

# Study suggests computational role for neurons that prevent other neurons from firing

9 January 2017, by Larry Hardesty



"There's a close correspondence between what you need for communication in rapidly changing networks and information processing in the brain," professor Nancy Lynch says. "We're trying to find problems that can benefit from this distributed-computing perspective, focusing on algorithms for which we can prove mathematical properties." Credit: MIT News

Researchers at MIT's Computer Science and Artificial Intelligence Laboratory have developed a new computational model of a neural circuit in the brain, which could shed light on the biological role of inhibitory neurons—neurons that keep other neurons from firing.

The model describes a neural circuit consisting of an array of input neurons and an equivalent number of output neurons. The circuit performs what neuroscientists call a "winner-take-all" operation, in which signals from multiple input neurons induce a signal in just one output neuron.

Using the tools of theoretical computer science, the researchers prove that, within the context of their model, a certain configuration of inhibitory neurons

provides the most efficient means of enacting a winner-take-all operation. Because the model makes empirical predictions about the behavior of inhibitory neurons in the brain, it offers a good example of the way in which computational analysis could aid neuroscience.

The researchers will present their results this week at the conference on Innovations in Theoretical Computer Science. Nancy Lynch, the NEC Professor of Software Science and Engineering at MIT, is the senior author on the paper. She's joined by Merav Parter, a postdoc in her group, and Cameron Musco, an MIT graduate student in electrical engineering and computer science.

For years, Lynch's group has studied communication and resource allocation in ad hoc networks—networks whose members are continually leaving and rejoining. But recently, the team has begun using the tools of network analysis to investigate biological phenomena.

"There's a close correspondence between the behavior of networks of computers or other devices like mobile phones and that of biological systems," Lynch says. "We're trying to find problems that can benefit from this distributed-computing perspective, focusing on algorithms for which we can prove mathematical properties."

## Artificial neurology

In recent years, artificial neural networks—computer models roughly based on the structure of the brain—have been responsible for some of the most rapid improvement in artificial-intelligence systems, from speech transcription to face recognition software.

An artificial neural network consists of "nodes" that,

like [individual neurons](#), have limited information-processing power but are densely interconnected. Data are fed into the first layer of nodes. If the data received by a given node meet some threshold criterion—for instance, if it exceeds a particular value—the node "fires," or sends signals along all of its outgoing connections.

Each of those outgoing connections, however, has an associated "weight," which can augment or diminish a signal. Each node in the next layer of the network receives weighted signals from multiple nodes in the first layer; it adds them together, and again, if their sum exceeds some threshold, it fires. Its outgoing signals pass to the next layer, and so on.

In artificial-intelligence applications, a neural network is "trained" on sample data, constantly adjusting its weights and firing thresholds until the output of its final layer consistently represents the solution to some computational problem.

### Biological plausibility

Lynch, Parter, and Musco made several modifications to this design to make it more biologically plausible. The first was the addition of inhibitory "neurons." In a standard artificial neural network, the values of the weights on the connections are usually positive or capable of being either positive or negative. But in the brain, some neurons appear to play a purely inhibitory role, preventing other neurons from firing. The MIT researchers modeled those neurons as nodes whose connections have only negative weights.

Many artificial-intelligence applications also use "feed-forward" networks, in which signals pass through the network in only one direction, from the first layer, which receives input data, to the last layer, which provides the result of a computation. But connections in the brain are much more complex. Lynch, Parter, and Musco's circuit thus includes feedback: Signals from the output neurons pass to the inhibitory neurons, whose output in turn passes back to the output neurons. The signaling of the output neurons also feeds back on itself, which proves essential to enacting the winner-take-all strategy.

Finally, the MIT researchers' network is probabilistic. In a typical artificial neural net, if a node's input values exceed some threshold, the node fires. But in the brain, increasing the strength of the signal traveling over an input neuron only increases the chances that an output neuron will fire. The same is true of the nodes in the researchers' model. Again, this modification is crucial to enacting the winner-take-all strategy.

