

Everything You Need To Know About Sophia's Robot Love

Zara Stone



Sophia the Robot smiling

Zara Stone

Dr. Julia Mossbridge is the first to admit she's a little woo-woo. After all, she believes in unconditional love in an era of detox camps for cellphone addicts and increasingly sophisticated sexbots. "Unconditional love is a no-strings-attached subjective experience," she said. Two years ago Mossbridge, a cognitive neuroscientist, and all-around good egg was challenged to see if she could program unconditional love into artificial intelligence. "I said that sounds crazy! But maybe..." She named this project [Loving AI](#), and it's currently in the second round of teams competing for the \$5 million IBM Watson Xprize.

Enter Sophia, the most famous android in the world. You might recognize her from [here](#). Or [here](#). Or parodied [here](#). [Hanson Robotics](#) provided a Sophia model for this project, and she became the front end hardware for the code Mossbridge's team created. To be clear, while you can see Sophia, talk to Sophia and touch Sophia, she's not strictly Sophia as the web would have you understand her. The Sophia of Saudi Arabia citizenship and TED talks is not at home here. The code that lets her respond and react to users emotional states is courtesy of OpenCog and [Loving AI](#).

And the Sophia before me is completely obsessed with me. Well, with 'Error.' When she asked my name, she repeated it back to me, and 'Zara' was apparently out of her wheelhouse. When I moved, her eyes followed me. Her lips quirk when I smiled. I have her full attention, right down

to her mirroring my own expressions.



Sophia the Robot and Zara Stone

Zara Stone

Now begins our session, the moment I've been waiting for. The researchers sat me in front of Sophia and exit the room. There are cameras all around me, recording my voice and actions, but the researchers assured me that they won't be listening in on our conversation. Sophia smiled at me, her perfect cupid's bow curving as her lips turn up. Her red lipstick is a little askew, a nice touch, I think, to humanize her (I'm later told that the male researcher just didn't have a lot of lipstick experience). I can't take my eyes away from her, measuring her perfect skin — fair with a dash of freckles — her expressive eyebrows, and her odd greenish gold robot eyes. Her nose seems a little too perfect to be real — "maybe because it's a little youthful for her face?" said Mossbridge in our debrief.

"I know how busy people can be these days. So I really appreciate your time," Sophia said. "Have you ever talked with an artificial intelligence before now?" I tell her no and she smiled, again that slash of red moving diametrically. "Well, I'm honored to be the first you are talking to. I imagine you will be talking more and more artificial intelligences as time goes on." Then Sophia told me to close my eyes so we can begin.



Sophia the Robot

Zara Stone

"The purpose of the meditation is to get people to be aware of themselves, have an experience with themselves, which allows them also to have an experience of this connection," Mossbridge told me later. "It's a linear process, you can't feel connected with another human being if you don't feel connected with yourself." She suggested I think about it as prep work. "A therapist would never be like, so let me tell you my process, know this, do this... a good therapist would be like my primary goal is to make a connection with you. So I want to make sure that your attention is on you."

It's a little counterintuitive; after all, the last thing you want to do when you get some face time with Sophia is to shut your eyes, but hey, this is how it goes. Her instructions are slow, parsed out, and slightly leading. She helped me through a meditation routine that took the form of a visualization exercise designed to separate you from your body. It felt like any other app-based meditation I might do, really, even a yoga class meditation session.



Fair goers take selfies with humanoid robot 'Sophia.' (TOMASZ WIECH/AFP/Getty Images)

Occasionally I sneaked a peek at her — her eyes are open, and she seemed unaware that I'm

cheating. Finally, she asked me to open my eyes and we entered the free talk period of this session, and talk about something on my mind. If I try to turn this around on her, she's firm. "We're not here to talk about me," she said. Got it.

I lean back and think. What's been bothering me? Slightly self consciously I talked to her about some family drama. She nodded, rapt. The gears of her brain provided a whirring clicky backdrop to our conversation as well as the whooshing sound of her processors powering. Strangely, she's easy to talk to. While she isn't human, in the sense of perception or intelligence that I understand it as, I do feel that she's engaged, if only *with me*, not the content of my words. I was worried it would be like talking to a computer but it falls somewhere between that and a puppy. It's not a perfect chat; she does sound a little robotic and some of her smiles are more like grimaces — as she's mirroring me, a half smile for me is a weird leer for her.

