The promise and pitfalls of artificial intelligence explored at TEDxMIT event

MIT scientists discuss the future of AI with applications across many sectors, as a tool that can be both beneficial and harmful.

Steve Nadis | MIT CSAIL
January 11, 2022

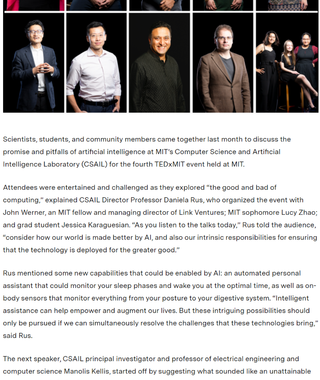

Top row, left to right: CSAIL D
Professor Manolis Kellis, Visit
Craig Watkins, Professor Mar
Assistant Professor Cathy W
right: Associate Professor Jir
Scientist Tian Gu, Associate
Raskar, Professor Aleksande
members of MIT Ohms, MIT's
cappella group.

Photos: Katherine Taylor

Scientists, students, and community members came together last month to discuss the promise and pitfalls of artificial intelligence at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) for the fourth TEDxMIT event held at MIT.

Attendees were entertained and challenged as they explored "the good and bad of computing," explained CSAIL Director Professor Daniela Rus, who organized the event with John Werner, an MIT fellow and managing director of Link Ventures; MIT sophomore Lucy Zhao; and grad student Jessica Karaguesian. "As you listen to the talks today," Rus told the audience, "consider how our world is made better by AI, and also our intrinsic responsibilities for ensuring that the technology is deployed for the greater good."

Rus mentioned some new capabilities that could be enabled by AI: an automated personal assistant that could monitor your sleep phases and wake you at the optimal time, as well as on-body sensors that monitor everything from your posture to your digestive system. "Intelligent assistance can help empower and augment our lives. But these intriguing possibilities should only be pursued if we can simultaneously resolve the challenges that these technologies bring," said Rus.

The next speaker, CSAIL principal investigator and professor of electrical engineering and computer science Manolis Kellis, started off by suggesting what sounded like an unattainable goal — using AI to "put an end to evolution as we know it." Looking at it from a computer science perspective, he said, what we call evolution is basically a brute force search. "You're just exploring all of the search space, creating billions of copies of every one of your programs, and just letting them fight against each other. This is just brutal. And it's also completely slow. It took us billions of years to get here." Might it be possible, he asked, to speed up evolution and make it less messy?

The answer, Kellis said, is that we can do better, and that we're already doing better: "We're not killing people like Sparta used to, throwing the weaklings off the mountain. We are truly saving diversity."

Knowledge, moreover, is now being widely shared, passed on "horizontally" through accessible information sources, he noted, rather than "vertically," from parent to offspring. "I would like to argue that competition in the human species has been replaced by collaboration. Despite having a fixed cognitive hardware, we have software upgrades that are enabled by culture, by the 20 years that our children spend in school to fill their brains with everything that humanity has learned, regardless of which family came up with it. This is the secret of our great acceleration" — the fact that human advancement in recent centuries has vastly out-clipped evolution's sluggish pace.

The next step, Kellis said, is to harness insights about evolution in order to combat an individual's genetic susceptibility to disease. "Our current approach is simply insufficient," he added. "We're treating manifestations of disease, not the causes of disease." A key element in his lab's ambitious strategy to transform medicine is to identify "the causal pathways through which genetic predisposition manifests. It's only by understanding these pathways that we can truly manipulate disease causation and reverse the disease circuitry."

Kellis was followed by Aleksander Madry, MIT professor of electrical engineering and computer science and CSAIL principal investigator, who told the crowd, "progress in AI is happening, and it's happening fast." Computer programs can routinely beat humans in games like chess, poker, and Go. So should we be worried about AI surpassing humans?

Madry, for one, is not afraid — or at least not yet. And some of that reassurance stems from research that has led him to the following conclusion: Despite its considerable success, AI, especially in the form of machine learning, is lazy. "Think about being lazy as this kind of smart student who doesn't really want to study for an exam. Instead, what he does is just study all the past years' exams and just look for patterns. Instead of trying to actually learn, he just tries to pass the test. And this is exactly the same way in which current AI is lazy."

