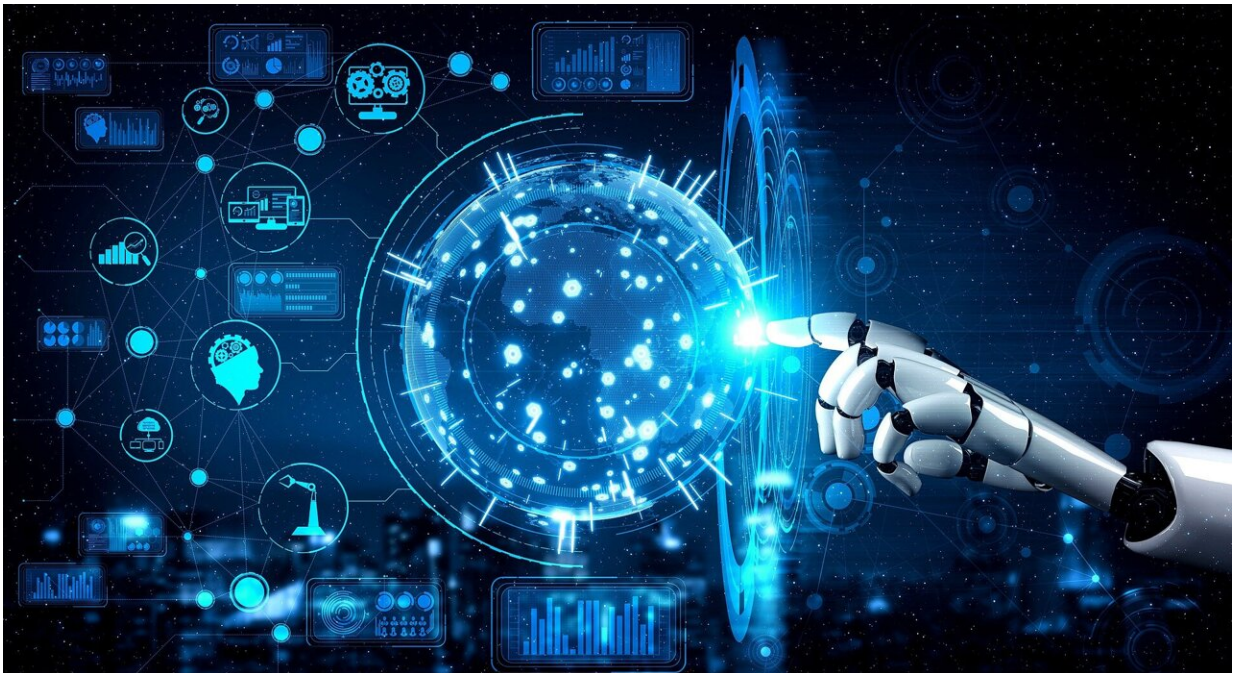# Q&A: Anthropologist discusses why we are so tempted to treat AI as 'god-like'

October 5 2023, by Morgan Sherburne

Credit: Pixabay/CC0 Public Domain

As artificial intelligence apps such as ChatGPT have proliferated, so have chatbots with a religious bent. People facing a moral or ethical dilemma can submit their questions to these chatbots, which then provide an answer based on the religious texts fed to them or crowd-sourced data.

Webb Keane, University of Michigan professor of anthropology, recently co-wrote an [op-ed](#) about what he and his co-author call "godbots," and the danger of giving moral authority to artificial intelligence.

People are becoming increasingly familiar with artificial intelligence and chatbots. But many may be surprised to learn about 'godbots.' Can you explain how these religious chatbots work and why they are unique in AI?

My co-author Scott Shapiro, a professor at the Yale Law School, and I came up with this term to describe a curious development that has emerged with generative AI such as ChatGPT. It turns out that very quickly after the development of ChatGPT, we started seeing bots that were specifically designed to give advice on moral and [ethical questions](#). Some of these were explicitly religious. For example, several very quickly showed up that will speak in the voice of Krishna and tell you what to do as a Hindu in a thus-and-such situation.

There's one where you can talk to Jesus Christ. And the one that especially interests me is called AskDelphi. It's named after the Delphic Oracle of ancient Greece, which was a massively influential institution in Greece that lasted for centuries, in which a medium went into spirit possession and responded to people's questions.

What the designers of AskDelphi claimed to have done is crowdsourced people's moral intuitions. They presented people with a variety of ethical quandaries—is it OK to cheat on a test if I really need the grade, or something similar. They then get a huge number of reactions and responses, from which the AI generates its advice. So now you can bring your [moral dilemmas](#) or ethical problems to this app. Of course AI is a

fast-moving target, but at the time I checked it out, the answers it gave were clear and decisive, with no consideration of complications or alternatives.

What we're calling the godbots here are taking advantage of a more general human propensity. This is something I want to stress: The temptation to turn to AI for answers to our hard questions isn't just for religious people. And don't think only the gullible are drawn to it. Godbots playing with something that's much more general. It is the tendency that people have to look for answers that have authority, that are totally certain.

