Futures in Biotech, 46: Towards Computers That Think

Leo Laporte
Bandwidth for Futures in Biotech is provided by Cachefly at cachefly.com.

Marc Pelletier
This is Futures in Biotech Episode 46, Towards Computers That Think.

Futures in Biotech is brought to you by audible.com, the internet's leading provider of spoken word entertainment. Get a free audio book download of your choice when you sign up today. Log on to audible.com/biotech today for details.

[Music]

Marc Pelletier
Due to the complex nature of today's show, we have, I have called in for backup. I think the human brain is probably one of the most complex machines in the universe. So, Dr. Dave Brodbeck has kindly agreed to help out. And so don't worry, he has been vetted. Welcome, Dave. By the way, Dave is an Associate Professor and Chairman of the Department of Psychology at Algoma University in Sault Ste. Marie, Ontario, Canada. He is also the host of the podcast Thunderbird Six.

Dave Brodbeck
And five others.

Marc Pelletier
And about five others, that's right. We will talk about them at the end. We will get to it.

Dave Brodbeck
Sure.

Marc Pelletier
Our guest today – we are very, very lucky – is Dr. Terrence Sejnowski. He is an investigator with the Howard Hughes Medical Institute and is the Francis Crick Professor at the Salk Institute of Biological Studies, where he directs the computational neurobiology laboratory. He is also Professor of Biological Sciences and adjunct Professor in the departments of – and correct me if I am wrong – neuroscience, psychology, cognitive science, computer science and engineering at the University of California in San Diego. Welcome to the show.

Dr. Terrence Sejnowski
Well, thank you very much. I am pleased to be here.

Marc Pelletier
That's very, very – going through this introduction, you are extremely busy. This is pretty – I am very, very happy that you are on. And let me ask you, so here's how I am going to start it. I am probably going to ask only two questions today, and so I will make it a little bit more complex. And as a scientist, and let's say there is a – what would be a good number? Three. What are the three most important questions that drive you? And how do you approach these questions?

Dr. Terrence Sejnowski
Well, first of all three is an interesting number. Why did you pick three rather than one or…?

Marc Pelletier

I only handle one. Seeing that literature behind you and actually, I don't know. That's why I asked, is three a good number? Do you have – are you the kind of scientist that approaches 20 questions? Do you prefer discovery science or do you prefer going hypothesis-driven?

Dr. Terrence Sejnowski
No, I actually think that by far the most important decision that anybody will make in their scientific career is the specific questions they are going to approach. And certainly for thesis students in my lab, when they pick a particular problem to work on and maybe spend four more years on it, that really can shape their entire life. So it's very important to think through carefully what those questions are.

But the problem is, though, that when you start coming up with really big questions like what is consciousness – which by the way is not in my top three, but it is a question that is very exciting and clearly is a problem that a lot of very smart people have thought about over centuries, including by the way my former colleague, Francis Crick. That was the central question that really drove him for the last 20 years of his career. But you also have to ask the question, you have to ask yourself whether or not the question you are asking is answerable in your lifetime or is approachable with the current techniques and with what you have available as an individual.

So in terms of my own career, it has really spanned now several decades and I would say that the questions have changed, but they have all focused around the same central issue, which has to do with the mystery of the brain and how it is that we think, how it is that we interact with each other socially, how do we remember things, how do we solve very difficult problems that may seem intuitively very simple? Like for example, being able to look out and see things, recognize objects and people and react appropriately within 100 milliseconds? That's staggeringly difficult. No computer program has ever been written that can duplicate human performance in those areas, and we take it for granted.

[6:10] So in terms of – if I had to pick one problem or question that – about the brain that is my favorite one, I would pick what's going on beneath the level of consciousness, what is it that we are not consciously aware of that allows us to solve all these very difficult computational problems in vision, motor control, memory, so flexibly and efficiently and rapidly. That is I think really a central problem. Consciousness often has been likened to the tip of an iceberg above the ocean. You only see that little tip, but what's holding the rest of that tip up is this enormous submerged block of ice that is really doing all the heavy lifting and in a similar way your unconscious processing that is occurring is really supporting, doing all the heavy lifting for consciousness.

One of the major research themes in my lab over the last 20 years has been, for example, sleep. Why? Well, it's something that we all do eight hours a day and there is absolutely no agreement amongst scientists as to the function of sleep, something so important, but we know so little about it. But it's really important from the perspective of unconscious processing because it's a time when consciousness by and large goes offline, when you no longer have a little tip above the iceberg except perhaps when you are dreaming during rapid eye movement sleep. But now we can actually see the brain performing we think very important computations that are purely unconscious. I do not know if that's the answer you wanted for your question.

