Wikipedia used to give AI context clues

Date:
September 19, 2017
Source:
Brigham Young University

Summary:
A team of computer scientists is teaching artificial intelligence agents how to interact with the world in a way that makes sense.

Share:
FULL STORY

Walk into a room, see a chair, and your brain will tell you that you can sit in it, tip it over or lift it up, but you wouldn't even consider drinking it, promoting it or unlocking it. As humans, explains Brigham Young University computer science professor David Wingate, we know intuitively that certain verbs pair naturally with certain nouns, and we also know that most verbs don't make sense when paired with random nouns.

"Consider the monitor on your desk: you can look at it, you can turn it on, you can even pick it up or throw it, but you cannot impeach it, transpose it, justify it or correct it," said Wingate. "You can dethrone a king or worship him or obey him, but you cannot unlock him or calendar him or harvest him."

That intuition, for the most part, doesn't exist with computer artificial intelligence agents, who are good at identifying objects but less so in knowing what to do with them. So Wingate and three student researchers, including lead author Nancy Fulda, developed a method for teaching agents about affordances -- the set of actions that can be done with an object. They recently presented their work at the International Joint Conference on Artificial Intelligence.

The team's ambitious end goal is to help build androids that can walk around the world and interact with it intelligently. Such an android "has incredible potential to do good, to help people," said Fulda, who is finishing her Ph.D. in computer science. An example she gives is elderly care: a robot who is told to "get me my glasses" could figure out what glasses look like, where they're likely to be, how heavy they are, how best to lift them and how to get them to the person requesting them.

As it stands right now, explains BYU computer science undergrad and research co-author Ben Murdoch, there are plenty of artificial intelligence agents who can identify what they're looking at, but they can't go beyond the ID: they might know they're looking at a phone but don't necessarily know what a phone is good for.

"When machine learning researchers turn robots or artificially intelligent agents loose in unstructured environments, they try all kinds of crazy stuff," said Murdoch. "The common-sense understanding of what you can do with objects is utterly missing, and we end up with robots who will spend thousands of hours trying to eat the table."

Because the hand-coding needed to help an agent understand which verbs make sense with which nouns would be a laborious, slow-moving process, BYU's research team instead found a way to put linear algebra and Wikipedia to use. Understanding that Wikipedia offers a vast corpus of mostly up-to-date language use, they downloaded it, ran it through a previously established algorithm that looks at words in their contexts, "and voila! The computer is equipped with common-sense knowledge about things that make sense," said Wingate, who recently received a National Science Foundation Career Award to help fund his artificial intelligence work.

For this project, the team tested their method in a series of text-based adventure games, which allow a player and agent to have back and forth text interactions, with the agent offering a situation and the player responding with a written phrase. Their method improved the computer's performance on 12 out of 16 games.

Even with help from Wikipedia, the team's agent makes mistakes with its language use. But he's made progress: unlike he tried to do in one of their early games, "he doesn't bulldoze Santa," said Fulda. "I love it when the agent does something that surprises me in a good way. I'm like, it did it: it figured it out all by itself. It's kind of like watching your child take their first steps."

There's plenty of work still to be done before reaching the team's end goal of having a functioning android, said co-author and BYU computer science master's student Daniel Ricks. "But it's really exciting to see the progress we've made."