We all know that when we're faced with really troubling or puzzling dilemmas, especially moral quandaries, it's comforting to have someone you can turn to who's going to tell you what the answer is. And when we face ultimate questions, we may want something more than just a friend's advice. A godbot is just a very extreme case of this, a source that gives you an authoritative answer, which comes from something beyond us, something that surpasses human limits.

We argue that this is why even rationalistic and secular people so easily talk about AI in religious terms, as if it was some kind of divine or magical source of wisdom. This is why you have Elon Musk calling AI "godlike" and the historian Yuval Noah Harari saying it will create a new religion.

Can you talk about what makes us susceptible to desiring such concrete answers?

The question we ask is, "What is it about the chatbot that makes it seem like a good place to turn for answers?" Our answer is that the design of the chatbots invites us to treat them like more-than-human oracles. Why? First of all, they're opaque. They don't show you their workings.

And so they can trigger a cognitive response in people that actually has a very long history. They do what oracles, prophets, spirit mediums and divination practitioners have always done. They have access to a source that is totally mysterious. It can seem to be tapping into something that just knows more than I do. A source like that seems more than human. It can appear to be divine.

If you go through the history of human divination techniques, you see this is repeated over and over again, whether it's ancient Chinese casting the I-Ching or Yoruba casting cowrie shells. An example we use in our article is sacrificing animals and then studying their entrails to find marks that come from the spirit world, a very common practice found from ancient Rome to many contemporary societies. Or the Delphic Oracle, who seems to have been a spirit medium, someone who went into a trance and whose words, sometimes quite enigmatic, seemed to come from elsewhere.

You don't have to believe in divine authority for this to work. You just need to feel that AI surpasses humans. The urge to turn to it for answers can start with no more than that. I really want to stress this point: we are not saying "Well, some suckers will fall for this." The godbots are just an extreme case of something that's actually much more common. People who pride themselves on their scientific rationality are susceptible, too.

Now, the second aspect of the chatbots is that they're designed to give you one answer and to give it with complete authority, without any doubts. When Harry Truman was president, he supposedly complained about his economic advisers: "When I ask them for advice, they say, 'Well, on the one hand, this, on the other hand, that.'" Truman said, "I want someone to find me a one-armed economist!"

Right now that's what chatbots do. This is one way they're more dangerous—and perhaps more appealing—than, say, the Google search

function. Google says, "Well look here's a whole bunch of sources." So it's at least implying that there isn't necessarily just one answer. Look at all these different sources! And if you want, you can go further into them, even compare them to one another.

Chatbots in their current state aren't like that. In effect they say, "I'm not going to tell you where I got the answer. You just have to accept it. And there's only one answer." Life is complex, is often bewildering and there's an irresistible attraction to things that promise to make it simpler.

And again, it's the design of the chatbot that because of its opacity, on the one hand, it has all the authority of crowdsourcing. For better or for worse, we've come to place a huge amount of faith in the wisdom of the crowd, and then projected it onto chatbots. As a result, it seems to know more than any human could possibly know. So, how can you doubt it?

And its inner workings are opaque—even computer programmers will tell you that some of the things going on in these algorithms are just too complex to explain. It's not necessarily that they don't understand their own devices, but that the explanation can be just as complicated as the thing it's meant to explain.

How are these chatbots designed? How do they gather data?

I'll use the example of something called the Moral Machine Project, based at MIT. As [self-driving vehicles](link) proliferate, the risk that they'll make bad decisions in a pinch is growing. What if they have to choose between hitting a pedestrian or turning into oncoming traffic and possibly killing their passengers? So the Moral Machine project aims to design an algorithm that will solve this problem. They created a computer game with the whole series of scenarios involving choices

between different fatal outcomes. And they were very pleased. They got more than a million people to play it.

This is big data, that looks like it gives us real answers about the best, or at least the most universal, human intuitions. But then if you start to look into the details and ask who these million people are, they turn out to be far from a representative sample. They were overwhelmingly males under the age of 35 who like to play computer games. And that means they were the sort of folks who have easy access to computers and with the leisure time to play these games. And who finds games like this fun to do? So, is this a good sample of humanity? No. But because they have them in large numbers, a million answers, it looks like really solid stuff.

These are the kinds of problems you run into when you rely on crowdsourcing for something like ethical problems. For instance, this approach tends to reduce moral dilemmas to something like a crossword puzzle or a video game where you show how clever you can be at reasoning out the answer. It removes the player from the kind of context in which real life moral dilemmas actually happen, which often involve real relationships to other people, fraught with emotional turmoil, confusion and so forth. That's a very distorted way to think of ethics.

Is there a danger in relying on this kind of AI?

Really, what I am worried about is the way in which these godbots push us in the direction of thinking about life's dilemmas as being like algorithms or games, things you can resolve with clever calculations. This has a distorting and very limiting effect on what we understand ethics to be. Second, it also encourages us to think there's always going to be a single, right answer. Third, it's giving authority to a machine, and tempts us to forget that, at the end of the day, the data come from human beings. If a person tells me, "You should do thus and such," I can just say "Well, I know who you are. I know where you're coming from. If we

have a history with one another, I know how that might shape your answer, too."

But if it's coming from an algorithm, it seems to have this cool, objective superiority. It hides its human sources. What most worries us is how it may be displacing our authority over our own thought processes and moral intuitions.

Provided by University of Michigan