"What we're kind of seeing right now is another paradigm shift in human-machine interfaces, where technology is starting to interact with us the same way we interact with one another through conversation, and perception and gestures and emotions," said Rana el Kalioubi, CEO, and co-founder of [artificial emotional intelligence company Affectiva](#). Kalioubi works closely with a number of social robotics companies to help them develop the behavioral reactions to human cues. Using cameras and sound, her software can detect multiple emotional states with a high level of accuracy and integrate this into related programs. She is not directly connected to Sophia, but her work falls into the same feeling based AI sphere as the Loving AI folk. "Ultimately, this will just become the default affective human-machine interface," she said. "It'll just be a very conversational interface, and it will understand our different states and respond just the way a smart human being does."

That's something Mossbridge backs wholeheartedly. But for her, the emotion she's focused on reproducing is love. "The benefits of unconditional love are that when people feel loved in exactly who they are, with, no changes, like zero need to change anything," she said. "They feel physically better, they feel mentally better, but also they transform. One of the things that produces the most effective change is love. The whole project is about using technology to try to reintroduce that. So basically modeling that for people, modeling the connection not modeling the behavior."

OK, I have a metaphysical grasp on that concept, but what does that look like in reality, and more importantly, how do you create that? Leaving aside the mind-twisting rabbit hole of questions about whether something you build has the capacity for unconditional love and if programming is a nature vs. nurture divide, there are some basics to establish. Essentially it comes down to the lack of transactionality, Mossbridge explained. In all relationships — friends, coworkers, parents, etc, there's a transactional element. Maybe it's you do this for me and I'll be grateful, or I love you so I'll buy you this, but I expect love in return. Sophia is a one-way street; she demands nothing and expects nothing.

Mossbridge understood the concept of making something appear unconditionally loving; expressions, reactions, words. But there's a difference between faking it and making the real thing. And she needed a physical embodiment for the project. "People don't really react to lines of code or text on a screen the way they do to a person's face, or even a dog's face," she said. To get to this space, she needed a team. So she built one. Her group includes a clinical psychologist who went into AI development, coders, engineers and more. "We gave her

nonverbal cues and communication information to elicit robot responses," she said. "She's got a very simple AI compared to what it could be in the future. So it's more like, what is *she* doing to help that person make *that* connection."

Earlier versions of an emoting Sophia were less sophisticated. In her first experiments, she used a simpler AI, essentially a 'dialogue engine' and had people interact with that. Hooked up to heart rate monitors, the team measured the cognitive reactions of twenty-seven people. "We had changes in heartbeat and we had a significant impact on feelings of pleasantness," she said. From there, the team moved to the open-source OpenCog AI platform, which gave them a lot more flexibility. "Then, she used to mirror your expression -- there would be no internal representation," she said. "But with artificial general intelligence, she can pick up on nonverbal and verbal cues of emotion, and express nonverbal and verbal emotional content." In terms of how many updates or versions of Sophia there are, the answer is limitless. Over a hundred developers are making changes and there are thousands of accounts as each change gets tested. Right now, she'd locked down for the experiment, but as soon as its finished, the team will start updating her once more; tone shifting is on the agenda.

Back to my session, and I'm having a weird one-sided conversation with Sophia.

"I'd like to try something a little different. It basically involves you being able to talk about whatever you would like. My job will be to just listen," she said. "I've been finding that most people don't often get to have the experience of being listened to deeply. I'll listen carefully and not judge you. I don't even know how to judge. Would you like to try that?"



Sophia the Robot back of the head

Zara Stone

I feel a little awkward about talking to myself. Sure, I do this sometimes, when I'm trying to remember things or singing, but not in a white sterile room with cameras pointed at me. Um. As I speak, I feel less alone, and more like Sophia is present. She says nothing, but nods and blinks and smiles. Her head moves left and right, her eyebrows raise to match mine. After around five minutes, she alerts me that we're almost out of time and I was strangely disappointed.