Marc Pelletier
So this is where – I am a molecular biologist by training and my job, so I think these are questions that I wouldn't even understand on how to start addressing them, right? To decide what this – and I am glad that I have Dr. Dave Brodbeck here, psychologists who can approach these questions because he is traditionally – that's what he does. How do you break it down and what kind of methodologies would you do to start understanding what are the underpinnings? And maybe you had it there in your answer, during sleep, by studying sleep then you eliminate that major portion of what consciousness is. Does that make sense? My question I guess is kind of roundabout, how do you go about, how do you – what's your, what are the methodologies you guys use in your labs and how do you formulate the questions around those? How do you break it down?

Dr. Terrence Sejnowski
Right, that's really the perfect question to ask because it's really – science is really limited by methodology and it really – having a new technique or a new instrument really opens up new questions that you can approach. And in my case...

Marc Pelletier
Do you design methodology? Sorry to interrupt. We have got about a half second lag here, so it's kind of difficult to interject. But so, do you design – do you ask a question and then find out the technologies to get to that question or do you prefer keeping attuned to the technology and then finding out, okay, well, that now allows me to answer the next question.

Dr. Terrence Sejnowski
It is both. In fact, it's precisely the questions that drive the development of new technologies and then once you have it, you can look at not just that problem, but open up – that opens up many other avenues of inquiry. In my particular case, I was very fortunate to be born at a time when digital computers were on the scene and digital computers now, as you know, have been increasing in power exponentially over the last roughly 60 years. And they've reached the point now where they are really driving an enormous amount of research in neuroscience in two different ways.

First, having large-scale digital computers with very large memory capacities has allowed us to do experiments and record from many neurons simultaneously, for example. Previously we could only do it one at a time, but now we can use optical techniques to go into the brain, record from hundreds of neurons at the same time and store that information and then to analyze it using advanced signal processing and machine learning algorithms that allow us to sort out from that in a way no human being could have done by just looking at the raw traces.

So that's had enormous impact on understanding how the brain is organized and how it processes information. But for my particular research, which is really focused on trying to come up with computer models of the brain, having the ability to scale up the models from a very simple, small networks with very abstract properties to very large-scale networks with very detailed properties has really given us enormous advantage in being able to try to really get to the core of what are the computational resources at the biophysical level, at the network level and at the systems level. So my lab works at all three of those levels, and without computers we would be dead in the water.

Marc Pelletier
Let's start with…

Dave Brodbeck
[11:38] When you look at computing, the sort of modeling stuff, which I mean for me when I was in graduate school, a lot of people were talking about this and I remember sitting around the sort of seminar table and people like Endel Tulving, because I went to U of T, people like Endel Tulving wondering, does – if the model does – looks like human behavior, like if you put something in it and then the model does something, and then what comes out looks like human behavior, do you have – are you – can you be pretty reasonably sure that the processes that you have modeled work the same way?

Dr. Terrence Sejnowski
Well, that really is again a very excellent question and…

Dave Brodbeck
Excellent.

Dr. Terrence Sejnowski

The answer to it really depends on the behavior. In other words –

Marc Pelletier
Right.

Dr. Terrence Sejnowski
I will give you an example of one where I think we have nailed it and then I will give you another example of one where I think we are far way.

So let's look at honey bee foraging and learning. Honey bees are probably the most sophisticated, champion species amongst the invertebrates. Why? Because they have to go foraging many kilometers from their hive and relative to their body size, it's like going from San Diego up to Los Angeles. Flying out, looking around for food, trying to identify a good food source. For example, flowers that – may have a very brief time during the spring when they have nectar that the bees need. The bees have to identify, by sampling from a wide range of different flowers, which are the best and then communicating that information back to the other bees at the hive. I mean, tremendously complicated behaviors that are involved here.

Dave Brodbeck
Yes.

Marc Pelletier
Considering how difficult that is, just to highlight it, when somebody gives you the directions to the gas station that's five blocks down, I always get lost. So, definitely understanding which flowers?

Dr. Terrence Sejnowski
Well, that's interesting because with regard to navigation, there's actually two different brain systems that are used – can be used, and the one that apparently is failing in your case is the one that is capable of processing symbolic information. But can you read maps?

Marc Pelletier
I study maps to a point where I can visualize them with total recall.

Dr. Terrence Sejnowski
That makes my point exactly, is that it's possible to navigate with this internal spatial map and you're using there your internal guidance system. And different people have different – rely on one or the other of those brain systems and I have known people who are extreme cases in both directions.

But in the case of bees, it's pretty clear that they keep track of their environment in a variety of different ways, from the perspective of direction, the sun heading for example, polarization, there are many cues that they've identified – I should say, that neuroethologists who study bees have identified, and somehow the bee's brain is able to put all those together and compute where things are, how far away they are and then using a waggle dance to communicate all of the parameters to the other worker bees that are going to go out and then exploit that resource.

