

<https://www.cbsnews.com/news/google-duplex-assistant-artificial-intelligence-cnet-exclusive/>

By RICHARD NIEVA CNET May 9, 2018, 10:02 PM

Google's Duplex could make Assistant the most lifelike AI yet

- [Share](#)
- [Tweet](#)
- [Reddit](#)
- [Flipboard](#)
- [Email](#)

The conversation seems mundane. At least, at first.

In a building called the Partnerplex on [Google's](#) sprawling campus in Mountain View, California, I've been invited to hear a 51-second phone recording of someone making a dinner reservation.

Person 1: *Good evening.*

Person 2: *Hello?*

Person 1: *Hello.*

Person 2: *Hi, um, I'd like to reserve a table for Friday the third.*

Person 1: *OK, hold on one moment.*

Person 2: *Mm hmm.*

Person 1: *OK... hold on one second.*

Person 2: *Mm hmm.*

Person 1: *So Friday November third. How many people?*

Person 2: *For... two people.*

Person 1: *Two people?*

Person 2: *Yeah.*

As I listen to what sounds like a man and a woman talking, Google's top executives for [Assistant](#), the search giant's digital helper, watch closely to gauge my reaction. They're showing off the [Assistant's new tricks](#) a few days before [Google I/O](#), the company's annual developer conference that starts Tuesday.

Turns out this particular trick is pretty wild.

That's because Person 2, the one who sounds like a man, isn't a person at all. It's the [Google Assistant](#). And it (or "he"? That's its own debate) doesn't sound at all like the semirobotic, disembodied voice you usually hear out of a [Google Home](#) smart speaker when it's giving you updates on the weather or telling you how long it'll take you to commute to work. (You can hear the [Assistant's conversations](#) for yourself in the video below.)

This could be the next evolution of the Assistant, Google's rival to [Amazon's Alexa](#), Apple's Siri and [Microsoft's Cortana](#). It sounds remarkably -- maybe even eerily -- human, pausing before responding to questions and using verbal ticks, like "um" and "uh." It says "mm hmm" as if it's nodding in agreement. It elongates certain words as though it's buying time to think of an answer, even though its responses are instantaneously programmed by algorithms.

With this new speaking ability, Google Assistant gets that much closer to hitting a milestone in the evolution of computing: passing the Turing test. Proposed by English computer scientist Alan Turing in 1950, it's a way of potentially evaluating a machine's ability to demonstrate intelligent behavior. To pass the Turing test, a computer's natural language responses would have to sound just like a human's.

Built with [technology Google calls "Duplex"](#) -- and developed by engineers and product designers in Tel Aviv, New York and Mountain View -- the AI sounds as though the future of voice assistants has arrived.

Well, almost arrived.

The demo is part of what Google calls an "experiment" it plans to launch this summer. With Duplex, a "small" number of people will be able to book restaurant reservations and hair appointments and to check holiday hours, stuff you'd normally do over the phone. All of that back and forth happens on the back end -- between Google Assistant and, say, the restaurant. You won't even hear the voice chat taking place. It'll come from an unspecified phone number, not your own.

Google is coy about the size of the rollout, but says it'll be limited -- all the company will say is it won't be available to everyone using Assistant today. Nick Fox, vice president of product and design for Google Assistant and Search, and Yossi Matias, Google's vice president of engineering, say the search giant wants to "proceed with caution" because it's such a new technology.

Google also won't say if its goal is to bring this natural-sounding Assistant to more products -- like hearing it come out of your [Google Home](#).

Here's how the reservation feature works: Say you want to book a haircut appointment. Just ask the Assistant from your Google Home, or text it from your [Android phone](#) or iPhone app, telling it where and when. If you want something at noon on Sunday, the Assistant will ask you for backup options just in case that time slot isn't available -- it knows that getting a range between, say, noon and 2 p.m. may be the smarter way to go.

If the business uses an automated online booking service like OpenTable, the Assistant will default to using that system for the appointment. But if the business doesn't take automated reservations that way -- and many small businesses don't -- Google's helper will make the call for you. The Assistant will text you back confirming your appointment and add it to your calendar.

In a different recorded conversation, I hear two female voices setting up a hair appointment. One is the Duplex technology. "She" sounds naturally human. She even does that thing with her voice that annoys some people, where her inflection goes up at the end of a sentence, even when not answering a question.

"We think of it as, 'What's a perfect assistant?'" says Fox, a 15-year Google veteran. "When I want an assistant with me, it's there. It's always ready to help. I don't worry about 'Can my assistant do it or not'? My assistant can do anything I want it to do."

Fox, who's spent most of his time at Google working on search ads, is now in charge of product for one of Google's landmark projects. It's been almost exactly two years since Google announced the Assistant at I/O 2016. At the time, CEO [Sundar Pichai sat down with CNET](#) and explained his vision: The Assistant is the essence of Google, set free from the search box on the company's iconic, sparse homepage.

