A Conversation with Dora Angelaki

INTERVIEWER: GARY STIX

Senior Editor, Scientific American

Dora Angelaki is the Wilhelmina Robertson Professor and Chair of the Department of Neuroscience, Baylor College of Medicine, with a joint appointment in the Departments of Electrical and Computer Engineering and Psychology, Rice University.

Gary Stix: You gave a talk here yesterday at the conference on how neurons that process sensory information can predict how an animal behaves during a simple perceptual task—I think it was a monkey moved in the world? Can you tell us briefly what the field of multisensory integration is attempting to do? Understanding how one sense works is difficult enough, so I assume that trying to integrate information from multiple senses is even more difficult. Are neuroscientists so attracted by this problem because it's a really tough challenge?

Dr. Angelaki: This is only part of the answer. In the laboratory, we like to simplify the enormous task of understanding how the brain works. Traditionally, neuroscience has studied one sensory system at a time. But in the real world, this is not what happens. As we talk now, I look at you, I see your mouth moving and I hear your voice. This experience is multisensory and our brain is optimized to process all these signals simultaneously.

Gary Stix: You don't stop and say I'm about to look at you.

Dr. Angelaki: No, it all happens together and my perception of you, of the world, of what you say, is based on many sensory signals all acting in unison simultaneously. How that happens in the brain, where it happens in the brain, what computations the brain uses—these are the challenges we seek to identify and solve. There are also different layers of understanding: different people have approached this question from the point of view of psychology and psychophysics, physiology, and/or normative theory.

How can the brain abstractly combine multiple sensory signals into a unified experience? How neurons or a network of neurons actually do this is largely unknown. We work at all levels simultaneously. We take advantage of lessons learned from theory and psychophysics, and try to apply them to neurophysiology.

Gary Stix: In your talk, you described experiments in multisensory integration that looked at the integration of visual and vestibular information as a monkey is moved. Vestibular information is information that we use to sense how our bodies are moving and how we can retain our sense of balance. Can you tell us how you went about doing that experiment?

Dr. Angelaki: When we do simple things like riding a bike or driving a car, we use many sensory signals, including input from our vestibular (also known as the balance system), somatosensory, and visual systems. As we move in the world, our retina sees the velocity of things moving relative to us, our accelerometers in the inner ear are activated, and we receive proprioceptive information from our muscles and joints. For example, when we are in a plane, we feel that it is moving even without looking outside the window.

Gary Stix: Or even just walking around.

Dr. Angelaki: Absolutely. But we don't separate the perceptions—oh, this is my vestibular sense, or this is my proprioceptive sense.

Gary Stix: We wouldn't be able to function in the environment.

Dr. Angelaki: Not without the brain being able to put all of these signals together to try to figure out how we move in the world. This is important, not just to perceive the world but to interact with it. Sometimes when we walk on the sidewalk, we turn to look at a shop window. We do this and still maintain our heading. We can talk to our friend, move our head around, and the brain manages to do all of this because of the interaction of many sensory and motor systems.

In the lab, we have set up virtual reality systems. We study human perception, as well as those of animals, macaques and rodents, to try to understand how neurons in the brain combine multiple sources of complementary, but also often redundant, information. In fact, the kind of equipment we use is very similar to what pilots use during training.

Copyright © 2014 Cold Spring Harbor Laboratory Press; all rights reserved; doi: 10.1101/sqb.2014.79.02

Gary Stix: In a flight simulator?

Dr. Angelaki: It's very much like the flight simulators used for pilot training or those used in amusement parks, in which they accelerate you, as you experience visual stimulation, and you feel like you are moving long distances through space, when you are actually only moving a few centimeters. It is the combination of the stimuli sensed by the ear and body receptors, in combination with the visual stimuli sensed by the retina that gives you the illusion that this is very real. This is what we're doing in the lab. We record brain activity during virtual reality, which allows us to dissociate sensory signals, and thus to understand how we orient and navigate in the world.

Gary Stix: There was one thing that you talked about that I thought was particularly surprising. You looked at the activity of a single neuron and found it could predict what the subject might perceive.

