easy tweak. You just turn this parameter a bit, and all of a sudden they start remembering all the details, but they don't generalize. And that's what savants don't do- they don't generalize well. They don't get jokes. They don't see the context of things. But they remember the details extremely well. So that was a surprise. It's something that came out of experimentation which I didn't expect. And there were things like that which I enjoyed finding.

0:44:36
Ginger Campbell: I noticed on your website- the Numenta website, I mean, and I haven't really had the chance to explore it in any detail- that you have some software out now for people to use. Is that for developers, or...?

Jeff Hawkins: Yeah, it's computer science-y stuff, but it's actually... What this software lets someone do is to- and that computer scientist or electrical engineer type of person. It's not easy- I'm putting it that way. These are very technical tools. But it allows you to configure a memory system that's hierarchical and works on the principles that we think the neocortex works on, and allows them to train it with some sort of data. So they can make a vision system, they can feed in auditory data, they can feed in financial information, and so on. And what the system will do is it will build a model of that data and try to figure out what the underlying causes or structure that data is. It can then do what we call "inferences," pattern recognition. It can make predictions without the future. And these are very useful things. There's many many problems in the world where people have data- they have sensory data of some sort- and they don't know how to process it. They don't know how to extract useful information out of it. They don't know how to make predictions from it. And so, this software gives you tools to doing that in the way that the brain does. It's just getting started. We had our first set of tools out there, it appeared about a year ago. And we're pretty limited. But we have somewhere between a hundred and two hundred people using it pretty much on a regular basis. We have a new set of software coming out very shortly- within a month- which is much, much better and has much more capability. There's a lot of interest in it. There's many many people who are experimenting with this doctrine and are really looking for solutions to problems that brains can do but they haven't really been able to do them on computers before. And this fits the bill for many of them.

0:50:28
Ginger Campbell: It seems to me like your work is very complementary to the work of the people who are doing the kind of robotics that's based on the importance of embodiment, because that's another aspect of how brains really work...

Jeff Hawkins: Yes

Ginger Campbell: That they're in bodies.

Jeff Hawkins: That's right- it is complementary to that. There's a lot of roboticists who are interested in what we're doing. Yeah I think anybody who takes biology seriously and says, "You know what- if I'm going to build a robot I've got to know something about how biological systems do this, how neurons do this stuff and how brains do this stuff. Anybody who takes that philosophy, we have commonality between them, because that's what we're doing. I believe all the stuff we're building is very biologically realistic. It's not just plausible, it's realistic. It's actually modeling what's going on in the cortex. And so people who have this similar sort of philosophy, it is complementary in many ways. We personally- Numenta- had not focused on robotics, because to build robotic systems, there's a lot of other engineering that has to be done. And we're trying to do things which, things such as computer vision, understanding language, and so on, that doesn't require a body as much as like a robot does. We're picking problems that we think we can attack quicker and easier without a lot of other engineering.

0:51:48
Ginger Campbell: You're building tools that it sounds like eventually the people who *are* doing robotics would be able to use.

Jeff Hawkins: Yes that's right. And the theory- this was not in my book, because I didn't understand it at the time. But I understand it a lot better now and there's some information on our website about this. The theory of the memory prediction framework, which we call- today we have a different term for it- we call it hierarchical temporal memory. But it's the same thing. That theory can explain how it is we generate complex motor behavior. How I'm generating my speech right now. It's all part of the same thing. Your neocortex is involved in motor behavior. That's an area that we'd really love to explore more- embedding these artificial neocortices in behavioral bodies. But we haven't done it yet but we think we understand the theory well enough that we *could* do it in the not too distant future.

0:53:38
Ginger Campbell: So that's one of the things that you're kind of hoping you'll be able to do in the future. What's the unanswered question that occupies your mind now? Is there one or many?

Jeff Hawkins: There are many. Most of them are pretty detailed [laughs]. I'll tell you one that might be interesting to your listeners. This is a brain-related, as opposed to a mathematical type of thing. The theory that we're developing says a lot about how neocortex stores sequences and makes predictions, and so on. And those have to happen by neurons. That's all the brain has. It's got neurons. So all the components of this theory have to be implementing neurons, and one of the things I'm focusing on right now is- neurons have complex dendrites- what they call dendritic trees. This is where all the synapses are. And this is where memory is formed. And there are several ways neurons could implement the basics of the theory that we're working on. And I'm working on trying to really narrow that down. Really try to understand exactly how neurons work. And there are theories about neurons and how they work, but most of those theories are very vague about why neurons have so many synapses. A typical neuron might have a thousand or ten thousand or even fifty thousand synapses. And most neural network theories and most brain theories can't really explain why they have so many. Our theories can- they can explain very clearly why cells have this complex dendritic trees. So I'm trying to come to a better understanding right now, you know, a real theory of neurons. You know, why do they look the way they are, and can I make predictions about exactly how those dendritic trees work? And I think this can be done. We're not too far from it. But that's a problem I'm thinking about. I've been asked to write a paper for the Royal Society Biological Journal. And I want to write about that. So I've been focusing on that just recently.

