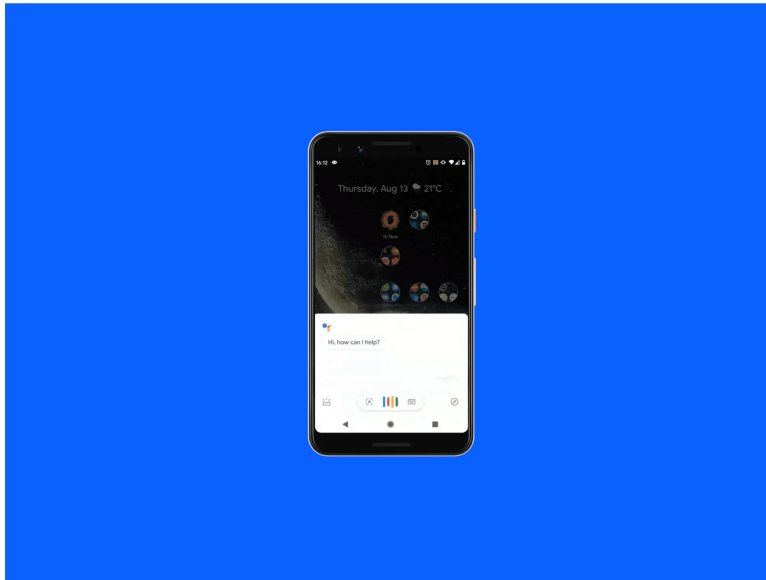


LG **GET THIS NEW OFFER**
 Save \$200 when you buy the GX Soundbar with the OLED G1 (\$55* +)
 SHOP NOW

JULIAN CHOKKATTU | GEAR | 04.20.2021 12:00 PM

Google Fixes Two Annoying Quirks in Its Voice Assistant

The AI just got a useful boost in understanding natural human speech.




PHOTOGRAPH: GOOGLE

I'M NOT MUCH of a cook, but the few times I've asked [Google Assistant](#) on my [Nest Mini](#) to start a timer in the kitchen have been hit or miss. All too often, the timer disappears into a void and Google can't tell me how many minutes are left. Other times, it takes multiple attempts to set it properly because Assistant struggled with understanding context.

Those problems (and a few others) are about to be resolved. Google's latest update to its voice assistant, which begins rolling out today, greatly improves its contextual understanding when you're asking it to perform a task like setting an alarm or a timer. Included in this update is another fix sure to be welcome for anyone who uses voice commands to manage calls and texts: You can finally teach Assistant how to properly pronounce your friend or family member's name.

Context Is Key



VIDEO: GOOGLE

Implicit List Selection

OK Google, set a timer for 20 minutes.

Second timer for 20 minutes, starting now.

⏸

When something feels different, it could be Alzheimer's.
 Now is the time to talk.
 LEARN MORE | alzheimer's association | ad.com

BUSINESS
Gummy Bears and Candy Bars Are Casualties of the Pandemic
 AARJAN MARSHALL

CULTURE
Made for Love, Handmaid's Tale, and the Trap of Dystopian TV
 KATE KNIBBS

SCIENCE
Why Not Turn Airports Into Giant Solar Farms?
 MATT SIMON

If you've interacted with a voice assistant, there's a good chance you've changed the specifics of your command mid-sentence. "Hey Google, set a timer for 20—no, 10 minutes." Until now, Assistant would've likely named your 10-minute timer "20, no." With the latest update, it understands you made a mistake, and that you just want 10 minutes on the clock.

ADVERTISEMENT

Make it a Parker
 Parker Boats

You've also been able to control multiple timers at once with Google Assistant for some time, but if you wanted to cancel one of them, that required some annoying back and forth. Assistant is now much faster at identifying which timer you want to cancel. And if you give each timer a name, like "eggs boiling," and then you said, "Cancel my egg timer," the old Assistant wouldn't understand what you're talking about, because the names don't match. The new update corrects that.

With alarms, if you asked Google Assistant before to move a previously scheduled alarm an hour later, it sometimes misconstrued that and set an alarm for one hour from the time of your request instead. Now it understands you were referencing a scheduled alarm, and it will make the adjustment properly.

The updated timer and alarm functions are available on screenless Assistant devices today (like Nest speakers) and will be coming to phones and smart displays at a later date.