Dr. Angelaki: It is indeed surprising. We knew from other studies that in trained animals, one can find neurons in the brain that are as sensitive as our perception. So if a single cell can do it so well, why do we have so many? In our experiment, rather than recording only from the cortex, we also recorded from the actual receptors from the inner ear, from the vestibular nerve, as well as neurons in the brainstem and cerebellum. We found that even single sensory afferent fibers are as sensitive to motion as our perception. We have 3000 of these neurons carrying information from each motion receptor in the ear to the brain. If one of them is so good, then why do we have so many?

Gary Stix: A single neuron could predict what?

Dr. Angelaki: There are two types of analysis. What I've described so far is how the neuron responses correlate with the stimulus. But, in addition, we record from neurons while the animal performs a task. Then we can also correlate the activity of those neurons with the perceptual choice that the animal makes.

Gary Stix: How the brain decodes the information?

Dr. Angelaki: That's it exactly. The first process is what we call encoding—how the neurons map the sensory world into neural activity. The second process is decoding, which is how the brain takes the activity of all the neurons and generates our perception of the world—critical to orient and be able to initiate an action. We don't know where this happens and any of the details, but that is what we are looking for.

Gary Stix: From that one neuron in the inner ear all the way into the cortex, you could predict what was going to happen?

Dr. Angelaki: Not exactly. It happens only with neurons in the central nervous system. All the neurons are nearly as sensitive as the behavior, but signals about decoding are found only centrally. What is interesting is that neurons with significant correlations with the behavioral choice are not only found in the cortex, but also lower in the brainstem and cerebellum, which are the very early sensory areas that process vestibular signals. You can listen to these neurons and predict what the animal's choice will be. This has many implications about how our brain works and why it doesn't do better. I presented a model, actually several classes of models.

Gary Stix: You're interested in theory and how it might apply to neuroscience and how it might help neuroscience understand the complexities of the brain.

Dr. Angelaki: Absolutely, because I strongly believe that without theory we can never understand how the brain works.

Gary Stix: Some neuroscientists are wedded only to experimentation but you seem to be doing both.

Dr. Angelaki: Yes. For me, it's extremely important to involve theory in all aspects of neuroscience. Not only do I think it's extremely important to use theory to understand how behavior is generated, but I also think that theory is important to connect the different levels of analysis from molecules to the network, to the brain and behavior.

Gary Stix: There is a lot of the emphasis now in neuroscience on gaining this big picture. Do you feel it's necessary to emphasize theory in the projects being set up here and in Europe to try to gain this bigger picture?

Dr. Angelaki: Absolutely. The existing approaches focus primarily on technology development so that we can see as big a piece of the brain at once as possible. That's extremely important and a very valuable goal. Some people make a parallel between neuroscience and astronomy. It was only after the discovery of technically excellent telescopes that we were able to make major advances in our exploration of the cosmos.

This is what is needed in neuroscience. But I don't think astronomy would have been as successful if there hadn't been the appropriate theories able to know what to search for and how to interpret massive amounts of data. Without knowing what to look for and what hypothesis to test, massive data can be overwhelming and not appreciated to their true potential.

Gary Stix: One other thing related to the theory. Turning off a neuron to see whether it causes a deficit in the animal is a standard technique in biology and it's heavily relied upon by neuroscience to prove cause and effect for a particular phenomenon. Yet, you said that, in some cases, that doesn't necessarily tell you that much about the function of a neuron. That seems really surprising.

Dr. Angelaki: That was indeed a provocative statement. I want to emphasize that there are many cases where this approach is extremely helpful, particularly when a causal effect on behavior can be demonstrated. However, for negative results (i.e., those where turning neurons on or off lead to no measurable effects on behavior), we cannot necessarily conclude that these neurons do not participate in the behavior. It's a negative result I emphasized because I presented an example of such a situation in which we inactivated a group of neurons and, using theory, we could predict both neural and behavioral results regardless of whether or not these neurons participated in the particular task we studied.

Gary Stix: I imagine that provoked a lot of discussion after your talk.

Dr. Angelaki: A lot, and many students came to talk to me afterwards and some of them disagreed, which is great. This is what science is about.