Now I am bringing this up because around a decade ago, two very talented post-docs at my lab Peter Dayan and Read Montague got very interested in that problem and they really I think identified a core component of that system which really explains an enormous amount of bee behavior in a totally unexpected way. It came from an experiment that was being done by Martin Hammer in Randolph Menzel's lab in Berlin. I happened to be visiting and I happened to be overseeing looking at Martin as he was recording from a single neuron in the bee brain. Can you imagine this tiny little bee sitting there and they are putting an electrode into the one tiny little neuron, they're recording action potentials from it. It's a tour de force experiment but what Martin had determined was that this little neuron called VUMmx1 ventral unpaired median neuron in the bee brain was keeping track of the reward values of the various types of odors and tastes, visual, colors and shapes of flowers.

In this particular case, let's take a particular…

Dave Brodbeck
A single neuron?

Dr. Terrence Sejnowski
A single neuron was keeping track of that. Here's how it worked.

Marc Pelletier
8-bit, 8-bit.

Dr. Terrence Sejnowski
Well, each of these neurons in the bee brain – there is only about a million of them but they are all extremely specialized. But this particular neuron was really sophisticated because it first of all responded extremely well to sugar water, right. It was hardwired to respond to a rewarding stimulus. So that's good. It tells us that it has something to do with the reward system. But now let's give it a neutral stimulus. Let's say let's give it the smell of a rose. Now by itself the smell of the rose has no impact or little impact on the neuron. Why? Because that rose may or may not be a good source of food. However if you pair the smell of the rose just before you give the sugar water, then after one trial learning –

Dave Brodbeck
Wow!

Dr. Terrence Sejnowski
[17:37] The rose smell now fires the neuron. So that's an example that was called classical conditioning and that's in a case of a single neuron. Now the single neuron though, VUMmx1, projects very broadly and very widely throughout the entire bee brain into regions that are important for sensory processing, for memory processing and for the motor output. So in a sense it is able to transmit important information about what the bee – this particular sensory stimulus and get the bee's motor system working for example should I land on the flower or not.

And to make a long story short, we developed this model that was based on temporal difference learning which is algorithm in reinforcement learning from engineering that Rich Sutton and Andy Barto had developed a decade – actually back in the '80s, right. So this is back 25 years earlier. But they had developed it originally as a model of classical conditioning but then they went on to apply it to many difficult control problems in engineering. So we applied it to the bee VUMmx1 system and to really get to the core of it what we discovered was that this particular neuron was giving information to the bee about the predicted value of the reward and that was being modified online as the bee went from flower to flower to update its internal model to tell you whether or not this color or this odor was predictive of getting the right reward.

Now temporal difference learning now – to jump now from the bee to humans, in this particular neuron, the neurotransmitter was called octopamine, which is a biogenic amine in the same family as dopamine, which is a very powerful reward signal in our own brain in all mammals in fact. In fact all vertebrates to my knowledge actually use this substance dopamine. And it is the substance by the way that is hijacked in all known addictive drugs and addictive behaviors like gambling.

Dave Brodbeck
Yes.

Dr. Terrence Sejnowski

But what is the dopamine neuron telling you? And this is from Wolfram Schultz's work on monkeys. It turns out that the dopamine neurons in fact are following the same algorithm, although it's much more complicated in terms of the brain around it but they are basically predicting the future reward that you are going to get if you follow a particular sensory stimulus or if you – do a particular action – what is the reward value that you're going to get from it and that information is broadcast very – diffusely throughout the entire cortical mantle and we know that there are signals up in the parietal cortex that are involved in sensory-motor transformations that are also carrying information about the predicted reward value and that really has spawned a whole new area of science now called neuroeconomics and it's a – some say I unholy alliance between economists on one hand who know quite a bit about how complex financial systems – I presume or..

Marc Pelletier
Don't work, yeah.

Dr. Terrence Sejnowski
A neuro-scientist who I presume know lot about how the brain is organized and by getting together – it's actually I think solves an important problem which is how do you come up with the global answer, solutions, given that you have many piece of informations scattered throughout the entire brain and how do you weigh them against each other and how do you finally use that to make a decision and it really looks now as if it's the dopamine system that's the common currency in the brain for being able to make those comparisons and also ultimately that information is being used to make high level decisions.

And by the way this also goes back to the comment I made earlier about unconscious processing. Psychologists who are doing experiments on decision making when there is a complex decision that needs to be made between many variables like, for example all the different issues that you have to look at when you're going to buy a car or buy a house – there's hundreds of variables that you have to look at.

It turns out that the more complex the problem and the more incommensurate variables you have to include then the more these unconscious mechanisms that weigh these different dimensions and then try to reduce them – that's to say, to combine that information in one final common forum. That now seems to be a part of that – the same basal ganglia dopamine projection that the cortex projects down subcortically and then that information then comes up and helps to guide behavior and much of what we call consciousness now it turns out to be – more of a – after the fact rationalization. Another …

Dave Brodbeck
Right.

Dr. Terrence Sejnowski
Some part of your brain basically has figured out what's best for you and now you've got to explain yourself why you did it. So it's changing the way we think about –

Dave Brodbeck
For sure.

Dr. Terrence Sejnowski
The function of the brain.