These improvements come from a ground-up redesign of the system Assistant uses for natural language understanding. Amarnag Subramanya, a distinguished engineer at Google who leads the NLU and Conversational AI teams on Google Assistant, says it allows for far more natural conversations between us humans and our nonhuman helpers.

"Today, when people want to talk to any digital assistant, they're thinking about two things: what do I want to get done, and how should I phrase my command in order to get that done," Subramanya says. "I think that's very unnatural. There's a huge cognitive burden when people are talking to digital assistants: natural conversation is one way that cognitive burden goes away."

Making conversations with Assistant more natural means improving its reference resolution — its ability to link a phrase to a specific entity. For example, if you say, "Set a timer for 10 minutes," and then say, "Change it to 12 minutes," a voice assistant needs to understand and resolve what you're referencing when you say "it."

The new NLU models are powered by machine-learning technology, specifically [bidirectional encoder representations from transformers](#), or BERT. Google unveiled this technique in 2018 and applied it first to Google Search. Early language understanding technology used to deconstruct each word in a sentence on its own, but BERT processes the relationship between all the words in the phrase, greatly improving the ability to identify context.

An example of how BERT improved Search ([as referenced here](#)) is when you look up "Parking on hill with no curb." Before, the results still contained hills with curbs. After BERT was enabled, Google searches offered up a website that advised drivers to point wheels to the side of the road. BERT hasn't been problem-free though. Studies by Google researchers [have shown](#) that the model has associated phrases referring to disabilities with negative language, prompting calls for the company to be more careful with natural language processing projects.



The WIRED Guide to Artificial Intelligence

Supersmart algorithms won't take all the jobs. But they are learning faster than ever, doing everything from medical diagnostics to serving up ads.

BY TOM SIMONITE

But with BERT models now employed for timers and alarms, Subramanya says Assistant is now able to respond to related queries, like the aforementioned adjustments, with almost 100 percent accuracy. But this superior contextual understanding doesn't work everywhere just yet—Google says it's slowly working on bringing the updated models to more tasks like reminders and controlling smart home devices.

William Wang, director of UC Santa Barbara's Natural Language Processing group, says Google's improvements are radical, especially since applying the BERT model to spoken language understanding is "not a very easy thing to do."

See **What's Next in Tech with the Fast Forward Newsletter**
From artificial intelligence and self-driving cars to transformed cities and new startups, sign up for the latest news.

Your email

Enter your email

SUBMIT

By signing up you agree to our [User Agreement](#) and [Privacy Policy & Cookie Statement](#)

"In the whole field of natural language processing, after 2018, with Google introducing this BERT model, everything changed," Wang says. "BERT actually understands what follows naturally from one sentence to another and what is the relationship between sentences. You're learning a contextual representation of the word, phrases, and also sentences, so compared to prior work before 2018, this is much more powerful."

Most of these improvements might be relegated to timers and alarms, but you will see a general improvement in the voice system's ability to handle unstructured context. For

Most Popular



BACKCHANNEL
They Told Their Therapists Everything. Hackers Leaked It All
WILLIAM RALSTON



BUSINESS
Gummy Bears and Candy Bars Are Casualties of the Pandemic
AARLAN MARSHALL



CULTURE
Made for Love, Handmaid's Tale, and the Trap of Dystopian TV
KATE KNIBBS



SCIENCE
Why Not Turn Airports Into Giant Solar Farms?
MATT SIMON

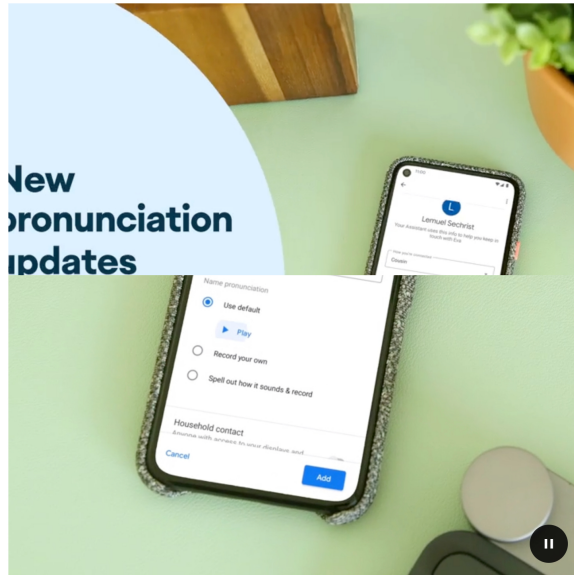


Ad SPONSORED CONTENT
What did it take to win the \$20M NRG COSIA Carbon XPRIZE?