A Conversation with Cori Bargmann

INTERVIEWER: JAN WITKOWSKI

Executive Director of the Banbury Center
at Cold Spring Harbor Laboratory

Cori Bargmann is Torsten N. Wiesel Professor and Associate Director of the Shelby White and Leon Levy Center for Mind, Brain and Behavior at the Rockefeller University and an Investigator of the Howard Hughes Medical Institute.

Jan Witkowski: Tell us about the system you're interested in.

Dr. Bargmann: My work is with the simplest animal being discussed at this meeting, the nematode worm *Caenorhabditis elegans.* It has only 302 neurons but is nevertheless a real animal. It moves around, it decides what it likes and doesn't like, it learns from past experience and uses that to affect future decisions. And we can understand those kinds of basic processes using this very simple nervous system, and relate what's going on in an individual's brain to its behavior.

Jan Witkowski: So, what aspect of its behavior do you study?

Dr. Bargmann: Most of my lab studies the response to odors, and that's because odors are the things that worms are most interested in. They are so well specialized for detecting odors that 10% of their genome encodes G protein–coupled olfactory receptors. To give you a comparison, humans are thought to be able to smell a trillion different odors, and we do that with about three or four hundred different olfactory receptor proteins. Worms have 2000 olfactory receptor proteins, so we don't even know how many things they can sense. But in any case, this matters to the worm, and because what it smells and what it likes matters to the worm, we study those things to try to gain access to its brain. We have a first-order understanding of the worm's olfactory system. We know what the molecules are that detect odors, we know what neurons those molecules are in, we know which odors the animal finds attractive and why it finds them attractive. So we're now starting to ask the second order of questions, which is why even these simple animals don't act the same every time. When given a particular set of choices, sometimes they'll pick one and sometimes they'll pick the other one. This is true even if you have animals kept under identical circumstances, even if they're getting the identical odor. Even if you have many animals of identical genotypes, you still see this variability. So the question is, where does the brain generate such variability?

Jan Witkowski: So the worms are not robots that always respond in the same way to a given stimulus.

Dr. Bargmann: No, and one of the most striking features of behavior in any animal is that it's variable, that different animals are doing different things. Some of the time they're wandering around, some of the time they seem to be engaged in purposeful action. They're tuning in and out of different behavioral programs.

Jan Witkowski: So what's a worm's response to odor? Presumably, if it's an attractive odor it tends to follow the gradient.

Dr. Bargmann: Yes, that's exactly right. The odors we're studying are attractive odors and one of the responses we look at is whether or not the animal tries to move toward the odor. So, if you take the odor away, for example, it will change its direction. It will look for the odor that is missing, but not always. Sometimes you take the odor away and it doesn't seem to respond at all. It just keeps going on as though it didn't notice. And so, the first question we asked is, well, did it notice? Did its brain even know that the odor was taken away? And since the worm is transparent, we can use fluorescent markers of neural activity in a live animal, without perturbing it at all, to look at the activity of different neurons within the animal's brain. And the answer to the first question is, the worm knows perfectly well that the odor used to be there and now it's not there anymore, because we can see the olfactory neurons generating very strong signals when the odor appears or was removed, and we see the olfactory neurons responding whether or not the animal generates a behavioral response. So, the variability is not because of the ability to detect. The variability is really a choice inside the worm's brain about whether or not to respond to what has just been detected.

Copyright © 2014 Cold Spring Harbor Laboratory Press; all rights reserved; doi: 10.1101/sqb.2014.79.04

Jan Witkowski: With only 302 neurons, what is the decision process that gives rise to this variable response?

Dr. Bargmann: Since you only have 302 neurons, you have to be able to use the property of those neurons to explain the variability of the behavior. And so we've been working inward, on the one hand, from the olfactory neurons, to look at their activity and the activity of their targets. And on the other hand, we've been working back from the motor neurons that generate the changes in direction, the changes in behavior, and asking: Since these two can be far apart, where do we start to see those two signals separating from each other?

