


BEHAVIOR & SOCIETY

Pushy AI Bots Nudge Humans to Change Behavior

Researchers use artificially intelligent bot programs to stimulate collaboration and make people more effective

By Richard Conniff on May 17, 2017



Credit: Sarah Holmlund Getty Images

When people work together on a project, they often come to think they've figured out the problems in their own respective spheres. If trouble persists, it's somebody else—engineering, say, or the marketing department—who is screwing up. That “local focus” means finding the best way forward for the overall project is often a struggle. But what if adding artificial intelligence to the conversation, in the form of a computer program called a bot, could actually make people in groups more productive?

This is the tantalizing implication of a study published Wednesday in *Nature*. Hirokazu Shirado and Nicholas Christakis, researchers at Yale University's Institute for Network Science, were wondering what would happen if they looked at artificial intelligence (AI) not in the usual way—as a potential replacement for people—but instead as a useful companion and helper, particularly for altering human social behavior in groups.

First the researchers asked paid volunteers arranged in online networks, each occupying one of 20 connected positions, or “nodes,” to solve a simple problem: Choose one of three colors (green, orange or purple) with the individual, or “local,” goal of having a different color from immediate neighbors, and the “collective” goal of ensuring that every node in the network was a different color from all of its neighbors. Subjects' pay improved if they solved the problem quickly. Two thirds of the groups reached a solution in the allotted five minutes and the average time to a solution was just under four minutes. But a third of the groups were still stymied at the deadline.

The researchers then put a “bot”—basically a computer program that

can execute simple commands—in three of the 20 nodes in each network. When the bots were programmed to act like humans and focused logically on resolving conflicts with their immediate neighbors, they didn't make much difference. But when the researchers gave the bots just enough AI to behave in a slightly “noisy” fashion, randomly choosing a color regardless of neighboring choices, the groups they were in solved the problem 85 percent of the time—and in 1.7 minutes on average, 55.6 percent faster than humans alone.

Being just noisy enough—making random color choices about 10 percent of the time—made all the difference, the study suggests. When a bot got much noisier than that, the benefit soon vanished. A bot's influence also varied depending on whether it was positioned at the center of a network with lots of neighbors or on the periphery.

So why would making what looks like the wrong choice—in other words, a mistake—improve a group's performance? The immediate result, predictably, was short-term conflict, with the bot's neighbors in effect muttering, “Why are you suddenly disagreeing with me?” But that conflict served “to nudge neighboring humans to change their behavior in ways that appear to have further facilitated a global solution,” the co-authors wrote. The humans began to play the game differently.

Errors, it seems, do not entirely deserve their bad reputation. “There are many, many natural processes where noise is paradoxically beneficial,” Christakis says. “The best example is mutation. If you had a species in which every individual was perfectly adapted to its environment, then when the environment changed, it would die.”

Instead, random mutations can help a species sidestep extinction.

“We’re beginning to find that error—and noisy individuals that we would previously assume add nothing—actually improve collective decision-making,” says Iain Couzin, who studies group behavior in humans and other species at the Max Planck Institute for Ornithology and was not involved in the new work. He praises the “deliberately simplified model” used in the *Nature* study for enabling the co-authors to study group decision-making “in great detail, because they have control over the connectivity.” The resulting ability to minutely track “how humans and algorithms collectively make decisions,” Couzin says, is “really going to be the future of quantitative social science.”

But how realistic is it to think human groups will want to collaborate with algorithms or bots—especially slightly noisy ones—in making decisions? Shirado and Christakis informed some of their test groups that they would be partnering with bots. Perhaps surprisingly, it made no difference. The attitude was, “I don't care that you're a bot if you're helping me do my job,” Christakis says. Many people are already accustomed to talking with a computer when they call an airline or a bank, he adds, and “the machine often does a pretty good job.” Such collaborations are almost certain to become more common amid the increasing integration of the internet with physical devices, from automobiles to coffee makers.

Real-world, bot-assisted company meetings might not be too far behind. Business conferences already tout blended digital and in-person events, featuring what one conference planner describes as “integrated online and offline catalysts” that use virtual reality,

augmented reality and artificial intelligence. Shirado and Christakis suggest slightly noisy bots are also likely to turn up in crowdsourcing applications—for instance, to speed up citizen science assessment of archaeological or astronomical images. They say such bots could also be useful in social media—to discourage racist remarks, for example.

But last year when Microsoft introduced a twitter bot with simple AI, other users quickly turned it into epithet—spouting bigot. And the opposite concern is that mixing humans and machines to improve group decision-making could enable businesses—or bots—to manipulate people. “I’ve thought a lot about this,” Christakis says. “You can invent a gun to hunt for food or to kill people. You can develop nuclear energy to generate electric power or make the atomic bomb. All scientific advances have this Janus-like potential for evil or good.”

The important thing is to understand the behavior involved, “so we can use it to good ends and also be aware of the potential for manipulation,” Couzin says. “Hopefully this new research will encourage other researchers to pick up on this idea and apply it to their own scenarios. I don’t think it can be just thrown out there and used willy-nilly.”



ADVERTISEMENT | REPORT AD

ABOUT THE AUTHOR(S)



Richard Conniff

Richard Conniff is an award-winning science writer for magazines and a contributing opinion writer for the New York Times. His books include *House of Lost Worlds* (Yale University Press, 2016) and *The Species Seekers* (W.W. Norton, 2010).

Credit: Nick Higgins

Scientific American is part of Springer Nature, which owns or has commercial relations with thousands of scientific publications (many of them can be found at www.springernature.com/us). Scientific American maintains a strict policy of editorial independence in reporting developments in science to our readers.

© 2017 SCIENTIFIC AMERICAN, A DIVISION OF NATURE AMERICA, INC.

ALL RIGHTS RESERVED.