

The Atlantic

Chatbots Have Entered the Uncanny Valley

When virtual assistants almost pass as human, they only seem more robotic.



KAVEH WADDELL | APR 21, 2017 | TECHNOLOGY

Like *The Atlantic*? Subscribe to [The Atlantic Daily](#), our free weekday email newsletter.



When a robot almost looks human—almost, but not quite—it often comes across as jarringly fake instead of familiar. Robots that are clearly artificial, like WALL-E or R2-D2, don’t have this problem. But androids [like this one](#) that imperfectly mimic human mannerisms and facial expressions are weird enough to be haunting.

This phenomenon is known as the [uncanny valley](#). It’s a major obstacle for designers who try to make their robots look like people—and it may be just as much of a hurdle for developers who are creating bots that *talk* like people, but that don’t have a body at all.

These bots are all over the place: They’re the digital assistants that talk to us from smartphones and Amazon Echoes, as well as the chatbots that are proliferating on platforms like Skype, Telegram, and Facebook Messenger.

Compared to their ancestors, these assistants are quite advanced. Remember Clippy, the [animated paperclip](#) that lived in Microsoft Word until 2004? Its prompts were more nuisance than assistance. Amazon’s Alexa, by contrast, can dim the lights, pull the shades, and put on music, all in response to voice commands.

But rather than making easy conversationalists, these bots have entered their own version of the uncanny valley. As interacting with them approaches the experience of talking with another human, their robot-ness becomes accentuated.

For both androids and chatbots, the uncanny valley problem might have something to do with how we perceive their capabilities. In [a study](#) published in *Social Cognitive and Affective Neuroscience* in 2012, researchers monitored people’s brain activity as they watched humans and robots move around in a video. They found that when participants watched a human moving fluidly and a very mechanical-looking robot moving jerkily, their brains reacted similarly. But when they watched a robot that *looked* human but didn’t *act* like one, their brain activity flared: They experienced what the researchers called a “prediction error.”

The same mechanic could be at play when humans interact with virtual assistants. When a bot is clearly a bot, the person interacting with it generally knows how limited its functions are. Take, for example, the digital assistant that answers when you call an airline’s 1-800 number: It asks you a few questions and tries to get you the information you’re looking for. If it can’t, it connects you with a human who can. The bot’s narrowly defined purpose guides the human that’s interacting with it.

By contrast, a smooth-talking virtual assistant that tries to mimic human speech, whether out loud or on a screen, can create different assumptions. “The more human-like a system acts, the broader the expectations that people may have for it,” said Justine Cassell, a computer-science professor at Carnegie Mellon University.

Cassell told me about a [talking robot](#) she and a colleague built in the ’90s that had a 2-D animated head on a computer screen. The head made certain human-like motions as it talked, like shifting its gaze and nodding at appropriate moments in the conversation. Cassell found that when users saw those cues, they unconsciously began talking to the computer more quickly and less clearly, the way they might talk to an actual person. That caused problems. “When people used it that way, they actually made it harder for the system to function effectively,” Cassell said.

Siri, of course, doesn't have a face, and Facebook and Skype bots don't even have a voice. But modern bots try to trick users into interacting with them as if they're human in other ways: by bantering and using humor, speaking (or writing) conversationally, and learning to parse free-form questions and answers.

"This creates a perception that if you say anything to this bot, it should respond to it," said Nikhil Mane, an engineer developing conversational AI at Autodesk. That sets the bot up for failure, Mane said. Without a sense of a bot's limits, a user is liable to overstep the bounds of its ability. When the bot replies to an out-of-scope question in an odd or confusing way, it's a dissonant reminder of its artificial nature.

Mane pointed to Slack, the popular workplace messaging app, for an example of a better approach. The app's resident helper, Slackbot, prompts users to ask it "simple questions" about how Slack works. "I'm only a bot, but I'll do my best to answer!" Slackbot says. "If I don't understand, I'll search the Help Center."

I threw some basic questions at Slackbot, and it fielded them deftly. But ask it something outside of its purview—"Slackbot, what's the meaning of life?"—and it responds, "I'm sorry, I don't understand! Sometimes I have an easier time with a few simple keywords."

Compare that to how Poncho, a weather-app company that created one of the first Facebook chatbots, set expectations about itself when it was first launched. In [an interview](#) with *Business Insider* last year, the company's CEO said he was aiming to build "the first bot you want to be friends with." He wanted users to be able to "shoot the breeze" with Poncho on just about any topic, and hired writers to program non-weather answers into the chatbot. But Poncho wasn't even very good at delivering weather forecasts, its core function.

The bot's gotten better since then, and a little more humble. It answered my weather-related questions with few mistakes, and made some puns when I asked it about its favorite music and movies. Stray too far from the script, though, and the bot hits a wall, with a descriptive message: "So, I'm good at talking about the weather. Other stuff, not so good. If you need help just enter 'help.'"

The modesty was effective. Rather than trying to keep up the charade—one of Poncho's old answers when it didn't understand a prompt was "Sorry, I was trying to charge my phone. What were you trying to say?"—it acknowledged its fallibility. The message moved Poncho further away from science fiction—and out of the uncanny valley.

LATEST VIDEO



Rob Reiner on the Burden of His Name

The award-winning director on confronting his father's legacy and leaving his children to confront his own.

ABOUT THE AUTHOR



KAVEH WADDELL is a former staff writer at *The Atlantic*.

[Twitter](#) [Facebook](#) [Email](#)