The goal for Pichai and his team is to have Google's artificial intelligence work alongside you as you go about your day: With the Assistant today, you can get morning news headlines from your Google Home, point your phone's camera at a flower to discover what species it is or turn up your thermostat so the house is a comfortable temperature when you get home.

But if Google's experiment becomes reality, the Assistant will truly be your digital assistant. "Think of it as building your own individual Google," Pichai told us two years ago.



Google has been expanding its line of Home products. The \$400 Home Max, powered by the Assistant, has higher-end audio than the regular Home.

Josh Miller/CNET

But as impressive as the Duplex phone recording is, we have to take the company's word for it. Google declined to do a live demo of the call during my briefing.

In the meantime, Google is trying to win more users with an [updated Assistant that speaks in six new voices](#) beyond the one female and one male it already has. The company is also making it easier for you to ask follow-up questions, and do things like send someone your ETA when you're driving and using Google Maps.

That personal touch

If the goal is for the Assistant to be Google personified, it still has a lot of evolving to do.

So far, Google's helper -- and honestly, Alexa, Siri and Cortana, too -- haven't lived up to their promise.

"We're still developing this technology, and we're actually working hard to get this right," Pichai said in Tuesday's keynote presentation at Google I/O. "We have many of these examples when the calls don't quite go as expected. But the Assistant knows the context and nuance. It knew to ask for wait times in this case, and it handled the interaction gracefully."

If Google ever decides to expand the Duplex technology to [Google Home and other consumer-facing products](#), that could change. Based on the short conversations I heard, it could feel like you're talking to an administrative assistant at the office or maybe even Samantha, the OS from the [Spike Jonze movie Her](#).

Of course, Google wasn't the first to the assistant market. Amazon, in a coup, surprised the world in 2014 when it released its Echo smart speaker, given voice by its Alexa software. For right now, smart speakers remain the best-known way to interact with a digital helper, and Amazon owns 69 percent of the market in the US. Google Home, released in 2016, has 31 percent, [according to CIRP, a research firm](#).

Both Google and Amazon have been beefing up their product lines. Last year, Google introduced the [\\$40 Home Mini](#) and [\\$400 Home Max](#), with high-end audio. Amazon has [10 Echo products](#), ranging in price from \$40 to \$230. Apple joined the battle in February with the [\\$350 HomePod](#).

But there's potentially bigger money to be made when digital assistants connect with the internet of things -- [controlling everything from your car's dashboard to your washing machine to your light switches](#). People will spend \$1 trillion on the internet of things by 2020, [according to Gartner](#).

That's why Google has been working hard to catch up. Today, the company announced that 500 million devices have shipped globally with the Assistant on them, including

speakers, [phones](#) and [TVs](#). Google also said the Assistant can speak in 30 languages and will be in 80 countries by the end of the year, including seven new ones: Denmark, Korea, Mexico, the Netherlands, Norway, Spain and Sweden. Amazon's Echo is also available in [more than 80 countries](#).

And last week, Google said the Assistant now partners with [5,000 household connected devices](#), up from 1,500 in January. That's an improvement, but still less than half of the 12,000 Alexa-enabled devices.

Google's Fox tries to downplay to the rivalry with Amazon.

"We think of it less in a competitive way," Fox says. "The bigger opportunity is for people who aren't using anything yet."

Still, both tech giants haven't been afraid to [play hardball and undercut each other](#) for [smart home](#) sales. Google in January cut off [Amazon Echo](#) and [Fire TV access to YouTube](#), which it owns. Amazon won't [sell some gadgets made by Google and Nest](#), the smart home maker [Google bought in 2014](#).

"We don't know all the answers"

Google has been working on its new natural-sounding voice technology "for years," says Matias, who leads the R&D center in Tel Aviv. The company says Duplex isn't broad AI technology for your phone. Instead, it's geared to specific tasks (like booking reservations, for starters). Matias readily admits that the tech isn't perfect and that Google is moving carefully because there are lots of thorny issues that come with a more human-sounding AI.

Like, what happens if the software starts talking nonsense? Matias tells me the tech won't go rogue and leave a trail of confused small-business owners talking to robots. When the feature actually rolls out, he says, the software will likely tell the person on the other end of the line that he or she is talking to a digital assistant -- even though the Assistant didn't do that during the recordings I heard.

"We need to be thoughtful about how we have this interaction while we're experimenting with it," Matias says. "We believe it's right to do this with transparency."

If a conversation does go awry, the software will fall back on what he calls "real-time supervised training." In that scenario, the software would turn over the conversation to a human to finish the call. Because calls aren't normally monitored by people, they'd have to begin from scratch. Google wouldn't say whether that person would have a log of what's already been discussed on the call.



Nick Fox, VP of product and design for Google Assistant and Search: "We don't want to force people into, 'This is what an assistant should sound like.'"

James Martin/CNET

The Duplex technology is the stuff of sci-fi. And that raises questions, both practical and philosophical: How should a person treat a lifelike AI? What's the proper human-robot etiquette or does that need to be developed? What's the best way for an AI to reveal it's a robot when talking to a human -- without disrupting the flow of a conversation?