0:54:39
Ginger Campbell: That sounds like a real interesting question. I look forward to hearing about what you discover. I hope you'll keep me informed. Are you getting pressed for time?

Jeff Hawkins: No, not particularly- no. Not at the moment. But I would like to mention that if people do want to know more about what we're doing, if they read my book and they want to know more about what's happened since the book came out, they should go to the Numenta website and there's a lot of information up there. There are talks that we have up there, there are white papers, there's a lot of description of the stuff that we've been talking about. Some of it's a fairly high level description, some of it's pretty detailed. Anyone can sign up for our newsletter, which you're going to get about once a quarter. It basically just says what's new in Numenta and where we are with our releases and so on. And so people who are interested in that should check that out.

0:56:49
Ginger Campbell: I will put a link to your website in my show notes [add link] so we'll definitely do that. Now I had one sort of off the side question I wanted to ask you if you had time.

Jeff Hawkins: Sure.

Ginger Campbell: I'm sure you're familiar with Antonio Damasio's work with emotions and the importance of emotions for people in making good decisions. I just wonder, since your focus is sort of staying away from that part, and you did a really good job in the book of explaining that you're not trying to make a human brain, you're trying to make an intelligent machine, which are two different things. You're not trying to model “the brain.”

Jeff Hawkins: I'm trying to model neocortex, not an entire brain, yes.

Ginger Campbell: Right. So obviously to make a soluble problem you've had to leave out parts.

Jeff Hawkins: Yeah.

Ginger Campbell: Do you think that when you're trying to model decision-making with the neocortex, this issue of emotional inputs has any importance?

Jeff Hawkins: It surely does. Just to give the listeners some basics of this- there are several parts of the brain that are intimately involved in emotions. There's a part called the amygdala, which has to do with fear, for example. These are older parts of the brain, older than the neocortex. They're generally fairly small. So the amygdala is the size of an almond or something like that. And they interact with the neocortex. But the neocortex itself doesn't mediate these emotions. It's sort of like- when you are fearful, the amygdala emits certain patterns and chemicals and so on that makes your fearful. The basic theory we have is that the amygdala has to learn what are fearful things in the world. And then create actions that are appropriate. So it has to look at what's going on in the neocortex, because the neocortex is how the amygdala or any other part of the brain can know what's actually going on out there. And you see that. You see the neocortex... all different parts of the neocortex project to this little structure called, you know- the amygdala and other emotional centers of the brain. The point of this is that the neocortex itself does not do emotion. It interacts with emotional systems in the brain. It's sort of a rational part of your brain. It says, "I should do this. This is the way the world is working." And the neocortex itself doesn't make you hungry or fearful or lustful or something like that. So we have modeled emotions in our systems in a very very simple way. For example, when we train our neocortical models, when we train them, we have a switch that says, "This is important to learn." And then we can turn the switch off and say, "This isn't important to learn." We don't have to, you know, model all the stuff that these other parts of the brain do, trying to judge what's important to human or not important to human. We just say, "It's time to learn" and "It's not time to learn." And that works very well for what we're doing. If I want to make a vision system or a system to understand language, I can do it that way. Now, if I want to embed this in a product that's making decisions about the world, well- we'd have to have some sort of thing equivalent to those other emotional centers of the brain. If I'm making a system that works like the neocortex, and it's trying to understand, say, financial markets, and trying to make predictions about financial markets and maybe trying to decide whether you should make an investment or not. Well, you have to have some sort of emotional saliency to good investments and bad investments or things that go up or go down. But that's outside the neocortex and someone can build that any way they want. It's not so core to the theory of how brains represent knowledge. You'll never really need to model the specific emotional centers in the human brain, unless you wanted to make something that's like a human. But we're not trying to do that, and so I don't think you'll have to model them in the future. We do have to have some equivalent and today our equivalent is very simple- it's sort of engineered for being a switch.

59:22
Ginger Campbell: Thank you. That answers that question. Do you have anything else you'd like to tell my listeners before we sign off?