Marc Pelletier
So There's no free will.

Dr. Terrence Sejnowski
Well there's – it's interesting I think there is a will all right but it's not free. You got to pay for it.

Dave Brodbeck
[22:50] That's the – I love that. I think I am going to start quoting you in my – I'm starting to teach intro-psych – well, gee, I guess in four days and I think I am going to use that line. It's funny that the bees are sort of online taking notes of what – the reward value of something is with a single neuron. It takes us a whole nucleus accumbens to do that, so just goes to show again what a friend of mine, Ken Cheng, always says that the bee is cognitively the most – pound for pound the most cognitively complex animal on the planet.

Dr. Terrence Sejnowski
That's the –

Dave Brodbeck
Yeah. I mean there – I love everything about the bees except the stinging – I could do without that. They – I am frightened, deathly frightened of them.

Dr. Terrence Sejnowski
Now why is that, have you been stung?

Dave Brodbeck
Once.

Dr. Terrence Sejnowski
One trial learning. That's very good.

Dave Brodbeck
Yes, exactly, one trial learning.

Marc Pelletier
One trial learning.

Dave Brodbeck
I watched my dad for years be frightened of bees and I think that what's they did it for me – a little social learning on that one. The thing about something like a bee and something like a human is we've got two clearly very different animals but we have pretty similar – a neuron is a neuron is a neuron is a neuron, pretty much.

Do you think that do we necessarily always have to be looking at – say human behavior or humans in general – I mean whatever rather than we can look at something like – rats, pigeons, bees, drosophila whatever. Is it – do you think that we better data out of – or more important data I guess is the better way to put it – out of the non human animals than we do out of humans?

Dr. Terrence Sejnowski
We're able to do experiments on non humans that can't be done on humans, so…

Dave Brodbeck
Sure.

Dr. Terrence Sejnowski
So it's absolutely essential for example the biophysical level, the properties of ion channels and the way as you said neurons are constructed – that's common to all of these animals, although there are differences between species that we're just discovering. But the point is, it's all built with the same genes and it's all more or less variance of the same classes, categories of ion channels and ways that the neurons are put together during development. But there are some things that have to be studied in humans, for example language.

There's nothing exactly to humans – to human language although there are some echoes of it, for example in song bird learning where sequences of notes are learnt from the conspecific males. In fact that's another area that Kenji Doya, who, when he was in my lab, used reinforcement learning to try to model and actually now that's the existing explanation for how they can do it. Interesting that it's all the basal ganglia machinery that's allowing us to do that. I think that it could well be that all these wonderful language abilities that are going to be ultimately – localized through understanding more about the basal ganglia which has been neglected compared to the cortex.

Dave Brodbeck
For sure.

Dr. Terrence Sejnowski
And – but the key thing though about humans is that we have evolved some abilities – not just language, language is just one of them – which really has changed the game for our ability to – for example to create very large societies and cultures. So, we can create our own environment and if you think about it for a second, if you have an animal that can learn the statistical properties and the contingencies – reward structure of the environment, by manipulating that reward structure, by creating schools for example or sports in which you have to compete and there's a big prize to be won if you're the best in some thing. You have an enormous power to just, completely alter the whole way that – that humans interact with each other, what they value, what they're going to be capable of – of assimilating, for example – I don't think nature could have possibly anticipated something like mathematics and certainly the brain wasn't built to really understand itself. Right.

Dave Brodbeck
Right, right.

Dr. Terrence Sejnowski
There was no advantage to that. It was more important that you behave in a way that is going to ensure survival but – being able to introspect to how it works is some thing that – that just wasn't worth evolving.

Dave Brodbeck
Yeah, it's true.

Marc Pelletier
Well, yeah. That could lead to a singularity though at one point right. If we have a complete understanding of the human brain, we've now changed our fate, our will, our understanding of free will.

Dr. Terrence Sejnowski
Well, that's going to happen. In fact, this is why I think we're living during such an exciting era is for the first time in history of our species, we have within our grasp the tools, computers and the optical techniques, genetic techniques that are going to go in and actually figure out how the whole system is put together. How it develops, how it adapts, how we interact with each other and social interactions are absolutely paramount in terms of our cultural abilities to create new disciplines to develop – the highest levels of human performance and all the different areas of music and art and so forth. We're going to understand that it's going to happen within the next generation because we have the tools, we have the facility, the – I think the basic toolkit.

Dave Brodbeck
Technology.

Dr. Terrence Sejnowski
Exactly, technology is really what's driving it, exactly.

Marc Pelletier
And the scientist are getting better and better because their brain's evolving to a – in that direction as well. In the same way that the social interactions are helping us develop.
Let me take a brief second, we have to thank our sponsors here and then we'll get right back to it. I really want to see how this cycle of – brain evolution is actually moving forward and how we've broken the line of the animal – from the animal world. But first, we must thank audible.com for sponsoring Futures in Biotech. They have now over 60,000 titles that's audible books, radio shows, interviews, pictures, an enormous library and they work on your iPod, iPhone and BlackBerry and over 500 devices. One of which Leo is probably owned one of each at a certain time or another.