What our results suggest is that whether the animal responds or not depends on what is going on in its brain, its own endogenous brain activity when the signal arrives—the internal patterns of brain activity the animal moves through spontaneously, just as part of its locomotor pattern. If signals arrive at certain times, then they're detected. If signals arrive at other times, they're ignored. And we all know that when your mother's shouting at you, sometimes you respond and sometimes you don't. And essentially that is what the nervous system of the worm is doing. At certain times it's available or accessible to external stimuli and at certain other times it's driven by internal stimuli and will, for just a few seconds, not incorporate external stimuli into its decision-making.

Jan Witkowski: Can you explore what those internal stimuli are?

Dr. Bargmann: There are different forms of internal states of the brain that my lab is interested in. One kind is generated by neuromodulators, molecules that represent certain kinds of motivational and emotional states. So, for example, an animal that is hungry responds very differently to a stimulus than an animal that's well fed. Those kinds of differences are to a large degree represented in the neurochemistry of different kinds of neuromodulators operating on certain synapses to make them more or less sensitive. But in the case of this very specific decision, the level at which we're looking is faster than the neuromodulatory systems, and we think it's transient network state. The animal will respond on minute one, respond on minute two, not respond on minute three, respond again on minute four. We think that those variable responses are reflecting shifts in the activity state of the internal circuits. What people have seen for many years in more complex brains, and we are now seeing in worm brains as well, is that different neurons often have collective patterns of activity rather than completely independent activity. One of the terms used to describe this, based on work from John Hopfield and David Tank in the 1980s, is an "attractor state," when groups of neurons tend to become coactive, then maintain each other's activity. So a group of neurons will seize control of the nervous system and maintain that control through their activity. These sorts of activity patterns are what we think allow animals to be responsive or nonresponsive, when temporary patterns of activity representing groups of neurons with collective activities are sensitive or insensitive to other kinds of input. Our work points in that direction. We see that for a smallish number of neurons. Work from other groups, Alipasha Vaziri and Manuel Zimmer's labs, has allowed them to look at the activity of the entire worm brain simultaneously. And they see strong evidence for large groups of neurons that become coactive and then become inactive, perhaps remaining active or in an active state for tens of seconds before flipping down into an inactive state.

Jan Witkowski: What's the functional significance of these states?

Dr. Bargmann: I think that's a question we need to explore. But one influence on us is the idea that at any given time, you want to select one action and suppress other actions. In any given circumstance, the worst thing you can do is essentially seize up and try to do three things at once. So perhaps the nervous system commits to particular kinds of actions, suppresses alternative actions, and then over some period of time will allow new actions to take place. And this concept of action selection can be useful in a variety of different circumstances.

Jan Witkowski: So, changing the subject, you've been involved with the Brain Initiative Project.

Dr. Bargmann: A year ago President Obama announced that there was going to be a grand challenge in neuroscience to understand the brain, analogous to the space program or the war on cancer, or the human genome project. When that was announced, the National Institutes of Health director Francis Collins decided that, rather than simply jump in, he wanted to have a rigorous scientific planning process. So, he asked a group of 15 external scientists in different areas of neuroscience but also engineering, clinical science, chemistry, physics, and genetics to get together and to ask, what are the important problems that we need to study in the brain? What are the best approaches to take to solving them? How would you do that? How much would it cost? I'm co-chair of that planning committee together with Bill Newsome from Stanford University, and we have spent the past year consulting with ourselves and consulting with many, many other neuroscientists and scientists in general to try to answer those questions.

The first year the NIH committed 40 million dollars to starting this new project. This is nothing to sneeze at. The second year the NIH has said that they will spend $100 million on this project, so they like the plan that we came up with in year one enough to double their investment. And we're just about to turn in a final report after a year of planning and we'll see if the NIH continues to show enthusiasm. But I think there's been quite broad support for the idea that we're at a time in neuroscience where we can study brains at a level that we've never been able to study them before. For 50 years we've been able to look at individual neurons and their activity, and for 30 years we've been able to look at whole brains by fMRI, and by imaging. Now it may be possible to look at the networks of neurons and the communicating circuits of neurons that transmit information at high speed, so that we can look at the brain both at high resolution and look at the big picture of many neurons at once. There's a sense that this kind of circuit activity is going to be very important in understanding normal brain function, and also that it may provide a new way to think about what to do for neurological, psychiatric, and brain injury disorders. So, there seems to be broad support for the idea that this is the time to tackle this problem.

