

Grow the Program?

Simply making the program much bigger could help. The feed-forward architecture has roots in the 1950s, when MIT's Marvin Minsky first simulated cortical function by hooking together simulated neurons to form neural nets.

In those days, the limited speed and memory of computers could handle only a small number of neurons and neural connections. Consequently, the neural nets were applied only to very simple problems. The performance of these neural nets was not good, or the problems they solved were trivial. But the proponents of neural nets have claimed ever since that scaling-up the size of the system by adding more neurons to include more feature detectors and more connections would help the simulations learn more about the world and thereby improve their performance to the point that it might eventually rival that of biological cortical material.

"With Roadrunner, we can actually test this hypothesis for the first time," says Bettencourt. Another member of the Synthetic Visual Cognition Team, Steven Brumby, ran PetANNet on a standard workstation and found that it took about 38 seconds to process a black-and-white image of 320 × 240 pixels. If the model's parameters were scaled up to human values—for example by increasing the number of feature detectors—it would take about a day to process a color image of a million pixels, which means that such a simulation could process only 300 or 400 scenes per year! And even if scaling-up significantly improved object-identification accuracy, the software would be much too slow to be useful.

Roadrunner, however, is fast enough to simulate the operation of the entire visual cortex in real time. There are about 10 billion neurons in the human visual cortex, and each neuron is connected to about 10,000 others. Each neuron also fires about 10 times per second, which for a computer means about 10 "floating-point operations" per second (called "flops"). Multiplied together, these numbers give a quadrillion flops per second, or one "petaflop" per second. The speed record the team set with Roadrunner last summer was 1.14 petaflop per second.

So, Roadrunner has what it takes to prove whether scaling-up a feed-forward neural net will improve the software's accuracy to human levels. If scaling-up is the answer, Roadrunner will also be able to identify objects as quickly as humans do.

However, the main research challenge in simulating a system as complex as the visual cortex is teaching the simulation about the visual world. Scaling up means that the representations of the visual world, especially in the upper layers of the visual cortex, can be more numerous and more precise. However, these representations are constructed only when the simulation actually observes the visual world. So, to fully realize the potential for creating more representations that have greater precision, the simulation must also be exposed to the visual world as widely as possible. Thus, the "training set" of visual images used to develop those representations must be as large and diverse as possible.

We also note that humans take months to start seeing well and years to understand what they see. Roadrunner will be able to test new ideas of how the human brain learns about the visual world and how it organizes itself, by making neural connections, to recognize features and to abstract meaning from what it sees.