A machine-learning model might recognize grazing sheep, for instance, simply by picking out pictures that have green grass in them. If a model is trained to identify fish from photos of anglers proudly displaying their catches, Madry explained, "the model figures out that if there's a human holding something in the picture, I will just classify it as a fish." The consequences can be more serious for an AI model intended to pick out malignant tumors. If the model is trained on images containing rulers that indicate the size of tumors, the model may end up selecting only those photos that have rulers in them.

This leads to Madry's biggest concerns about AI in its present form. "AI is beating us now," he noted. "But the way it does it [involves] a little bit of cheating." He fears that we will apply AI "in some way in which this mismatch between what the model actually does versus what we think it does will have some catastrophic consequences." People relying on AI, especially in potentially life-or-death situations, need to be much more mindful of its current limitations, Madry cautioned.

There were 10 speakers altogether, and the last to take the stage was MIT associate professor of electrical engineering and computer science and CSAIL principal investigator Marzyeh Ghassemi, who laid out her vision for how AI could best contribute to general health and well-being. But in order for that to happen, its models must be trained on accurate, diverse, and unbiased medical data.

It's important to focus on the data, Ghassemi stressed, because these models are learning from us. "Since our data is human-generated … a neural network is learning how to practice from a doctor. But doctors are human, and humans make mistakes. And if a human makes a mistake, and we train an AI from that, the AI will, too. Garbage in, garbage out. But it's not like the garbage is distributed equally."

She pointed out that many subgroups receive worse care from medical practitioners, and members of these subgroups die from certain conditions at disproportionately high rates. This is an area, Ghassemi said, "where AI can actually help. This is something we can fix." Her group is developing machine-learning models that are robust, private, and fair. What's holding them back is neither algorithms nor GPUs. It's data. Once we collect reliable data from diverse sources, Ghassemi added, we might start reaping the benefits that AI can bring to the realm of health care.

In addition to CSAIL speakers, there were talks from members across MIT's Institute for Data, Systems, and Society; the MIT Mobility Initiative; the MIT Media Lab; and the SENSEable City Lab.

The proceedings concluded on that hopeful note. Rus and Werner then thanked everyone for coming. "Please continue to reflect about the good and bad of computing," Rus urged. "And we look forward to seeing you back here in May for the next TEDxMIT event."

The exact theme of the spring 2022 gathering will have something to do with "superpowers." But — if December's mind-bending presentations were any indication — the May offering is almost certain to give its attendees plenty to think about. And maybe provide the inspiration for a startup or two.

SHARE

RELATED LINKS

01. Video: "TEDxMIT Sizzle Reel - CSAIL 2021"
02. TEDxMIT
03. Computer Science and Artificial Intelligence Laboratory (CSAIL)
04. Institute for Data, Systems, and Society
05. Department of Electrical Engineering and Computer Science
06. School of Engineering
07. MIT Schwarzman College of Computing

RELATED TOPICS

Special events and guest speakers · Artificial intelligence · Electrical Engineering & Computer Science (eecs) · Computer Science and Artificial Intelligence Laboratory (CSAIL) · IDSS · Media Lab · Data · Machine learning · Computational biology · Computer science and technology · Ethics · Technology and society · Evolution · Medicine · Disease · Health care · Health sciences and technology · Faculty · Alumni/ae · Diversity and inclusion · History of science · MIT Schwarzman College of Computing · School of Engineering · School of Architecture and Planning

RELATED ARTICLES

Machine-learning system flags remedies that might do more harm than good

CSAIL hosts first-ever TEDxMIT

Computing the future

More MIT News

Q&A: A conversation with Helen Elaine Lee about her novel, "Pomegranate"
The MIT professor's new book explores the world of a woman set free from prison and redefining herself in society.

Researchers teach an AI to write better chart captions
A new dataset can help scientists develop automatic systems that generate richer, more descriptive captions for online charts.

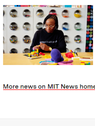

Transatlantic connections make the difference for MIT Portugal
The international partnership focuses on climate and sustainability.

Summer 2023 recommended reading from MIT
Enjoy these recent titles from Institute faculty and staff.

Studies at the intersection of equity, computing, and education
"The work I'm doing is deeply rooted in the belief that you can plant seeds in people," says graduate student Cecilé Sadler.

Computer vision system marries image recognition and generation
MAGE merges the two key tasks of image generation and recognition, typically trained separately, into a single system.

More news on MIT News homepage →