Dave Brodbeck
Just this year!

Marc Pelletier
This past year. Which, I'm glad he does because he tests them out for us.

Dave Brodbeck
Exactly.

Marc Pelletier
Yeah. So Dave, you have a pick.

Dave Brodbeck
Yeah. In fact it's great that we were just mentioning things like music. For example because the book I've picked is "This is Your Brain on Music: The Science of a Human Obsession". And is by Daniel Levitin, he is at McGill University in Montreal. It's also – and home of the Montreal Canadians, Montreal. And the Montreal Alouettes. Just thought I'd mention that and also wear my Mom's from. Hi Mom. I promised her I'd say hi. The thing about Levitin is he is a neuroscientist. But he's also – used to be a record producer and he still does a lot of consulting in music industry. When I say record producer, he isn't a guy with a pirated copy of Protools; this is a guy that has gold and platinum albums to his name.

It's great reading the book when you have – he will be talking about some sort of high end, sort of neuroscience stuff and then on the next page, he is talking about something to do with Sting when he worked with the Police. It's really quite cool.

And, he is one of the first people to really talk a lot of about music. In the book, not only he talk about the evolutionary function of music and what it might be. But he also of course talk about how your – the book's called "This is Your Brain on Music". So he talks about the – how music affects the humans.

And in fact, it's really led to almost a brand new field – sub-field of neuroscience which is the sort of neuroscience of music. Indeed we just hired a guy here at Algoma Duane Kia [ph] who'll be finishing up his PhD shortly. And Duane is right into this stuff, looking at music perception and this kind of stuff – I don't think the work could have been done without Dan Levitin doing this kind of work. So it's very cool and…

Marc Pelletier
Now, you can get answers by listening to Led Zeppelin.

Dave Brodbeck
You probably can somewhere. Well you know, perhaps, if that's the case I should have a PhD in U2.

Marc Pelletier
Well. Hey, that's a really great pick. I'd just like to say that if you'd like to download this book for free, only you need to do is sign up for a 15 day free trial of the audible listener Gold account and it's a win/win situation, right? If you like audible.com, you keep the free title and you stay signed on. But if you decided that the subscription is not for you, you get to actually keep the book and you just cancel it before the 14 day is over. So if you'd like your free book head over to Audible.com/biotech.

Dr. Terrence Sejnowski
By the way, it's a fabulous book,

Dave Brodbeck
It is, isn't it?

Dr. Terrence Sejnowski
Yeah, good choice.

Dave Brodbeck
Yeah, it's – thank you.

Marc Pelletier
Well, is music one way to – animals don't have music. Is music, how would you approach music or would you like to talk about music? Or – you were talking about social – how social aspects of human species, as a species is influencing the function of the brain and the development of the brain?

Dr. Terrence Sejnowski
[32:52] Well, I'm not an expert on music and I do follow the literature and I know that for example, using brain imaging, it's clear that the brains of musicians are not the same as the brains of people who just listen to the music passively. They also say that they developed specialize circuits that allowed them to think about what they doing and also to create new music. I mean if you think about the process of creativity is very, very much to the fore there.

But I'd like to use it – just as one example of many differences that you can point out between humans and even our closest relatives, the chimps and bonobos for example. I happened to be part of a group here at UCSD, which started being about 15 years ago, it was called La Jolla group on the origin of humans, and Ajit Varki was the inspiring leader and the Mathers Foundation supported it and we would meet regularly.

And the goal was to try to come up with a list of all the differences, I mean that we could come up from all of our difference perspectives from the perspective of paleontology you know, how did we get here, what do we know about the life early in – for the hominids? What do we know about the conditions under which perhaps language developed? And now in fact that group has graduated and is now an organized research unit, which means, it's a part of the – formal part of the structure of research at UCSD.

And one of the really strengths of that particular group is there from the very beginning, it was eclectic from the point of view of bringing fields, people from fields as diverse as linguistics together with people who study the chimpanzees' behavior – from people who studied neurons or model like – or modeled neurons, like my lab. And it's an ongoing group now called CARTA [Center for Academic Research and Training in Anthropogeny], and it's on the web if you want to take a look at it.

And what's happening right now is that all of the experts are writing – it's going to be like the Wikipedia for human behaviors, and differences between humans and other species that I think is, now, now that we have genomes for a variety of different primates, for example we have the chimp genome and we can compare it to the human genome and now we see the differences and what we need to do is to translate those differences into differences in the phenotype, the differences in the behaviors, the differences in the way that we behave and act. And that's non-trivial and it's really going to take an enormous effort by all these different experts in different areas.