Jeff Hawkins: The only thing I would like to suggest is that I believe that the kind of work we're doing at Numenta- and there's other people doing similar work, but the kind of work that we're doing at Numenta is really the beginning of an entire new era in both brain science and computing science. It's sort of like when people built the first computer 60 years ago. We're at that era. We have the basic ideas about how these systems should work and the kinds of problems they can solve, and we're just getting working, and I believe there's going to be many decades of exciting advancements going on in the area. And 30 years from now, or even 20 years from now, this is going to be a huge industry and our knowledge about how these systems work will be so far advanced from where they are today. I mention that because if there are young people listening who are so inclined, whether you are in computer science or engineering or neuroscience, if you're thinking about entering this field, I would encourage you to do that. I would say, "Here's one of the most exciting fields to work in whether you're coming from the neuroscience world or from the computer science world." And if you're just thinking about your career, study up, read your stuff, and I would encourage you to get involved. We could use lots of smart people working on these things.

1:00:37
Ginger Campbell: Do you have any specific advice for students that think they might want to get involved in this area?

Jeff Hawkins: Well it depends on whether they're going to get involved in the neuroscience side or in the computer science side. It's easier from the computer science side. I'd just say, become knowledgeable about these things. Study a little bit about neuroscience. Learn some probability theory, because there's a lot of Bayesian mathematics that we use in these things. But mostly you can just start with my books, start with Numenta, read up- we have lots of links to other resources they can read. You can study cognitive theory in college these days, you can study neuroscience. You need to learn a lot of different things. But it's essentially the beginning of the computer era and there's no career path that's well-defined. You just have to say, "I'm going to work on these things. I want to get involved." You know we've had a lot of people who read my book contact me. Still, there's lots of people around the world. And many of them have started an initiative based on the book. Some of them want to work here. We've actually hired some interns who basically said, "I went to college and studied this because of your book." Now they're going to come work at Numenta. So there's lots of opportunities. If you're excited about it, you think it's great, just study what's available and I'm just saying that I believe it would be a great career choice.

Ginger Campbell: Thank you and I really appreciate you taking the time to come on the podcast and I hope I'll be able to talk to you again.

1:03:24
Ginger Campbell: I really enjoyed getting to talk to Jeff Hawkins and I will have a link to the Numenta website in my show notes at **brainsciencepodcast.com**. Don't forget to check out the latest episode of ***Books and Ideas***, which is episode 20 at **booksandideas.com**, which is the interview with Delany Dean about meditation. Next month, we'll be having the interview with Michael Arbib, and also our semi-annual 6-month review episode. I would like to encourage those of you who haven't been to the website at **brainsciencepodcast.com** to visit the website. It's really a blog so you can actually subscribe to the feed from the website separately and you will automatically not only get show notes, but also posts that I put up in between episodes, such as links to other good science podcasts. There's also information on the website for how you can contribute to supporting the ***Brain Science Podcast***. There are links to our Facebook group and also to our **Brain Science Podcast Community on Flickr** for those of you who like to share your photography. I've been running into quite a few listeners on Facebook and I would like to remind you if you send me a friend request, it's a good idea to send me a message telling me that you are a ***Brain Science Podcast*** listener. Same thing if you send me e-mail at **docartemis@gmail.com**. Be sure to put something like Brain Science Podcast in your subject line. That way your e-mail won't get thrown into junk by my e-mail filter. I love getting e-mails and I appreciate people who post comments on the website, but the best place to leave your feedback is at the discussion forum at **brainscienceforum.com**. It's a great place to meet other listeners and share your ideas, and it's much better for getting a conversation going. I've mentioned in the past that Diane Jacobs has been doing some transcribing of episodes and she did György Buzsáki episode, which was episode 31. She's just recently finished a transcription of the interview of Kristof Koch, so if you are one who felt that he talked rather rapidly and had an accent, you might want to check out that transcript. I have added a link to her transcript in the show notes for that episode. That was episode 22. Last but not least, let me remind you about the new website, **sciencepodcasters.org**, which is a group blog for science podcasts of all sized. We're still looking for new members, so if you have a favorite science podcast or you're a science podcaster, go to **sciencepodcasters.org**. Check out the requirements for joining, and if you're interested, send me an e-mail at **docartemis@gmail.com**. I want to thank those of you who have

participated in the audience survey. I did get some feedback from someone who found the survey to be quite boring, which I can't really disagree with since it's really aimed at the kind of demographic information that potential advertisers are interested in, but I still appreciate those of you who have done that, and would ask you that when you visit the website, if you haven't done the survey, click on the survey button, which you'll see on the left side bar at **brainsciencepodcast.com**. I appreciate your listening and I look forward to talking to you again in a couple of weeks.