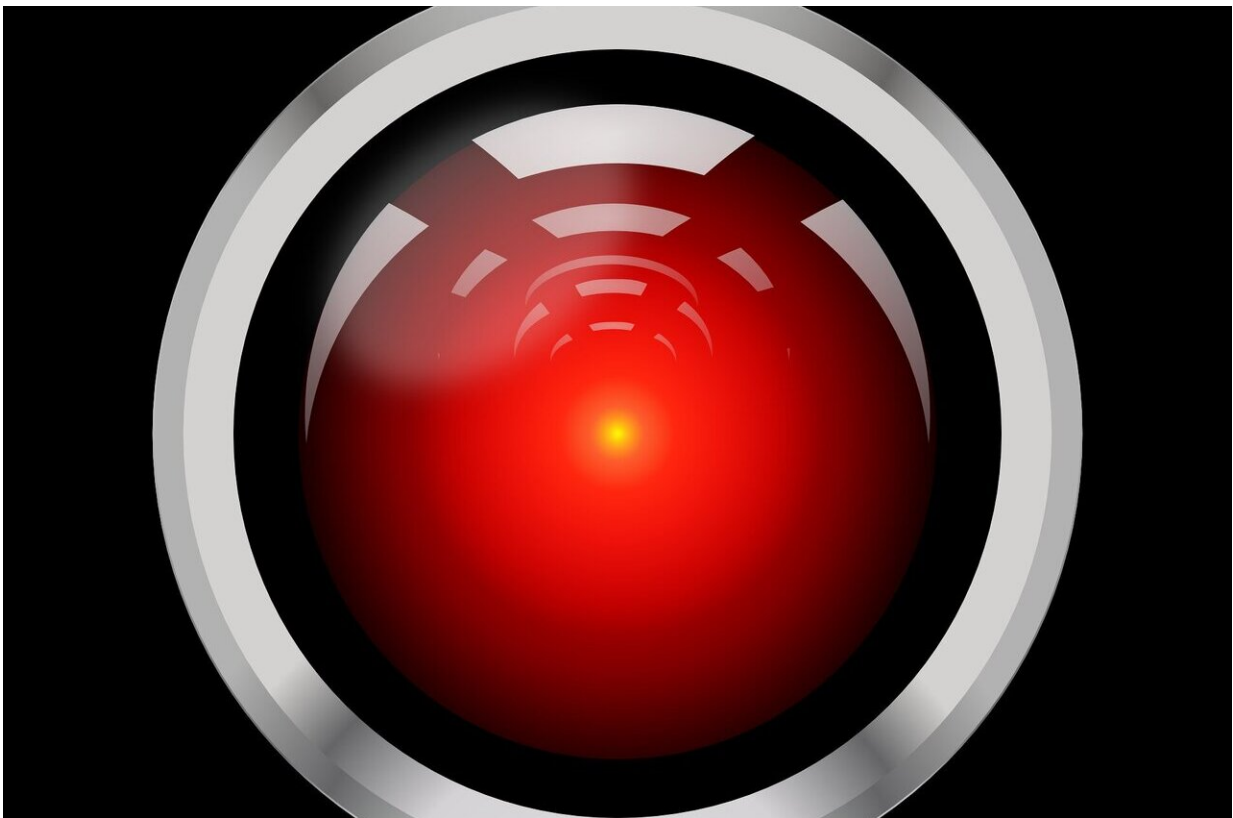# New report assesses progress and risks of artificial intelligence

September 16 2021

Credit: Pixabay/CC0 Public Domain

Artificial intelligence has reached a critical turning point in its evolution, according to a new report by an international panel of experts assessing the state of the field.

Substantial advances in language processing, computer vision and pattern recognition mean that AI is touching people's lives on a daily basis—from helping people to choose a movie to aiding in medical diagnoses. With that success, however, comes a renewed urgency to understand and mitigate the risks and downsides of AI-driven systems, such as algorithmic discrimination or use of AI for deliberate deception. Computer scientists must work with experts in the social sciences and law to assure that the pitfalls of AI are minimized.

Those conclusions are from a report titled "Gathering Strength, Gathering Storms: The One Hundred Year Study on Artificial Intelligence (AI100) 2021 Study Panel Report," which was compiled by a panel of experts from computer science, public policy, psychology, sociology and other disciplines. AI100 is an ongoing project hosted by the Stanford University Institute for Human-Centered Artificial Intelligence that aims to monitor the progress of AI and guide its future development. This new report, the second to be released by the AI100 project, assesses developments in AI between 2016 and 2021.