Jan Witkowski: And so is it a technical project, in the sense of looking at neurons, looking at their physical connections, at the local level and the brain level, analogous to the way the human genome project was a technical setting of a foundation for future developments in human genetics?

Dr. Bargmann: Yes, I think what inspires the idea of the Brain Initiative now is truly remarkable technical advances that have occurred over the past five to ten years. The first has been the ability to record from many neurons at once, hundreds of neurons at once instead of just one. The second has been the development of optogenetics, methods for perturbing neuron activity and not just watching the activity, but actually linking it to causes. And the third has been remarkable advances in computing technology that let you take very large data sets and make sense of them, and also give you theoretical and modeling capability. But, all of those things are still pretty far from what they would need to be to really address how a circuit works. When you're talking about the circuits that are involved in memory formation, or perception, it's not a hundred neurons you might be thinking about, it might be a million neurons. The first stages of thinking about the Brain Initiative are how to make these methods more powerful, more scalable, less expensive, disseminating them broadly, making them available to a large number of people—developing the technology that would make these large-scale views of brain activity much more powerful. For the first few years, that will be the main investment, and then over time, it will shift to using those tools to asking questions, with some of both at both times.

Jan Witkowski: How long will the project last?

Dr. Bargmann: We were initially told as a working group to think about the first year, then after we handed in the report we were told to keep thinking, and I think that there is a general sense in scientific planning that to really accomplish something in science takes more than 1 or 2 years. Real scientific progress is something that tends to take at least 5 years to get going because you need to set things up, you need to be able to take wrong turns and then correct them. And there's also a sense that if you're thinking out more than about 10 years, you're just going to be hopelessly out of date by the time it's done. And so, I think in real terms people have tended to try to think 10 years ahead, where you think in a lot of detail about what's going to happen in the next couple of years, you think in moderate detail about what might be possible in 5 years, and then you just give a very rough sense of what might be possible at the end of that.

A Conversation with James DiCarlo

INTERVIEWER: GARY STIX

Senior Editor, Scientific American

James DiCarlo is Head of the Department of Brain and Cognitive Sciences at the Massachusetts Institute of Technology McGovern Institute for Brain Research.

Gary Stix: Your research is focused on understanding the neuronal representations and computational algorithms that underlie visual object recognition in primates. You've asked some very basic questions recently, such as what is object recognition and why is it challenging?

Dr. DiCarlo: We all have an intuitive feel for what object recognition is. It's the ability to discriminate your face from other faces, this car from other cars, a dog from a camel. But making progress in understanding how our brains are able to accomplish this task is a very challenging problem.

Part of the reason it's challenging is defining what it is; and it's a challenging problem because we're using our brain machinery to solve the problem. It seems effortless to us. When I try to explain, for instance, to my mom, "Oh, I'm working on object recognition," she will reply "Well, I just see and I solve the problem," but of course, she is using the machinery of her brain that we are trying to reverse engineer and to decipher.

What makes the problem challenging is that each object presents an essentially infinite number of images to your retina, so you essentially never see the same object, the exact same image twice. The ability of the brain to deal with all those different images and still know that they're coming from the same object in the environment is a challenge our brain solved, we think, through evolution. But machines are still struggling to understand how to do it.

Gary Stix: This is one of the really big problems both in neuroscience and in machine learning.

Dr. DiCarlo: It's a really great problem because it's at the intersection of neuroscience and machine learning, and of psychology or cognitive science, because objects in the world are what we use to build higher cognition, things like memory or making decisions. Should I reach for this? Should I avoid it? Our brains can't do higher cognition without these foundational elements that we often take for granted unless you're a computer scientist and realize how challenging that problem is.

Gary Stix: Your research is focused on finding some of those foundational elements. There are areas of the brain, such as the inferior temporal cortex, where there might be some clues.