In the researchers' model, the number of input and output neurons is fixed, and the execution of the winner-take-all computation is purely the work of a bank of auxiliary neurons. "We are trying to see the trade-off between the computational time to solve a given problem and the number of auxiliary neurons," Parter explains. "We consider neurons to be a resource; we don't want to spend much of it."

### Inhibition's virtues

Parter and her colleagues were able to show that with only one inhibitory neuron, it's impossible, in the context of their model, to enact the winner-take-all strategy. But two inhibitory neurons are sufficient. The trick is that one of the inhibitory neurons—which the researchers call a convergence neuron—sends a strong inhibitory signal if more than one output neuron is firing. The other inhibitory neuron—the stability neuron—sends a much weaker signal as long as any output neurons are firing.

The convergence neuron drives the circuit to select a single output neuron, at which point it stops firing; the stability neuron prevents a second output neuron from becoming active once the convergence neuron has been turned off. The self-feedback circuits from the output neurons enhance this effect. The longer an output neuron has been turned off, the more likely it is to remain off; the longer it's been on, the more likely it is to remain on. Once a single output neuron has been selected, its self-feedback circuit ensures that it can overcome the inhibition of the stability neuron.

Without randomness, however, the circuit won't converge to a single output neuron: Any setting of the inhibitory neurons' weights will affect all the [output neurons](#) equally. "You need randomness to break the symmetry," Parter explains.

The researchers were able to determine the minimum number of auxiliary neurons required to guarantee a particular convergence speed and the maximum convergence speed possible given a particular number of auxiliary neurons.

Adding more convergence neurons increases the convergence speed, but only up to a point. For instance, with 100 input neurons, two or three convergence neurons are all you need; adding a fourth doesn't improve efficiency. And just one stability neuron is already optimal.

But perhaps more intriguingly, the researchers showed that including excitatory neurons—neurons that stimulate, rather than inhibit, other neurons' firing—as well as inhibitory neurons among the auxiliary neurons cannot improve the efficiency of the circuit. Similarly, any arrangement of inhibitory neurons that doesn't observe the distinction between convergence and stability neurons will be less efficient than one that does.

Assuming, then, that evolution tends to find efficient solutions to engineering problems, the model suggests both an answer to the question of why inhibitory neurons are found in the brain and a tantalizing question for empirical research: Do real inhibitory neurons exhibit the same division between convergence neurons and stability neurons?

"This computation of winner-take-all is quite a broad and useful motif that we see throughout the brain," says Saket Navlakha, an assistant professor in the Integrative Biology Laboratory at the Salk Institute for Biological Studies. "In many sensory systems—for example, the olfactory system—it's used to generate sparse codes."

"There are many classes of inhibitory neurons that we've discovered, and a natural next step would be to see if some of these classes map on to the ones predicted in this study," he adds.

"There's a lot of work in neuroscience on computational models that take into account much more detail about not just inhibitory neurons but what proteins drive these neurons and so on," says Ziv Bar-Joseph, a professor of computer science at

Carnegie Mellon University. "Nancy is taking a global view of the network rather than looking at the specific details. In return she gets the ability to look at some larger-picture aspects. How many [inhibitory neurons](#) do you really need? Why do we have so few compared to the excitatory [neurons](#)? The unique aspect here is that this global-scale modeling gives you a much higher-level type of prediction."

**More information:** Computational Tradeoffs in Biological Neural Networks: Self-Stabilizing Winner-Take-All Networks: [arxiv.org/pdf/1610.02084v1.pdf](https://arxiv.org/pdf/1610.02084v1.pdf)

*This story is republished courtesy of MIT News ([web.mit.edu/newsoffice/](http://web.mit.edu/newsoffice/)), a popular site that covers news about MIT research, innovation and teaching.*

Provided by Massachusetts Institute of Technology

APA citation: Study suggests computational role for neurons that prevent other neurons from firing (2017, January 9) retrieved 22 September 2021 from <https://techxplore.com/news/2017-01-role-neurons.html>

*This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.*