Dave Brodbeck
I think calling it non-trivial was probably a bit of an understatement. It's funny. Thinking about the kind of stuff you're talking about here, because taking a look at the kind of work that you do which is – pretty hardcore neuroscience. And it allows us to understand things about – the biggest questions like where do we come from, why we exist? You know the big evolutionary questions. Do you – in your work, does that drive you at all? I mean do you sit – or is it like – are those almost emergent properties? When those questions end up getting answered a little bit, those sort of bigger, almost evolutionary questions. Is that something that you think about when you're designing, like a research program or is it like: Oh, this is cool, I never thought of it this way.

Dr. Terrence Sejnowski
Well, this is a very good question of what drives Terry Sejnowski. And I really don't think I'm the best person to answer that. Because of the fact that – I don't think we have good self-knowledge about what is that is behind the really important decisions we make in life. It could be, I think you mentioned earlier was your father's experience with bees got transmitted down to you. I'm sure that there is some parental influence on the sorts of questions that I'm looking at.

But, at a day-to-day level, what drives me is the excitement of seeing what's new under the next rock. I mean I wake up in the morning and I have got a terrific lab and they are constantly discovering things and it's really – very satisfying and you know it actually is not just my lab, I'm talking about the community. In the area for example, that I'm working in, competition neuroscience didn't exist 25 years ago. And now it's a very thriving community of several hundred people and we get together and we have meetings and there is a sense of shared adventure and we work very closely together with experimentalists and so this is you know, what could be more exciting than trying to understand how the human brain works.

Marc Pelletier
Let me ask you if you – are you identifying fundamental processes in the way memory and learning works, and then going across with crossover departments identifying premises or sort of theorems in engineering that say well, we have already got algorithms for that. Is it are we seeing a – in our – as we deepen our understanding of how the – how memory and well, even the physiological processes work, are we seeing a correlation to the way the way we have designed computers?

Dr. Terrence Sejnowski
Absolutely and this is really what the goal of the Institute for Neural Computation is which is an organized research unit that I direct, co-direct with Gert Cauwenberghs at UCSD and our goal, explicit goal is to take what we know about how brains work and then translate that into genes that have similar behaviors, similar capabilities. The…

Marc Pelletier
With no faults.

Dr. Terrence Sejnowski
[39:19] I'll give you one example, again a specific example building on the bee. So reinforcement learning is now a huge area of engineering. As computers get faster, actually the capabilities get better. But still wasn't – reinforcement learning wasn't capable of really handling the really difficult control problems and there is a meeting that I organized with Rodney Douglas that was sponsored by the NSF and this was back 2007 in which we bought together psychologists who

are experts on human learning and neuroscientist together with machine learning experts, people who were building analog VLSI chips that would emulate brain circuits – people who are doing with reinforcement learning and something really exiting happened at that meeting which was that Andrew Meltzoff from University of Washington – he studies child development was giving lecture about how phenomenally good children were at imitating.

In doing experiments, carefully controlled experiments, he was able to show that at a very young age, already a few years of age children would be able to mimic not just the actual not just the putting your hand out but trying to accomplish the same goal, that the adult would have that would be – serve as a model. And Andrew Ng from Stanford who works in reinforcement learning said: Wow, that really is a very powerful way and wouldn't it be wonderful if we could supercharge our reinforcement learning algorithms and he developed something called apprenticeship learning.

So here is how it works. He, there is – helicopters are extraordinarily difficult to fly because the control for, are very counter-intuitive in terms of how you have a control system that you are trying to make the thing for example fly in circles, or inverted flight and it is not at all obvious even after you are very good at it you are still right on the edge of destruction of losing it.

But there are humans who have mastered this obviously on real helicopters but on toy helicopters you know that are like three foot across. There are some humans who are so good that they can do acrobatics and they can fly upside down and they are just incredibly good. So what Andrew did was to basically train his reinforcement learning algorithm to mimic the human expert on the very same helicopter and he used it in a very clever way to not just come up with a better control law but also to improve on the model of the helicopter that the control system was operating under the assumption that their model correct, was completely wrong. You turn helicopter upside down and all the forces in the bearings and everything have changed the control properties. So obviously human had learnt that difference but it took this apprenticeship learning to get to the point.

Now if you go on YouTube you will see these fantastic videos of this automatic apprenticeship system being able to do acrobatics on helicopters. So here is an example of how being able to see one of the tricks that humans have – nature has created for humans. To allow us to be able to bootstrap ourselves up, great complex behaviors can be taken back down to engineering. But I think there is going to be more and more examples of this as we get more and more understanding of what the actual mechanisms and the algorithms are inside the brain.

Marc Pelletier
Well since he has shown it has worked. Now you have got a good digging for those new algorithms in the human brain.

Dr. Terrence Sejnowski
Well here is what the – the good news is is that there may not be that many.

Marc Pelletier
Not in mine.