Dr. DiCarlo: It's been known for several decades that there's a portion of the brain, the temporal lobe, that when lost or damaged in humans and nonhuman primates leads to deficits of recognition. So we had clues that that's where the algorithms live. But just saying that that part of your brain solves the problem is not much more specific than saying your brain solves the problem. It's still a very large piece of tissue. But anatomy's told us more about the various visual areas that exist in the brain, a whole network of areas that exist there.

The tools of neurophysiology, and now more advanced tools, allow us to look more closely at neural activity, especially in nonhuman primates. Then we can begin to decipher the actual computations to the level, for example, that an engineer might need to be able to emulate what's going on in our head, or we might need if we wanted to replace those circuits or augment those circuits—to really understand the function at a detailed level, not in just the gross "this part of the brain" sense.

Gary Stix: How are you trying to drill down and find one of these foundational elements?

Dr. DiCarlo: The foundation of any science is the ability to have predictive models of a phenomenon. So for object recognition, as an engineer, if you want to emulate that, you first need to define what are we trying to predict. What would success look like? There are various levels that we might set there, but we've set one, a goal for ourselves, the ability to emulate what we do in core-object recognition.

Put simply, what that means is that I'm going to show you an image for about 200 milliseconds, which is about the time that your eyes dwell as they explore a scene. It's a time that's driven from the biology of how you explore the world with your eyes. We can do a lot with that short time window. We can easily recognize one or more

Copyright © 2014 Cold Spring Harbor Laboratory Press; all rights reserved; doi: 10.1101/sqb.2014.79.09

objects within that one-fifth of a second glimpse. It's not all of vision, but it's a defined space where we can start to get some traction on the problem.

Gary Stix: So you've got a predictive model, and then you want to test that model.

Dr. DiCarlo: In core recognition, images come in and are processed by the eyes, and then through a series of visual areas, processed somehow in ways that are murky, but we can record the neural activity along the processing pathway. Others have done that before us, and now we're doing it at a much larger scale. We can record neural activity, and we're especially interested in a place in the brain called the inferotemporal (IT) cortex, at the end of the temporal lobe, the highest level of this processing chain.

We've found that the population patterns of neural activity there, the firing patterns of all the neurons in that part of the brain, with a very simple model, an algorithm if you will, can predict very accurately the animal's perception and our own perception—our ability to do recognition in that core domain. So we have a model from neural activity in IT cortex to behavioral report.

Gary Stix: So you could predict, say, that a person or a monkey is looking at a tree in the background from examining that neural activity?

Dr. DiCarlo: That's exactly what I mean. The granularity with which we can do that is still being studied but I can tell if you're looking at a tree versus a dog, or I can tell you you're looking at a tree versus a car. I can tell if you're looking at one tree versus another.

We're now trying to see if we can do this on a trial-by-trial basis rather than on an average basis. Already the models are quite good for the tasks that we've tested. Our next step is to look even harder and build better models.

Gary Stix: The challenge with object recognition is that if I'm looking at a tree and then move slightly to the left or the right, the tree changes or I start to see another tree. Will your model still recognize that that's a tree or that it's the same tree?

Dr. DiCarlo: That's the largest thing the model has to deal with. When I say the model deals with that, the neurons up to the IT cortex when we record have dealt with that, so that part of the problem from an engineering perspective is solved.

So once I build a decoder on the IT cortex, a reading of the IT activity, if you will, then a new image of the tree will be properly decoded as a tree. It's a brand-new image, but the model will still make a prediction of what you will say, and the model will be quite accurate.

Gary Stix: What are some of the implications of this for machine learning and perhaps, one day, even for understanding problems that people have with disruptions to neural circuitry?

Dr. DiCarlo: From the machine-learning point of view, this neural activity in IT is something that machine-learning folks would call features. Those are features computed on the image. They're a very powerful set of features for the reason I described. What many people would love to do is to be able to have algorithms that produce those features. Much of machine learning is devoted to finding good features.

The brain, through evolution, has already found some good features, and that's essentially what we're reporting. Here are some nice features. Here's where they are. Here's our evidence that they're nice features. Now the machine-learning community is working, and we're working alongside of them, to help build what are called encoding algorithms that produce those features.

