

How to build a likable chatbot

January 12 2021, by Ranjay Krishna

Why do some chatbots, like Microsoft's Chinese chatbot Xiaoice, above, amass millions of monthly users, while others get quickly discarded? Credit: Microsoft

A few years ago, [Business Insider predicted](#) that 80% of enterprise applications would use chatbots by 2020. Today, the internet is flooded with millions of conversational artificial intelligence agents. Yet only a handful of them are actually used by people—most are discarded.

Even though the technical underpinnings of these agents continue to improve, we still lack fundamental understanding of the mechanisms that influence our experience of them: What factors influence our

decision to continue using an AI agent? Why, for example, did Microsoft's Chinese chatbot Xiaoice amass millions of monthly users, while the same techniques powering Microsoft's English version, Tay, led to it being discontinued for eliciting antisocial troll interactions?

Unfortunately, existing theories do not explain this discrepancy. In the case of Xiaoice and Tay, both agents were based on the same underlying technology from Microsoft, but they resulted in very different reactions from users. Many AI agents have received polarized receptions despite offering very similar functionality; for example, emotional support chatbots Woebot and Replika continue to evoke positive user behavior, while Mitsuku is often subjected to dehumanization.

Our team of researchers from Stanford was interested in studying the effects of an important and unexamined difference between these otherwise similar AI agents—the descriptions attached to them. Words are one of the most common and powerful means that a designer has to influence user expectations. And if words can influence our expectations, they can also impact our behavior and experiences with AI agents.

Descriptions, or more formally metaphors, are attached to all types of AI systems, both by designers to communicate aspects of the system and by users to express their understanding of the system. For instance, Google describes its search algorithm as a "robotic nose" and YouTube users think of the recommendation algorithm as a "drug dealer," always pushing them deeper into the platform. Designers often use metaphors to communicate functionalities of their systems. In fact, they have used metaphors for decades, starting with the "desktop" metaphor for personal computing to "trash cans" for deleted files, "notepads" for free-text notes, and analog shutter clicking sounds for mobile phone cameras (your phones certainly don't have to make that sound to take a photo).

Today, AI agents are often associated with some sort of metaphor. Some, like Siri and Alexa, are viewed as administrative assistants; Xiaoice is projected as a "friend," and Woebot as a "psychotherapist." Such metaphors are meant to help us understand and predict how these AI agents are supposed to be used and how they will behave.

In our recent preprint paper, my coauthors—HAI co-director and Stanford computer science professor Fei-Fei Li, communications professor Jeffrey Hancock, computer science associate professor Michael Bernstein, and Carnegie Mellon University graduate student Pranav Khadpe—and I studied how these descriptions and metaphors shape user expectations and mediate experiences of AI agents while keeping the underlying AI agent exactly identical. If, for example, the metaphor primes people to expect an AI that is highly competent and capable of understanding complex commands, they will evaluate the same interaction with the agent differently than if users expect their AI to be less competent and only comprehend simple commands. Similarly, if users expect a warm, welcoming experience, they will evaluate an AI agent differently than if they expect a colder, professional experience.

We recruited close to 300 people based in the United States to participate in our experiment where they interacted with a new AI agent. We described this new agent to each participant with different metaphors. After interacting with the agent to complete a task, participants were asked to report how they felt about the agent. Would they want to use the agent again? Are they willing to adopt such an agent? Will they try to cooperate with it?

Our results suggest something surprising and contrary to how designers typically describe their AI agents. Low-competence metaphors (e.g., "this agent is like a toddler") led to increases in perceived usability, intention to adopt, and desire to cooperate relative to high-competence metaphors (e.g., "this agent is trained like a professional"). These

findings persisted even if the underlying AI performed at human level. This result suggests that no matter how competent the agent actually is, people will view it negatively if it projects a high level of competence. We also found that people are more likely to cooperate with and help an agent that projects higher warmth metaphors (e.g., "good-natured" or "sincere").

Finally, with these results in mind, we retrospectively analyzed the descriptions attached to existing and past AI products, such as Xiaoice ("sympathetic ear"), Tay ("fam from the internet that's got zero chill!"), Mitsuku ("a record-breaking, five-time winner of the Turing Test"), and showed that our results are consistent with the user adoption and behavior with these products. Tay elicits low warmth and attracted a lot of antisocial users; Mitsuku projects high competence and was abandoned; Xiaoice projects high warmth and positively engages with millions of users.

Descriptions are powerful. Our analysis suggests that designers should carefully analyze the effects of metaphors that they associate with the AI systems they create, especially whether they are communicating expectations of high competence.

More information: Pranav Khadpe et al. Conceptual Metaphors Impact Perceptions of Human-AI Collaboration, *Proceedings of the ACM on Human-Computer Interaction* (2020). [DOI: 10.1145/3415234](https://doi.org/10.1145/3415234)

This story is part of [Science X Dialog](#), where researchers can report findings from their published research articles. [Visit this page](#) for information about ScienceX Dialog and how to participate.

Provided by Stanford University

Citation: How to build a likable chatbot (2021, January 12) retrieved 25 April 2024 from <https://sciencex.com/news/2021-01-likable-chatbot.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.