QnAs with Terrence J. Sejnowski

For most of his career, Terrence Sejnowski, a professor of computational neuroscience at the Salk Institute for Biological Studies and a Howard Hughes Medical Institute investigator, has peered at the brain with pin-sharp precision. By using simulations to make sense of experimental data, Sejnowski has helped link biophysical processes in the brain to human behavior. His research has revealed insights into a raft of phenomena from vision to sleep to brain disorders. They could lead to practical benefits: Bestriding the fields of computational biology, neuroscience, psychology, and education, Sejnowski and other researchers hope to usher the age of machine learning into the real world. Sejnowski tells PNAS how using machines to model and emulate human behavior could make a difference in our lives.

PNAS: How did you become interested in machine learning?

Sejnowski: One of the most challenging questions in neuroscience is how social behaviors emerge from brain processes underlying sensation, emotions, language, memory, and cognition. When we first set out to address this challenge, it occurred to us that one way by which physicists figured out phenomena like gravity and aerodynamics was by building devices that exploited those phenomena. So, we needed to build machines that work like the brain by using software and computer chips that would form circuits capable of interacting with humans through social signals. In collaboration with Paul Ekman, an expert on reading facial expressions, our goal was to make machines capable of interpreting facial expressions so that, someday, social robots could communicate with humans on their own terms.

PNAS: And where would we use these social robots?

Sejnowski: Javier Movellan, a computational neuroscientist at the University of California, San Diego's Institute for Neural Computation, has built a social robot he calls Rubi that interacts with toddlers who are just beginning to learn language. One of the challenges for preschool teachers is classroom control; the kids are running all over the place, so it's difficult for the lone teacher to help kids focus. Rubi engaged the kids, encouraged dialogue, and facilitated learning. So, the idea is to use robots as teaching assistants. But it's still early days.

PNAS: How do you make robots emulate human social learning?

Sejnowski: The first step is to get the child to accept the robot as a learning partner rather than as a toy. By using mathematical theory and demonstration, Javier showed that the

Terrence J. Sejnowski.

most crucial variable for interacting with humans is response time. If a robot does not respond to a child's question within a certain time window, the child loses interest. Also, a child will look at an object to which a teacher is pointing, so robots should be capable of shared attention, another hallmark of human learning. Robots must also be capable of other important features of human learning, such as empathy and imitation, which come from recognizing human emotions. But again, it's early days.

PNAS: All this smacks of artificial intelligence.

Sejnowski: This is very different from traditional approaches in artificial intelligence, where the goal is to create a cognitive machine that creates a model of the world and computes responses based on that model. That's not how the brain generates behavior. With its limited capacity, the brain selects only the most important sensory inputs to process and the most effective responses to store. Thanks to its capacity for learning and memory, the brain is able to interact in a social way with relatively low bandwidth, which is partly what makes social robots feasible. By emulating biological intelligence, machine learning is heralding a new era.

PNAS: To many, a robot in the classroom is the stuff of science fiction. How do you convince policymakers that the investment is worth the payoff?

Sejnowski: First of all, the cat's already out of the bag. It's now a question of optimizing the technology for our own benefit. For example, social robots can serve as personal cognitive enhancers. Second, the idea would not be to replace teachers but to provide them with assistants. Besides helping teachers to hold toddlers' attention in the classroom, social robots can stand in when teachers

need to be briefly absent. Robots can help relieve teachers of some of their mundane duties so that teachers can serve as role models and tailor attention to individual students. That said, we can't predict the full impact of these transformative technologies.

PNAS: Fair enough. So where's the rub?

Sejnowski: It's mainly in the resources. We've made sufficient progress in neuroscience and engineering to be able to overcome technical challenges to using machines in social contexts. But we need to scale up lab experiments, clearly calling for a major investment of resources. If we had a thousand Rubis, we could accelerate research and reduce costs. The other problem is societal. Will our institutions be able to adapt to the new environments that such endeavors will help create? That's an open question.

PNAS: How will the new environment help children improve their cognitive skills?

Sejnowski: There's a lot of emphasis on classroom learning of subjects like language, mathematics, and science, but to improve learning, we also need an emphasis on acquiring basic cognitive skills like attention, listening, and memory. We have evidence that social robots can help improve attention. Paula Tallal, codirector of the Temporal Dynamics of Learning Center (TDLC) in San Diego, has developed software already being used in classrooms across the country that can help children who have difficulties listening and hence, understanding language. Hal Pashler, also at TDLC, has studied a well-known phenomenon in memory research—the spacing effect—to find the optimal intervals for refreshing memory to help children retain learned material for many years. These are just a couple of examples of wide-ranging research in neuroeducation, a field dedicated to helping children become better learners.

PNAS: Your own work in the mid-1990s shed surprising light on reinforcement learning.

Sejnowski: We developed a computational model of the brain's dopamine system, involved in reward-based learning, to understand how the dopamine neurons learn to make predictions about future rewards. This computational model has been confirmed in a wide range of settings using brain imaging in humans. As they learn new facts about the world, children use the dopamine system as a guide to finding the best sequence of steps to solve problems and to reach a goal. We are just beginning to understand how the different learning systems in the brain work together to produce the astonishing range of behaviors humans are capable of.

—Prashant Nair, *Science Writer*

www.manaraa.co