There's a lot of exciting progress in the field, even in the last few years, driven by what are essentially brain-inspired models that are actually now some of the state-of-the-art computer vision algorithms. They were inspired by work of the type done by those who went before us, and I hope from some of our work, to build those kind of models. They turned out to be very powerful in the computer-vision and machine-learning community. Those are the most active and exciting models in computer vision.

Gary Stix: The grand vision of what you're doing is the ability to model this all the way from encoding to neural activation and then to the decoding and perception in the brain?

Dr. DiCarlo: That's exactly the grand vision. If we can do all of that, then we would have a complete end-to-end understanding of this domain of behavior.

Gary Stix: How long until you get there?

Dr. DiCarlo: Within basic-level core recognition, it depends on your level of detail, but in the next 10 years we will have a very good understanding of core basic-level object recognition to the degree that many engineers will be satisfied. We won't know it down to the synapse, but we will know it in the way that the algorithms are very predictive of the neural activity at various levels of the system.

Gary Stix: Do you think that this could provide some insight into what goes wrong when the circuitry in the visual cortex or the temporal cortex goes awry?

Dr. DiCarlo: The most common deficit that affects recognition is major damage to that part of the brain through stroke or lesion. You've taken out those neurons. Maybe our studies could lead to ways that you could sidestep the damages or replace the function.

There are other deficits to the temporal lobe where people have things like deficits in the ability to discriminate among faces or other types of objects. They're not very common, but this kind of work should bear on those deficits, as well. We hope it will also bear on things like how kids learn to read.

At the end of the day, whenever you're doing visual tasks, you're leaning on these kind of representations to do much of your vision. I think this work will help us understand these higher-level issues of, say, social cognition or things like dyslexia. Things that depend on these foundational representations will be better understood because of our understanding of these neural circuits.

Gary Stix: Despite what people see in movies, robots in the real world are still very limited in what they can do. One of their big limitations is the ability to recognize and process information they perceive. Do you think this work could help with that?

Dr. DiCarlo: Certainly. The computer-vision community is already using brain-like algorithms. The next frontier though is expanding the domain of task, not just what you can do in 200 milliseconds, but as you explore a scene with many eye movements or navigate a scene, accumulating information over time. There will be more feedback in the system.

Those are things that we've not yet begun to touch very much, but that's the next frontier as we understand this more core domain, this foundational domain. I won't say that if we do this work we will have robots doing everything you see Data on *Star Trek* doing, but it will be a foundation to enable those next steps.

Gary Stix: Is a larger conceptual framework necessary to create kind of a movie of this whole process, to go beyond 200 milliseconds and create something that can process the full variety of our visual environment?

Dr. DiCarlo: I think what you're asking is what goal we are seeking. Sometimes to know that we've met it, we have to just emulate the system. For instance, if we could build a robot that behaved as well as us and fooled us like Data on *Star Trek*, I think we would all declare that as evidence of success.

This will certainly require, at some point, really embodied systems to test. Vision is not an isolated sense. It has to interact with the motor system, so as these systems get better, you're going to see them on moving things like robots and other autonomous devices. This is already happening. Then the work of my lab and others will have to move beyond vision, and integrate with motor action and other senses.

Gary Stix: The complexity of that seems overwhelming. Don't you think it's going to take many generations to get to these ultimate goals?

Dr. DiCarlo: In that goal, probably. Again, it's hard to make predictions with confidence. In this ventral visual stream that ends in IT cortex is a series of cortical areas where each local bit of cortex looks anatomically very, very similar. One of the exciting ideas that's been in our field for a long time is that there's some interesting local learning algorithm that's of smaller scale in both size and computational power than the global algorithm I've been talking about.

One of the things many of us hope for is that if we could get some insight from this kind of work into that more local algorithm, that could generalize more broadly to things that we're not even thinking about, that might apply to audition or other senses, maybe to some of these higher-level cognitive things, but that's hard to see right now. So there's a hope for a shortcut, if you will, that arises from working on a clearly defined problem.

Gary Stix: You mentioned in a recent paper that there are 40,000 neuron building blocks. Is that what you mean?