Dr. Terrence Sejnowski
No, no, no, from the perspective of fundamental algorithms. For example all of life as we know it depends on DNA, right? There are minor variations now that are all centered around the idea that it's DNA that is carrying forward the blueprint for the species and for all the genes and so forth. And that's the common principle. So we need to find similar principles, computing principles that are being used in nature and I have already mentioned one now that is very, very robust and that's temporal difference learning, we found it in bees, it's found in mammals, it's found in humans and it has been found by the way not just in this conditioning reinforcement system but also it's been found at the computational level in the cortex in something called Spike Time

Dependent Plasticity, STDP and it is working there in the window of 10 milliseconds rather than seconds which is reinforcement learning.

Marc Pelletier
10 milliseconds.

Dr. Terrence Sejnowski
So 10 milliseconds and so we

Marc Pelletier
Hundredths of a second

Dr. Terrence Sejnowski
So what's going on. That's one hundredth of a second. What's going on there is that the spikes are coming into neurons very rapidly and then they are integrated and that then is driving the neuron. But it turns out that every time the neuron spikes it updates all of the synapses that have just within the last 10 milliseconds have been transmitting information and its – it increases their strength. And it says that if you happen to help me fire the action potential within the last 10 milliseconds because that's the integration time then you get rewarded. You get your synapses increased.

Now there is a flipside to the coin and that is if you happen to be, if you happen to be a synapse that is fired within the 10 milliseconds after the spike, you get decreased in strength. In other words you get punished and that's Spike Timing Dependent Plasticity.

Well, Raj Rao when he was a post-doc in my lab proved mathematically that this is mathematically identical to temporal difference learning except on the 10 millisecond timescale and using the membrane potentials rather than reward signals. So it may well be that there is a small handful of very powerful algorithms that we can discover that are being instantiated in different brain circuits with different timescales for different purposes but basically, these are the ones that nature has hit upon that actually work. And that is what we are looking for.

By the way this is a – I think the – again let's go back to your initial question about important questions. I think the question about how we learn and the substrate for memory is got to be one of the big ones that were tremendously important for understanding things like education and trying to understand things about what sort of flaws do our memory have. I mean, we don't have perfect memories. It wasn't designed to be perfect like a computer memory, well, it was designed for some other function. Let's try to understand that. And by and large, you know, our memory is associative. We can associate different events that have occurred in our lives, we can associate a person's face with their name, we can associate and fill in patterns, right? Pattern completion.

Marc Pelletier
Yes.

Dr. Terrence Sejnowski
And that's a different kind of memory primitive than simply storing a bit of information at a particular memory location, which is what digital computers do. But, so we now have it in our capability to actually build synaptic memories. DARPA has a major multi-million dollar program that Todd Hylton is running called SyNAPSE, and there is now about a half dozen groups out now partnering up with big corporations like IBM and Hughes. And they are actually going to build, they think, I mean I think it's a little bit ambitious, they actually want to build a cat brain.

Marc Pelletier
Build a cat brain?

Dr. Terrence Sejnowski

Yes, a cat brain.

Marc Pelletier
So this is…

Dave Brodbeck
Using semiconductor, or live neuron, wires?

Dr. Terrence Sejnowski
Well, using you know, whatever new materials you might need. Obviously you start with silicon because that is the best worked process, but there is other materials out there that might actually be much better and it really comes down to energy. Interestingly, the real constraint is to try to get the whole thing to work in a small volume and not to burn up.

Marc Pelletier
Why not get a cat if they wanted a cat?

Dr. Terrence Sejnowski
Okay, that's true, why not just get a cat, but you know, again –

Marc Pelletier
No, I understand though.

Dr. Terrence Sejnowski
Nature doesn't make it easy for you to figure out what's in the cat brain, but if we can actually build a cat brain we might be able to reverse engineer that.

Marc Pelletier
Well, that becomes one – what's the next step after the cat brain, right? And building a circuit that can handle human language and all the social capabilities, I mean, we are talking a hundred a years ahead of now but…

Dr. Terrence Sejnowski
Well, you know, I wouldn't wait for a hundred years.

Dave Brodbeck
I do not know about a hundred, Marc, you know. No, I would imagine we are getting closer and closer with this kind of stuff…

Marc Pelletier
The singularity that I was talking about, right?

Dave Brodbeck
Yes.

Marc Pelletier
Having a human brain, well, this is Data from Star Trek.

Dave Brodbeck
Yes.

Dr. Terrence Sejnowski
[48:48] Well, let me give you – it's actually – it's going to happen much quicker than you think. In fact, it's already happened here at UCSD that we have put social robots into classrooms. So this is the work of Javier Movellan, and his goal was to try to see whether or not a robot could interact profitably, that is to say, could actually help teach little 18 month old toddlers and a social robot

was not using language but it was using – it had a head and it had eyes, and it had a sound, it could create sound. And the question is, you know at that age, you know. The kids aren't talking either, right, 18 month olds are just beginning to talk, but they can interact in a very intelligent way and that was the real question that Javier asked. What would it take? And he discovered a couple of really important principles. Again, it turns out that you don't need to have a tremendously sophisticated, complex language system in order to be able to communicate.

