

AI can read! Tech firms race to smarten up thinking machines

January 23 2018, by Matt O'brien



Stanford doctoral student Pranav Rajpurkar, who helped develop the Stanford Question Answering Dataset, poses for photos in an office of The Gates Computer Science Building in Stanford, Calif., Tuesday, Jan. 23, 2018. Seven years after a computer beat two human quizmasters on a "Jeopardy!" challenge, the tech industry has pumped resources into training its machines to get even better at amassing knowledge and answering questions. (AP Photo/Jeff Chiu)

Seven years ago, a computer beat two human quizmasters on a "Jeopardy" challenge. Ever since, the tech industry has been training its



machines to make them even better at amassing knowledge and answering questions.

And it's worked, at least up to a point. Just don't expect artificial intelligence to spit out a literary analysis of Leo Tolstoy's "War and Peace" any time soon.

Research teams at Microsoft and Chinese tech company Alibaba reached what they described as a milestone earlier this month when their AI systems outperformed the estimated human score on a reading comprehension test. It was the latest demonstration of rapid advances that have improved search engines and voice assistants and that are finding broader applications in health care and other fields.

The answers they got wrong—and the test itself—also highlight the limitations of computer intelligence and the difficulty of comparing it directly to human intelligence.

ERROR! ERROR!

"We are still a long way from computers being able to read and comprehend general text in the same way that humans can," said Kevin Scott, Microsoft's chief technology officer, in a LinkedIn post that also commended the achievement by the company's Beijing-based researchers.

The test developed at Stanford University demonstrated that, in at least some circumstances, computers can beat humans at quickly "reading" hundreds of Wikipedia entries and coming up with accurate answers to questions about Genghis Khan's reign or the Apollo space program.





In this Thursday, Jan. 13, 2011, file photo, "Jeopardy!" contest Ken Jennings, who won a record 74 consecutive games, refers to his opponent, an IBM computer called "Watson," while being interviewed after a practice round of the quiz show in Yorktown Heights, N.Y. Seven years after Watson beat two human quizmasters on a "Jeopardy!" challenge, the tech industry has pumped resources into training its machines to get even better at amassing knowledge and answering questions. (AP Photo/Seth Wenig, File)

The computers, however, also made mistakes that many people wouldn't have.

Microsoft, for instance, fumbled an easy football question about which member of the NFL's Carolina Panthers got the most interceptions in the 2015 season (the correct answer was Kurt Coleman, not Josh Norman). A person's careful reading of the Wikipedia passage would have discovered the right answer, but the computer tripped up on the word "most" and didn't understand that seven is bigger than four.



"You need some very simple reasoning here, but the machine cannot get it," said Jianfeng Gao, of Microsoft's AI research division.

HUMAN VS. MACHINE



Stanford doctoral student Pranav Rajpurkar, who helped develop the Stanford Question Answering Dataset, poses for photos outside of The Gates Computer Science Building in Stanford, Calif., Tuesday, Jan. 23, 2018. Seven years after a computer beat two human quizmasters on a "Jeopardy!" challenge, the tech industry has pumped resources into training its machines to get even better at amassing knowledge and answering questions. (AP Photo/Jeff Chiu)

It's not uncommon for machine-learning competitions to pit the cognitive abilities of computers against humans. Machines first bested people in an image-recognition competition in 2015 and a speech recognition competition last year, although they're still easily tricked.



Computers have also vanquished humans at chess, Pac-Man and the strategy game Go.

And since IBM's "Jeopardy" victory in 2011, the tech industry has shifted its efforts to data-intensive methods that seek to not just find factoids, but better comprehend the meaning of multi-sentence passages.

Like the other tests, the Stanford Question Answering Dataset, nicknamed Squad, attracted a rivalry among research institutions and tech firms—with Google, Facebook, Tencent, Samsung and Salesforce also giving it a try.

"Academics love competitions," said Pranav Rajpurkar, the Stanford doctoral student who helped develop the test. "All these companies and institutions are trying to establish themselves as the leader in AI."





Stanford doctoral student Pranav Rajpurkar, who helped develop the Stanford Question Answering Dataset, poses for photos in an office of The Gates Computer Science Building in Stanford, Calif., Tuesday, Jan. 23, 2018. Seven years after a computer beat two human quizmasters on a "Jeopardy!" challenge, the tech industry has pumped resources into training its machines to get even better at amassing knowledge and answering questions. (AP Photo/Jeff Chiu)

LIMITS OF UNDERSTANDING

The tech industry's collection and digitization of huge troves of data,



combined with new sets of algorithms and more powerful computing, has helped inject new energy into a machine-learning field that's been around for more than half a century. But computers are still "far off" from truly understanding what they're reading, said Michael Littman, a Brown University computer science professor who has tasked computers to solve crossword puzzles.

Computers are getting better at the statistical intuition that allows them to scan text and find what seems relevant, but they still struggle with the logical reasoning that comes naturally to people. (And they are often hopeless when it comes to deciphering the subtle wink-and-nod trickery of a clever puzzle.) Many of the common ways of measuring artificial intelligence are in some ways teaching to the test, Littman said.

"It strikes me for the kind of problem that they're solving that it's not possible to do better than people, because people are defining what's correct," Littman said of the Stanford benchmark. "The impressive thing here is they met human performance, not that they've exceeded it."





Stanford doctoral student Pranav Rajpurkar, who helped develop the Stanford Question Answering Dataset, poses for photos outside of The Gates Computer Science Building in Stanford, Calif., Tuesday, Jan. 23, 2018. Seven years after a computer beat two human quizmasters on a "Jeopardy!" challenge, the tech industry has pumped resources into training its machines to get even better at amassing knowledge and answering questions. (AP Photo/Jeff Chiu)

© 2018 The Associated Press. All rights reserved.



Citation: AI can read! Tech firms race to smarten up thinking machines (2018, January 23) retrieved 21 May 2024 from <u>https://phys.org/news/2018-01-ai-tech-firms-smarten-machines.html</u>

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.