

Developing bots that talk more like people

January 24 2018, by Rob Matheson



Giant Otter Technologies, founded by artificial intelligence video game developer Jeff Orkin, an MIT Media Lab alumnus, has developed a platform that uses AI algorithms and crowdsourced annotators to build a natural-language database that helps chatbots expertly navigate human conversations. Credit: Chelsea Turner/MIT

Before coming to MIT, Jeff Orkin SM '07, Ph.D. '13 spent a decade

building advanced, critically acclaimed artificial intelligence (AI) for video games.

While working on F.E.A.R., a survival-horror first-person shooter game, he developed AI that gave computer-controlled characters an unprecedented range of actions. Today, more than 10 years later, many video game enthusiasts still consider the game's AI unmatched, even by modern standards.

But for Orkin, the game's development inspired a new line of interest. "A big focus of that game was getting squads of enemy characters to work together as a team and communicate constantly," Orkin says. "That got me really into thinking about, 'How do you get machines to converse like humans?'"

Following the game's release in 2005, Orkin enrolled in the MIT Media Lab, where he spent the next eight years tackling that challenge. Now, through his startup Giant Otter Technologies, he's using his well-honed AI skills to help chatbots expertly navigate tricky human conversations.

Giant Otter's [platform](#) uses AI algorithms and crowdsourced annotators to build a natural-language database, compiled "bottom-up," from archived sales and [customer support](#) transcripts. Chatbots draw on this robust database to better understand and respond, in real time, to fluctuating, nuanced, and sometimes vague language.

"[The platform] was inspired by the way episodic memory works in the human mind: We understand each other by drawing from past experiences in context," says Orkin, now Giant Otter's CEO. "The platform leverages archived data to understand everything said in real time and uses that to make suggestions about what a bot should say next."

The startup is currently piloting the platform with e-commerce and telecommunication companies, pharmaceutical firms, and other large enterprises. Clients can use the platform as a "brain" to power a Giant Otter chatbot or use the platform's conversation-authoring tools to power chatbots on third-party platforms, such as Amazon's Lex or IBM's Watson. The platform automates both text and voice conversations.

Benefits come in the form of cost savings. Major companies can spend billions of dollars on sales and customer support services; automating even a fraction of that work can save millions of dollars, Orkin says. Consumers, of course, will benefit from smarter bots that can more quickly and easily resolve their issues.

Human-machine collaboration

In conversation, people tend to express the same intent with different words, potentially over several sentences, and in various word orders. Unlike other chatbot-building platforms, Giant Otter uses "human-machine collaboration," Orkin says, "to learn authentic variation in the way people express different thoughts, and to do it bottom-up from real examples."

Giant Otter's algorithms comb through anywhere from 50 to 100 transcripts from sales and customer support conversations, identifying language variations of the same intent, such as "How can I help you?" and "What's your concern?" and "How may I assist you?" These are called "utterances." All utterances are mixed around into chunks of test scripts for people to judge for accuracy online.

Consider a script for a sales call, where a salesperson is selling a product while the prospect is pushing for a discount. Giant Otter's algorithms match and substitute one utterance in one script with a similar one from another script—such as swapping "I may be able to offer a discount"

with "I'll see if I can reach your price point." That version is uploaded to Mechanical Turk or another crowdsourcing platform, where people will vote a "yes" or "no" if the substituted sentence makes sense.

In another human task, people break conversations into "events." Giant Otter will lay out conversations horizontally and people will label different sections of the conversation. A salesperson saying, "Hello, thanks for contacting us," for instance, may be labeled as "call opening." Other section labels include "clarifying order," "verifying customer information," "proposing resolution," and "resolving issue."

"Between these two tasks, we learn a lot about the structure over how conversations unfold," Orkin says. "Conversations break down into events, events break down into utterances, and utterances break down into many different examples of saying the same thing with different words."

This builds a robust language database for chatbots to recognize anywhere from a few to more than 100 different ways to express the same sentiment—including fairly abstract variations. This is important, Orkin says, as today's chatbots are built top-down, by a human manually plugging in various utterances. But someone seeking an order status update could say, for instance, "My order hasn't come, and I checked my account, and it said to contact customer support."

"Nowhere does the person even say 'status.' If I was creating content for a bot by hand, there's no way I would have thought of that," Orkin says. The platform continues to learn and evolve after chatbots are deployed.

Path to chatbots

Giant Otter's origin goes back to 2005, when Orkin joined the Cognitive Machines Group led by Deb Roy, an associate professor of media arts

and sciences. Roy had just initiated the Human Speechome Project, his effort to gather data on how humans develop language by video and audio recording his newborn for three years.

Branching off from the project, Orkin developed The Restaurant Game, an online game that paired people online to have natural, text-based discussions as a customer and a waitress at a restaurant. "We hoped to get maybe 100 people to play. But not too long after [the game launched], we had data from 16,000 people," Orkin says.

Many people would order food, pay bills, and talk about the menu. Others, however, were more unorthodox, asking the waitress on a date, stealing the cash register, or stacking cakes up to climb onto the roof. All of that data was valuable. "Whatever they did, we had all that data of natural conversations between players," Orkin says. Soon, he hired people around the world through crowd-labor platforms, such as Mechanical Turk and Upwork, to ascribe context to the game's transcripts.

The game turned into a platform to collect and curate dialog to make AI conversations more natural. In 2013, Orkin met Geoff Marietta, a Harvard Graduate School of Education student studying how virtual worlds could facilitate learning and improve relationships. Using Orkin's platform, they developed a game called SchoolLife, where players assume roles of a bullying victim and a bystander. Players interacted with AI-controlled characters to come up with a solution to student conflict.

To commercialize the game, the two co-founded Giant Otter in 2013. For early support, Orkin turned to the MIT Venture Mentoring Service. Introductions through VMS have since led to a valuable partnership and potential pilot customers.

SchoolLife earned a Small Business Innovation Research grant from the National Science Foundation and was used in numerous schools in the region. But the budgeting cycles of school districts made it difficult for a startup to thrive in the education sector. Moreover, simulating bullying posed an issue, Orkin says. "With our platform, you need to capture data reflecting how people naturally converse. With bullying, we did come up with ways to record conversations from people role-playing in simulated scenarios, but it wasn't authentic data," Orkin says.

A third co-founder, Dan Tomaschko MBA '15, soon recognized the potential to impact the corporate world and guided a pivot to a training tool, called Coach Otto, that trained employees to deal with sensitive scenarios in the workplace. But last fall the startup realized the niche market of chatbot development for customer support was more profitable. (Companies, however, can still use the platform to practice phone conversations for sales and support training.)

Currently, Giant Otter is working on better integrating its [conversation](#)-authoring tools with third-party chatbot platforms. It's also developing its own automated customer support chatbot, powered by any company's call transcripts. "It's taken us years to realize where the most value is, but we're focused on ... [having] the right assembly line to crank enterprise phone call and live chat transcript data through our platform and turn it into something that can automate chatbot conversations," Orkin says.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Developing bots that talk more like people (2018, January 24) retrieved 17 July 2024 from <https://phys.org/news/2018-01-bots-people.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.