Asked about these issues, Fox acknowledged that the folks at Google are learning about these things as they go.

"We need to set the right expectations with everyone," he says. "These are implementation questions of the technology that I'd say, humbly, we don't quite know all the answers to yet, and need to be figured out as we see this in the real world."

Paul Saffo, a Silicon Valley futurist now with Stanford University, has some thoughts on how this should play out. The creators of any realistic AI have the obligation to disclose to anyone who interacts with it that they're talking to a machine, he says.

But as this sort of technology becomes more widely used, he thinks people will begin to recognize the "tells" in a bot. "As our machines evolve, there's the hope that people will start to figure out if they're talking to one," Saffo says.

The Assistant everywhere

While Duplex remains an experiment for now, the Assistant is getting upgrades in other areas, some as soon as today.

That includes more voice options. There are six new ones, both male and female, that all speak with an American accent. One female voice is soft and soothing. A new male voice is a little deeper with a slight rasp.

"My wife hates that one. She's like, 'Get that out of my room,'" Fox says, as he plays a male voice that booms like an anchorman. The expanded library of voices is made possible by technology called WaveNet, which dramatically reduces the time and audio samples needed by a voice actor recording each track.

The voices aren't labeled with specific names. Instead, each is designated by a color, like red, orange or green. You select the one you want by scrolling through your phone's settings. When you're connected to a Google Home, you just pick your favorite voice setting and the speaker will recognize your voice and respond with the Assistant voice of your choice.

Google also teased that it will [bring singer John Legend's voice to the Assistant](#), too.

The idea, as with most of these updates, is to make the Assistant more personal. "We don't want to force people into, 'This is what an assistant should sound like,'" Fox says.

Another update, due this summer, is a mode called [Pretty Please](#) that will require you to be polite when giving the Assistant commands. The goal is to reinforce good manners in children -- and adults, too -- by prompting them to say please and thank you.



Google VP of Engineering Yossi Matias says the search giant has been working on natural-sounding voice technology "for years."

James Martin/CNET

Amazon announced a similar feature last month for a [new line of Echo Dot speakers designed for kids](#). The child-friendly version of Alexa offers positive reinforcement to a kid who says "please" by responding, "Thanks for asking so nicely."

While Amazon announced its version first, Google says it's been working on Pretty Please for a while. "All of this is really nascent," says Lilian Rincon, director of product management for the Assistant. "It's not surprising we'd both be grappling with these questions."

There's plenty of other new stuff: You can now [ask two questions at a time](#), and the Assistant will answer both in a row. (For example: How old is Kendrick Lamar and what's the name of his latest album?). You can also ask follow-up questions without saying "Hey Google" again or repeating the subject of your question. The mic will stay hot for up to 8 seconds before the window closes. Before this, the mic went inactive after you asked your question.

The Assistant can text people your ETA when you're navigating with Google Maps while driving. It will show you more visual content when you access the Assistant on your phone -- things like reminders, upcoming flight information and to-do items. You'll also be able to set up a "custom routine," which triggers a set of actions right after you shut off your morning alarm.

That means you could automatically start your day listening to [headlines](#), a traffic update and a favorite song to pump you up.

It's all useful. But it also means the Google Assistant is getting to know you a lot more intimately.

That's worth talking about because user privacy is a major topic of discussion in tech. The data collection practices of Silicon Valley's biggest companies have been under extra scrutiny in the weeks since [Facebook's scandal with Cambridge Analytica](#) began in March. The digital consultancy got access to personal information from up to 87 million Facebook users without their knowledge. The scandal has [raised questions about how tech giants secure your data and how they make money](#). For the most part, Google, Facebook and others rely on troves of personal information they've collected about your likes, dislikes, shopping preferences, lifestyle and more to allow marketers to target specific ads to you.

Google made \$95.4 billion in ad sales last year, while Facebook made most of its \$40 billion in 2017 sales from ads.

"Our principles here have always been to give users control over that information," Fox says. "The last few weeks have underlined that this is incredibly important. But we've always known it's important. It's healthy for us to get that reinforcement."

Assistant queries, he says, are treated exactly the same way as regular search queries, and have all the same user settings. That means you can clear your search history, but that may affect which ads come your way.

All that segues into the real reason we're talking about the Assistant. As Pichai noted back in 2016, it's the next evolution of Google search, with people spending less time with their desktop computers and more time with their phones and other internet-connected things, such as smart speakers and cars.

"Done correctly, it will save time for people and generate a lot of value for businesses," Pichai said on stage.

As my interview comes to a close, I ask Fox if he thinks the Assistant is Google's most important project right now.

He pauses for 13 seconds before answering.

"The Assistant is a critical effort for Google," Fox says. "It's something really important to get right."

This article originally appeared on [CNET](#) and was first published May 8, 10:40 a.m. PT.

Update, May 9 at 11:28 a.m.: Adds comments from Google CEO Sundar Pichai.