


## Can a robot be too nice?

Smart machines need the right “personality” to work well—and experts are finding the best choice may not always be what we think we want

By Leon Neyfakh | GLOBE STAFF AUGUST 15, 2014

PRINT REPRINTS

ARTICLE



MINA PRICE FOR THE BOSTON GLOBE

The two robots, built by researchers in Singapore, had very different jobs. One was a nurse, designed to take a person’s blood pressure, provide basic medical advice, and offer to book an appointment with a doctor. The other was a security guard, equipped with a closed-circuit surveillance system, that alerted its human users to suspicious intruders and possible emergencies in the building.

The researchers, led by an engineer named Taezoon Park, gave them more than jobs. They gave them distinct personalities. It was an experiment: would humans react to the robots differently based on how they carried

speak loudly and quickly, use more animated hand gestures, and start conversations instead of waiting to be spoken to. The other personality was more reserved, speaking much more slowly and quietly, moving around less, and letting the user initiate communication.

What the researchers found, as they described in a recently published paper, was a striking difference between the two. When it came to the nurse robot, people preferred and trusted it more when its personality was outgoing and assertive. What people wanted in a security guard was exactly the opposite: the livelier, extraverted version clearly rubbed people the wrong way. Not only were they less confident in its abilities, and dubious that it would keep them away from danger, they simply liked it less overall.

The idea of programming a robot to have a specific personality might sound like science fiction; in a world where true artificial intelligence has yet to be achieved, a personality—an individual's distinct mixture of emotional response, attitude, and motivation—seems even more subtle and complex. But for computer scientists interested in social robotics, it has become a surprisingly immediate goal. As machines become more sophisticated, and integrated in new ways into human society, researchers have begun to realize that their effectiveness depends on how easily we relate to them.

#### Related

- [My day as a robot](#)
- [Pentametron, robot poet of Twitter](#)
- [Will a robot take your kid's job?](#)
- [The word: Robot wars](#)

“With technology that is genuinely going to live with us in an embedded way...and is going to be interacting with us in complex ways, personality is extremely important, the same way it is when you're dealing with people,” said computer scientist Peter McOwan, the coordinator of a major European research effort on social robotics that ran from 2008 until 2012.

What researchers are finding is that it's not enough for a machine to have an agreeable personality—it needs the right personality. A robot designed to serve as a motivational exercise coach, for instance, might benefit from being more intense than a teacher-robot that plays chess with kids. A museum tour guide robot might need to be less indulgent than a personal assistant robot that's supposed to help out around the house.

A growing body of research is starting to reveal what works and what doesn't. And although building truly human-like robots will probably remain technologically impossible for a long time to come, researchers say that imbuing machines with personalities we can understand doesn't require them to be “human-like” at all. To hear them describe the future is to imagine a world—one coming soon—in which we interact and even form long-term relationships with socially gifted devices that are designed to communicate with us on our terms. And what the ideal machine personalities turn out to be may expose needs and prejudices that we're not even aware we have.

\*\*\*

DESIGNERS TRYING to build character traits into machines face one immediate obstacle: People encountering artificial personalities have often hated them. The authoritative talking cars of the 1980s (“your fuel is low”) quickly became dated laughingstocks, while Clippy, the “helpful” Microsoft Office paper clip of the mid-1990s, can still be invoked as shorthand for “annoying and useless.”

Jaron Lanier has argued that making machines seem more “alive” would make people unduly deferential to their devices, and harmfully scramble their intuitions about the difference between a fellow human and an electronic device whose plug we shouldn’t find it hard to pull. In a widely circulated 1995 essay, Lanier called the issue of anthropomorphism “the abortion question of the computer world”—a debate that forced people to take sides regarding “what a person is and is not.”

Today the discussion surrounding anthropomorphism is much less heated—maybe because 20 years after Clippy, the closest thing many people have to a social robot in their lives is a Roomba. But the argument for going further with the effort is that, as technology becomes more advanced, we’ll actually need a more natural, human way to interact with our machines. If we just rely on buttons and typing, said Bernt Meerbeek, a Dutch researcher who studies personality in artificial agents: “what will happen is that they will be able to support us and do more than communicating with them will be too difficult.” In many cases, the most intuitive interface of all is going to be one that speaks to us in our own language, the way Siri, the “digital assistant” that lives inside people’s iPhones, takes voice commands and delivers replies aloud.

Many of the questions in personality design are technical, of course, but some of the most important ones are emotional: What kind of personalities do we actually want our various machines to have? The example of the security bot from Singapore suggests that in the wrong context, positive-seeming qualities can backfire. One study by Meerbeek, found that people want their automated vacuum cleaners to like routine, and to be “calm, polite, a cooperative.” Another, scheduled to begin in a month as a collaboration between Nissan, Wendy Ju of Stanford, the University of Twente’s Vanessa Evers, will investigate whether a self-driving car’s “personality”—as expressed the way it handles situations on the road—could cause safety issues if it is poorly matched to the personality of its occupant.

This kind of work sits at the intersection of computer science and social psychology, with researchers drawing as much on insights from the study of human personality and cognitive processing as engineering. Sometimes people in the humanities get involved, too, like a drama professor at Carnegie Mellon University who helped her colleagues from across campus build “storytelling robots.”