"In the past five years, AI has made the leap from something that mostly happens in research labs or other highly controlled settings to something that's out in society affecting people's lives," said Michael Littman, a professor of computer science at Brown University who chaired the report panel. "That's really exciting, because this technology is doing some amazing things that we could only dream about five or 10 years ago. But at the same time, the field is coming to grips with the societal impact of this technology, and I think the next frontier is thinking about ways we can get the benefits from AI while minimizing the risks."

The report, released on Thursday, Sept. 16, is structured to answer a set of 14 questions probing critical areas of AI development. The questions were developed by the AI100 standing committee consisting of a renowned group of AI leaders. The committee then assembled a panel of

17 researchers and experts to answer them. The questions include "What are the most important advances in AI?" and "What are the most inspiring open grand challenges?" Other questions address the major risks and dangers of AI, its effects on society, its public perception and the future of the field.

"While many reports have been written about the impact of AI over the past several years, the AI100 reports are unique in that they are both written by AI insiders—experts who create AI algorithms or study their influence on society as their main professional activity—and that they are part of an ongoing, longitudinal, century-long study," said Peter Stone, a professor of computer science at the University of Texas at Austin, executive director of Sony AI America and chair of the AI100 standing committee. "The 2021 report is critical to this longitudinal aspect of AI100 in that it links closely with the 2016 report by commenting on what's changed in the intervening five years. It also provides a wonderful template for future study panels to emulate by answering a set of questions that we expect future study panels to reevaluate at five-year intervals."

Eric Horvitz, chief scientific officer at Microsoft and co-founder of the One Hundred Year Study on AI, praised the work of the study panel.

"I'm impressed with the insights shared by the diverse panel of AI experts on this milestone report," Horvitz said. "The 2021 report does a great job of describing where AI is today and where things are going, including an assessment of the frontiers of our current understandings and guidance on key opportunities and challenges ahead on the influences of AI on people and society."

In terms of AI advances, the panel noted substantial progress across subfields of AI, including speech and language processing, computer vision and other areas. Much of this progress has been driven by

advances in machine learning techniques, particularly deep learning systems, which have made the leap in recent years from the academic setting to everyday applications.

In the area of natural language processing, for example, AI-driven systems are now able to not only recognize words, but understand how they're used grammatically and how meanings can change in different contexts. That has enabled better web search, predictive text apps, chatbots and more. Some of these systems are now capable of producing original text that is difficult to distinguish from human-produced text.

Elsewhere, AI systems are diagnosing cancers and other conditions with accuracy that rivals trained pathologists. Research techniques using AI have produced new insights into the human genome and have sped the discovery of new pharmaceuticals. And while the long-promised self-driving cars are not yet in widespread use, AI-based driver-assist systems like lane-departure warnings and adaptive cruise control are standard equipment on most new cars.

Some recent AI progress may be overlooked by observers outside the field, but actually reflect dramatic strides in the underlying AI technologies, Littman says. One relatable example is the use of background images in video conferences, which became a ubiquitous part of many people's work-from-home lives during the COVID-19 pandemic.

"To put you in front of a background image, the system has to distinguish you from the stuff behind you—which is not easy to do just from an assemblage of pixels," Littman said. "Being able to understand an image well enough to distinguish foreground from background is something that maybe could happen in the lab five years ago, but certainly wasn't something that could happen on everybody's computer, in real time and at high frame rates. It's a pretty striking advance."

As for the risks and dangers of AI, the panel does not envision a dystopian scenario in which super-intelligent machines take over the world. The real dangers of AI are a bit more subtle, but are no less concerning.

Some of the dangers cited in the report stem from deliberate misuse of AI—deepfake images and video used to spread misinformation or harm people's reputations, or online bots used to manipulate public discourse and opinion. Other dangers stem from "an aura of neutrality and impartiality associated with AI decision-making in some corners of the public consciousness, resulting in systems being accepted as objective even though they may be the result of biased historical decisions or even blatant discrimination," the panel writes. This is a particular concern in areas like law enforcement, where crime prediction systems have been shown to adversely affect communities of color, or in health care, where embedded racial bias in insurance algorithms can affect people's access to appropriate care.

As the use of AI increases, these kinds of problems are likely to become more widespread. The good news, Littman says, is that the field is taking these dangers seriously and actively seeking input from experts in psychology, public policy and other fields to explore ways of mitigating them. The makeup of the panel that produced the report reflects the widening perspective coming to the field, Littman says.

"The panel consists of almost half social scientists and half computer science people, and I was very pleasantly surprised at how deep the knowledge about AI is among the social scientists," Littman said. "We now have people who do work in a wide variety of different areas who are rightly considered AI experts. That's a positive trend."

Moving forward, the panel concludes that governments, academia and industry will need to play expanded roles in making sure AI evolves to

serve the greater good.

More information: Report: [ai100.stanford.edu/sites/g/fil …](ai100.stanford.edu/sites/g/fil) I100Report_MT_10.pdf

Provided by Brown University