Dr. DiCarlo: That's exactly what I'm referring to: the idea that you might have sub-modules of something of that scale, say, a millimeter of cortical tissue. Again, this is an idea that's been around for a while because each piece of cortex looks very similar. What is it doing? What's that fundamental operation?

We are working on it from the problem of vision, but we hope we will get some understanding of that because we set a clear goal of what success looks like. We're building models to emulate what the neurons do at all levels, and so we might start to glean some insight that could generalize outside of vision.

Gary Stix: So these mini modules might be present in other parts of the cortex that are processing incoming sensory information?

Dr. DiCarlo: That's right. They're known to be present. This is a hypothesis that's been in the field for a long time. Like the dream of discovering the DNA code, discovering the cortical processing algorithm, is a long-standing dream.

It will be a big jump and certainly not the whole story. Bits of cortex have to talk over long distances, so everything's not going to be defined as a local small module, but we might get a very big boost from understanding that kind of module.

Gary Stix: When you say an algorithm, is that made up of a kind of code that would be present, say, in the inferior temporal area?

Dr. DiCarlo: People use words like algorithm and models in different ways. If we talk about the algorithm of the ventral stream, that is maybe executed by many subalgorithms, that consist of these building blocks, but when put together, give rise to the global algorithm of the ventral stream. The hope is that the global algorithm of the ventral stream built by subalgorithms that are simpler and more primitive and more understandable and that might generalize well. That subalgorithm in a piece of cortex might take some set of inputs, some learning, local processing, and give some set of outputs. It's essentially a little box that does something clever, and when stacked horizontally across the visual field and then stacked vertically into a deep stack, visual area V1, V2, a series of areas along the temporal lobe, might give rise to the global temporal-lobe or ventral-stream algorithm that I was referring to earlier.

Algorithm is a word computer scientists like, but as neuroscientists, what we're imagining is not just a set of lines or codes that says, "If this, do that," in a way that those of us who write code are used to. It's also going to be a set of neurons that execute that function.

We're building algorithms that are actually already constrained somewhat by the biology. We're not just writing code and hoping that it predicts the responses. We're building systems that are modeled on the biology so that they will translate both to give constraint and to what this neuron does within the algorithm. That's a different level of how deeply you want to look at the system, how you want to describe it.

What's exciting about neuroscience and also challenging is that you have to work at all these levels—what's the goal of the system, how do we describe the algorithm, what are the mechanisms of the neurons and the synapses? We'd like to say we understand all those levels of the system.

Gary Stix: There has been work in the last few years on retinal prosthetics. One approach has been to implant codes into prosthetics so that they can process incoming photons in the same way the retina does. Does the kind of work you're doing connect in some ways with this?

Dr. DiCarlo: This is one of the things we're most excited about right now. There are visual prosthetics for people who, say, have lost a retina. The dominant approach is to try to bypass the retina and reinject a spatial pattern of activity, say, in an early visual area or the subcortical area that comes right behind the retina, called the LGN.

That makes sense from an engineering point of view. The downside is that if you want to get an image in, it's a very high-dimensional space, with many, many pixels that you'd want to play to make it feel like your normal vision would feel like.

We're working at the highest level where your brain has already reduced the dimensionality from millions of pixels to something that's more abstract. If we could have a hundred ways to push the neurons around, we don't know what it would feel like yet, but it might be very functional. It could be a better way to think about brain–machine interfaces, that you only need to have a hundred ways to inject signal, a hundred channels rather than millions, to make a rich perceptual space.

We're now testing what moving neurons around in a monkey does to that monkey's percept right now. This is a long way from a human, but as we start to understand what's happening in an animal, our work might provide a shortcut or a better prosthetic.

Gary Stix: To sum up, what your work is doing is taking something very basic that all of us can relate to and try to find the physical basis in our brains of how we see, how we recognize an object, and come to a fundamental physical and theoretical understanding of that really huge challenge.

Dr. DiCarlo: That's really been the goal of neuroscience since its formation. We believe the brain is a set of mechanisms that give rise to this amazing phenomenology that we feel. We are studying one example of that phenomenology but it's one that many of us can relate to and it would be a foundational success if we can create that end-to-end understanding—a large brick in the foundation of building towards understanding cognition.