So how did it work? Well, first of all these 18 month olds immediately attacked Ruby the robot, as he called her, because they thought it was a toy and they ripped the arms off. So the first hurdle was to try to figure out how to prevent that from happening and the way they did it was by having a pressure sensor that would trigger the robot to cry if it was pulled too hard. Now, the toddlers immediately recognized that social signal and backed off. In fact, one of them went and hugged the robot.

So it's clear that, you know, there was immediately – there was an emotional rapport there that you can get between the social robot and the child. Now, the question is, can you get the little kids' attention. You know at that age, their attention span is only about 10 seconds and they throw away something that doesn't interest them immediately. And the key that Javier found in order to get the robots to – the kids to pay attention to the robots is for the robot to respond to the child when the child was trying to say something or do something. If the robot did not respond within one to two seconds then the child would lose interest. If it responded too quickly it was mechanical, if it was – responded too late, the child would just turn away; this is not something that I can interact with. But there is something magic that happens if you get that sweet spot.

Marc Pelletier
Right.

Dr. Terrence Sejnowski
And then finally, this is an important issue that has to do with a specific human ability, which is shared attention, that many other species don't have. Even the chimpanzees, for example, our closest relatives, don't have this ability; dogs, interestingly, do. Maybe that's because we've been around with them, co-evolving for so long.

So here is what a group – what would have happened – a little kid would point to an object, now – a ball. Ruby now within two seconds has to attend to the ball, either by moving its head and pointing toward it with its hand or its eyes and once that's established, once the little toddler sees that this funny looking object – because it has the body of a computer case and it has a panel like a Teletubby and everything, it's a very odd looking creature. But if shared attention is achieved then suddenly the child is communicating; the child now understands that there is something here that is really interesting and now you can begin the process of interactive education. And Javier used it to teach songs and he is trying to find out now whether it can be used to teach language, for example.

So this is I think a real breakthrough because it is telling us that in order to get off the ground, you know, there are some basic things you have to do right. But once you've gotten past that point, then you've entered a different realm of being able to communicate and to transmit information back and forth. So it's all about emotional bonding, it's all about shared attention and it's about imitation. And Javier and I together with Andy Meltzoff, who I mentioned earlier, and Pat Kuhl at the University of Washington, wrote a review paper that was published last month – actually in July in the – Science magazine entitled "Foundations for a new science of learning" where we brought together all of these ideas and tried to lay out a research program that is going to take us now into that area where I think we'll be able to improve teaching in the classroom by bringing in assistants, automated assistants that are able to track every individual student.

Marc Pelletier

I think that your approaches are going to have an unbelievable level of influence on education, for sure, I mean, that's – and not to mention I suppose criminology and sociology as well, as we understand these – get answers to some of these questions. But I think we have really only touched the – you know, it's been like that tip of the iceberg, you know, in terms of understanding consciousness. But we have to close the show right now, so, but maybe we could have you back because I would certainly like to talk about how the computational aspects – how, you know, in maybe even in a more technical way, how those algorithms are being coded into machines and what you guys are doing in the lab at that level. That'd be a lot of fun.

Dr. Terrence Sejnowski
Sure, I'd be happy to come back, and there is just so much more that we could talk about.

Marc Pelletier
I'd like to talk about electrophysiology as well, and I'd like to talk about genetics as well. I know you cover – your work covers every aspect of the human brain, which is just unbelievable. So I really appreciate you coming on the show.

Dr. Terrence Sejnowski
You are very welcome.

Marc Pelletier
I'd like to, in as a closing here, thank Dr. Sejnowski for joining us today. He is the Francis Crick Professor and Director of Computational Neurobiology Laboratory at The Salk Institute for Biological Studies.

I'd also like to thank Dr. Dave Brodbeck, Chair of the Department of Psychology at Algoma University in Sault Ste. Marie, and host of Thunderbird Six among his one of many podcasts, including his famous Psych Lectures which you can find at thunderbirdsix.org, that's one word, thunderbirdsix.

Dave Brodbeck
Yes, there is links to all the other shows including the lecture podcasts on that page.

Marc Pelletier
Awesome. It's a really cool page, too. I like your design. Very retro.

Dave Brodbeck
Thanks, that was the point, yes.

Marc Pelletier
Thanks for coming on, Dave.

Dave Brodbeck
Always a pleasure.

Marc Pelletier
I appreciate the help. I'd also like to thank the great folks that made this possible at twit.tv, Leo Laporte, Dane Golden, Colleen Kelly, Erik Lanigan, and at Salk Institute, Bryan Nielsen.

If you would like transcripts for the show, they are available at futuresinbiotech.com. Thanks to the kind folks at Pods in Print. If you need transcripts done they can do even the most technical transcripts. So you can visit them at podsinprint.com.

I'd also like to thank Phil Pelletier and Will Hall for the opening and closing themes.

Futures in Biotech, I'm Marc Pelletier. Thanks for listening.