One line of thinking suggests that the artificial agents we interact with—whether they’re a voice inside our phone or an animated face on a screen asking you for ID when you walk into a building—will work best if they have personalities that correspond to our own. To investigate this, roboticists Maja Mataric and Adriana Tapus from the University of Southern California conducted a set of experiments involving robots that assisted people with physical rehabilitation exercises, like turning pages in a newspaper and putting books on a shelf. It turned out that “supportive” and “nurturing” robots, which said things like “I hope it’s not too hard” and “I’m here for you,” produced better results when dealing with introverts, while “coach-like” robots, which said things like “Move! Move!” and “You can do more than that!” in a more assertive tone, were more effective with extroverts.

Another approach to designing machine personality starts not with the user, but with the task being performed. A study by Jennifer Goetz and Sara Kiesler at Carnegie Mellon looked at robots leading a group of young, healthy subjects in a 20-minute breathing and stretching routine. Goetz and Kiesler found that people liked a “playful” robot that joked around and treated the task as fun—but it was ultimately less effective at getting people to complete than a “serious” robot that stayed focused and reminded users of the health benefits of what they were doing. “A

How we feel about an artificial personality may also depend on how much control we're going to have over the robot. Meerbeek, the Dutch researcher, conducted a study involving robots that helped people survey what was on TV and made suggestions about what to watch. A robot named Lizzy, designed to be friendly and extroverted, was compared to a robot named Catherine, who was less expressive and more formal. It turned out that when users were given less control over the robot—when it was programmed to take more initiative and talk in a more assertive way—they preferred the dynamic and chatty Lizzy; however, when the robot was programmed to ask more questions and wait for permission before doing anything, Catherine got better marks.


A similar lesson about authority emerged from a University of Arizona study involving a virtual border guard. The device, nicknamed AVATAR, was built to speed up lines at border crossings by interviewing people about their travels and analyzing their speech and body language for signs of deception or other irregularities. According to Aaron Elkins, one of the researchers involved in the project, a lab experiment determined that a smiling virtual guard was perceived as unserious and weak, while an unsmiling one projected power and authority. Elkins said his team has started experimenting with different scripts, in which the guard's demeanor gradually shifts from friendly to accusatory. One takeaway from this research: Even if artificial entities are here to serve humanity, it's possible that in some cases that means making them appear more dominant and powerful.

\*\*\*

DESIGNING ARTIFICIAL entities perfectly groomed to meet our emotional needs has an obvious appeal, like creating the exact right person for a job from thin air. But it's also not hard to imagine the problems that might arise in a world where we're constantly dealing with robots calibrated to treat us, on an interpersonal level, exactly the way we want. We might start to prefer the company of robots to that of other, less perfectly optimized humans. We might react against them, hungry for some of the normal friction of human relations. As Lanier worried, we might start to see the lines blur, and become convinced that machines—which in some ways are vastly inferior to us, and in other ways vastly superior—are actually our equals.

If it's any consolation, though, the march of science isn't exactly imposing this way of seeing the world: Our brains got there first. Even without a gifted programmer's help, we ascribe intent, motivation, and character to machines. According to one study, even a single dot could be moved around a screen in a way that made people react as if it were alive. This tendency to anthropomorphize objects has long been a driver in product design (marketers even use the phrase "product personality"). Cars can be designed to look friendly or mean; teapots can be made to look bashful or cute. We are, in short, already surrounded by subtly engineered, artificial personalities that our brains engage with naturally.

There is, however, at least one big difference between a tough-looking car and a cheerful domestic robot, and it underscores the potential of social robotics research to change human experience. As Peter McOwan points out, robot personality is portable: It is really just a set of algorithms that can be easily transferred from one device to



MINA PRICE FOR THE BOSTON GLOBE

day, “migrating” from one platform to another while maintaining a coherent and recognizable personality.

If that sounds too much like a “Her”-style dystopia, in which ultra-reliable, infinitely tolerant artificial spirits invade people’s lives so thoroughly that it becomes difficult to live without them, it’s possible that we’re shortchanging the ingenuity of human design. Our personalities, after all, are not static, nor are they consistent; we change based on where we are, who we’re with, what we’re doing, and how we’re feeling. Astrid Weiss, a postdoctoral research fellow at Vienna University of Technology, points out that robots can be programmed to do this, too—and that the people using them could, and should, maintain control over what aspects of their personality they’re expressing at any given moment. “Maybe you want your robot to treat you differently when you are home alone [and when] friends are over,” Weiss wrote in an e-mail.

That element of human control—the ability to turn the volume up and down, if you will—might be one way to keep us conscious of the fact that any artificial personality we deal with is nothing more than, as Weiss writes, its “pre-programmed social and emotional capabilities.” By making robots even more like humans, in other words—by giving them the flexibility to have lots of different personalities, and retaining the power to flip between them—we make it at least slightly more likely that we’ll hang onto the fundamental truth of how different they are from us.

**More coverage:**

[My day as a robot](#)

[Will a robot take your kid’s job?](#)

*Leon Neyfakh is the staff writer for Ideas. E-mail [leon.neyfakh@globe.com](mailto:leon.neyfakh@globe.com).*