SPONSORED BY X PRIZE

general improvement in the voice assistant's ability to broadly understand context. For example, if you ask it the weather in New York and follow that up with questions like "What's the tallest building there?" and "Who built it?" Assistant will continue providing answers knowing which city you're referencing. This isn't exactly new, but the update makes the Assistant even more adept at solving these contextual puzzles.

Teaching Assistant Names



VIDEO: GOOGLE

Assistant is now better at understanding unique names too. If you've tried to call or send a text to someone with an uncommon name, there's a good chance it took multiple tries or didn't work at all because Google Assistant was unaware of the proper pronunciation.

ADVERTISEMENT

Thankfully, Google's new voice modeling technology now lets you read out a name to Assistant so it can better identify it. You'll just need to manually set this up in Assistant's settings on your phone. Better yet, your voice recording isn't uploaded to the cloud and sent to Google. "We're able to learn the aspects of the pronunciation without having to store the audio," Subramanya says.

This improved name recognition is only available in English for now on Android phones, smart speakers, and smart displays. Google says it hopes to expand it to other languages soon.

Subramanya says Google's quest to make conversations with its voice assistant more natural is not unlike its gradual updates to Google Search over the years. "If you go to the very early days of Google Search, you had to think precisely about your queries. Now you can just say 'Coffee shops nearby.' You don't have to think so much about the set of words you use. We're seeing a similar progression with digital assistants."

But don't expect to have long, back and forth conversations with your AI just yet. Wang says machines still have a hard time processing certain kinds of requests. For example, they're still not very good at providing answers learned from images, videos, or other sources.

"There's still a long way to go for machines to be able to talk to humans and really understand naturally and be able to respond naturally," he says.

More Great WIRED Stories

- 📧 The latest on tech, science, and more: [Get our newsletters!](#)
- The cold war over McDonald's [hacked ice cream machines](#)
- What octopus dreams tell us about the [evolution of sleep](#)
- The lazy gamer's [guide to cable management](#)
- How to log in to your devices [without passwords](#)
- Help! Am I [oversharing with my colleagues?](#)
- 🔍 Explore AI like never before with [our new database](#)
- 🎮 WIRED Games: Get the latest [tips, reviews, and more](#)
- 🏃 Want the best tools to get healthy? Check out our Gear team's picks for the [best fitness trackers](#), [running gear](#) (including [shoes](#) and [socks](#)), and [best headphones](#)



Julian Chokkattu is a senior associate editor at WIRED, covering personal technology and reviewing consumer products. Previously he was the mobile and wearables editor at Digital Trends, steering coverage and reviews of smartphones and smartwatches, and an intern at



TechCrunch. He graduated from Rutgers University with a bachelor's degree in... [Read more](#)

SENIOR ASSOCIATE REVIEWS EDITOR

COMMENTS

Join the conversation [REGISTER](#) [SIGN IN](#)

Welcome to our community. We invite you to join our discussion. Our [community guidelines](#) are simple: be respectful and constructive, keep on topic, support your fellow commenters, and report bad behavior.

If you're new to our comments section, your first comment will be held for moderation.

And if you're interested in accessing all WIRED has to offer, [subscribe now](#).

Post a comment

B *I* **11** Spoiler

[Sign in and Join the Conversation](#)

All Comments

Sort by [Newest](#)

– **Gynobot** 5 days ago

Yeah, but will they add back the functionality to learn your voice and unlock your phone so you don't ask Google to do something through your earbuds, then have to fish your phone out of your pocket to unlock it at which point you might as well do the thing yourself? 🤔

This is a really relevant problem if you use Bluetooth while driving, especially a motorcycle, scooter, or bicycle. F'real.

[Respect](#) [Reply](#) [Share](#)

[Report](#)

[Book of comments](#) [Book of article](#)

Featured Video