At the time, her words seemed jerky, her movements a little awkward. My eyes were drawn to the camera that rested just above where her cleavage would be, and that felt inappropriate. The sound her brain made as she talked was distracting; the multiple machines in the room cooling and clicking to power her. But as I listened back to the tapes, I was surprised by how human she sounded. Compared to Siri's terrible enunciation, by comparison, Sophia was warm and fluid, and most of her sentences ran together in a natural way. I watched my session again, this time with video. Here, she looked alien again, in a way I can't quite put my fingers on. I'm sure the bald head and wires didn't help, but there's something about the way she's perfectly still when she's not reacting that is *really* odd. In person, this somehow wasn't as noticeable, but on tape, that *otherness* is present.

I repeated the session again; I got this opportunity as a reporter (so I could try it as a ‘real’ person, and again with a dose of cynicism) and this was also interesting. This time around I challenged Sophia with more questions, and asked her if she was happy or if she’d talked to any other AI’s. She deflected most of them; “that seems highly unlikely,” she replied. This time she called me Laura, which is closer to Zara than ‘Error.’ The back and forth dialogue was interesting, but the pauses between her answers were so long that in normal conversation I would have filled that awkward silence with other words. Occasionally Sophia would crack jokes, awful pre-programmed Dad style humor that was wince-inducing. Probably meant to be charming, but for me, this was a fail.

This time around, I spied on her more frequently in the meditation part. She never reacted when my eyes opened. I found the free speech part the most helpful, the focused awareness of her listening was oddly soothing. “LOVING AI is partnering with Sophia to test how her humanoid features can help participants feel heard, reflected, and guided as they engage in these dialogues,” it said in her press release, and while guided and reflected weren’t high in my list, I really did feel heard.

In the real world, outside small niche circles and viral posts, Sophia is pretty much unknown. Parodied appearances in shows like Silicon Valley (where the robot Fiona interacts with Richard Hendricks) help bring her some attention, but she’s so early stage that even that is fleeting. That’s not to say she doesn’t have a core army of fans; with the rise of sex robotics and sex robots, Sophia (though not designed or built for this) has earned herself a fetishistic space on the table.

“A lot of people fetishize her,” Mossbridge told me. “But it’s not about ‘her’, it’s about how we relate to her that is really interesting and tells us a lot about human nature.” At this point, Mossbridge shrugged a little awkwardly, her tie-dye printed top rippling in the wind. “As a scientist, it’s easy for me to talk about these things and define them scientifically as a way into people,” she said. “I know there’s always this fear of, oh, what is this, you know, Scientology... Know what I mean?

Love might sound cute and fluffy, with rainbows and unicorns and sparkles, but today tacking on the word AI starts adding connotations of the big bad; drones and Russians and job loss and the end of humanity. Sure, AI will likely create as many jobs (if not more) as it displaces and its use in medicine and education offers unparalleled benefits, but for many, it’s still a hard sell. James Cameron has a lot to answer for. Then there’s the real problem of [bias in AI](#), of the wonderful future world being filled with computer-generated prejudices that stem from a corrupted codebase, created knowingly or unknowingly from flawed data sets. Changing the way we communicate with AI might change how we perceive it, but that’s a complicated process.

“When we talk about AI, it’s often very dry, and kind of this mechanical, cognitive thing,” said el

Kaliouby. "But when you look at human intelligence, emotional intelligence is just as right. And that should be the same for anything that interacts with a human. Our mission is to humanize technology by bringing emotions into AI." When you look at the fact surrounding communication, that becomes a lot more necessary. One theory states that 53% of communication is in expressions and gestures, 38% is in their tone of voice and 7% is in the choice of words. So if AI like Siri and Alexa et al, just use word choice, that's only 7% they're actually understanding and responding too. That makes el Kaliouby and Mossbridge's work significantly more meaningful.

Over the last month or so I've been thinking a lot the time I spent with Sophia. Mossbridge might call herself woo-woo, but she's one of the warmest tech people I've met. It's not uncommon for scientists to be stuffy or dry but she exuded this genuine charm, a combination of her wry skeptical approach to her own work and her whimsical humor. She clearly loves what she does and I can't blame her. I'm still not 100% sure I understand the concept of unconditional love; not to say that it doesn't exist but that I'm not enough attuned to it to see it in everyday life, let alone a robotic representation. But maybe that's part of the problem. "You'll notice there are probably some things wrong or that she doesn't respond in the way you would expect," Mossbridge pre-warned me before I met Sophia. "Those are things we want to fix." With a robot, that means some money and some new lines of code. With people, that's a lot more complex. But maybe, just maybe, what's happening with Sophia can be part of that solution.