



ARTIFICIAL INTELLIGENCE

The man who helped invent virtual assistants thinks they're doomed without a new AI approach

Boris Katz has spent his career trying to help machines master language. He believes that current AI techniques aren't enough to make Siri or Alexa truly smart.

By Will Knight

March 13, 2019



MS. TECH

Siri, Alexa, Google Home—technology that parses language is increasingly finding its way into everyday life.

Boris Katz, a principal research scientist at MIT, isn't that impressed. Over the past 40 years, Katz has made key contributions to the linguistic abilities of machines. In the 1980s, he developed START, a system capable of responding to naturally phrased queries. The ideas used in START helped IBM's Watson win on *Jeopardy!* and laid the groundwork for today's chattering artificial servants.

But Katz now worries that the field suffers from a reliance on decades-old ideas, and that these ideas won't give us machines with real intelligence. I met with him to discuss the current limits of AI assistants and to hear his thoughts on where research needs to go if they're ever going to get smarter.

How did you become interested in making computers use language?

I first encountered computers in the 1960s as an undergraduate student at Moscow University. The particular machine I used was a mainframe called BESM-4. One could only use octal code to communicate with it. My first computer project involved teaching a computer to read, understand, and solve math problems.

Then I developed a poetry-writing computer program. I still remember standing in the machine room waiting to see the next poem generated by the machine. I was stunned by the beauty of the poems; they appeared to be produced by an intelligent entity. And I knew then and there that I want to work for the rest of my life on creating intelligent machines and finding ways to communicate with them.

What do you make of Siri, Alexa, and other personal assistants?

It's funny to talk about, because on the one hand, we are very proud of this incredible progress —everybody in their pocket has something that we helped create here many, many years ago, which is wonderful.



MIT

But on the other hand, these programs are so incredibly stupid. So there's a feeling of being proud and being almost embarrassed. You launch something that people feel is intelligent, but it's not even close.

There's been significant progress in AI thanks to machine learning. Isn't that making machines better at language?

On the one hand there is this dramatic progress, and then some of this progress is inflated. If you look at machine-learning advances, all the ideas came 20 to 25 years ago. It's just that eventually engineers did a great job of making these ideas a reality. This technology, as great as it is, will not solve the problem of real understanding—of real intelligence.

It seems like we are making progress in AI, though ... (see "10 Breakthrough Technologies: Smooth-Talking Personal Assistants")?

At a very high level, modern techniques—statistical techniques like machine learning and deep learning—are very good at finding regularities. And because humans usually produce the same sentences much of the time, it's very easy to find them in language.

Look at predictive text. The machine knows better than you what you are going to say. You could call that intelligent, but it's just counting words and numbers. Because we keep saying the same thing, it's easy to build systems that capture the regularities and act as if they are intelligent. This is the fictitious nature of much of the current progress.

What about the “dangerous” language-generating tool announced recently by OpenAI?

These examples are very impressive indeed, but I am not sure what they teach us. The OpenAI language model was trained on 8 million web pages in order to predict the next word, given all of the previous words within some text (which was on the same topic as the one the model was trained on). This huge amount of training certainly ensured local coherence (syntactic and even semantic) of the text.

Why do you think AI is headed the wrong way in language?

In language processing, like in other fields, progress was made by training models on huge amounts of data—many millions of sentences. But the human brain would not be able to learn language using this paradigm. We don't leave our babies with an encyclopedia in the crib, expecting them to master the language.

When we see something, we describe it in language; when we hear someone talk about something, we imagine what the described objects and events look like in the world. Humans live in a physical environment, filled with visual, tactile, and linguistic sensory inputs, and the redundant and complementary nature of these inputs makes it possible for human children to make sense out of the world, and to learn language at the same time. Perhaps by studying these modalities in isolation, we have made the problem harder rather than easier?

Why is common sense important?

Say your robot is helping you pack, and you tell it: “This book would not fit in the red box because *it* is too *small*.” Clearly, you want your robot to understand that the *redbox* is too small, so that you can continue to have a meaningful conversation. However, if you tell the robot: “This book would not fit in the red box because *it* is too *big*,” you want your robot to understand that the *book* is too big.

Knowing what entity in a conversation a pronoun refers to is a very common task that humans do every day, and yet, as you could see from these and other examples, it often relies on deep

understanding of the world, which is currently beyond the reach of our machines:
understanding of common sense and intuitive physics, understanding of beliefs and intentions of others, ability to visualize and reason about cause and effect, and much more.

You are trying to teach machines about language using simulated physical worlds. Why is that?

I have yet to see a baby whose parents put an encyclopedia in the crib and say, “Go learn.” And this is what our computers do today. I don’t think these systems will learn the way we want them to or understand the world the way we want to.

What happens with babies is they get tactile experience immediately of the world. Then babies start seeing the world and absorbing events and objects’ properties. And then the baby eventually hears linguistic input. And it’s this complementary input that makes the magic of understanding happen.

What is a better approach?

One way forward is to gain a greater understanding of human intelligence and then use that understanding in order to create intelligent machines. AI research needs to build on ideas from developmental psychology, cognitive science, and neuroscience, and AI models ought to reflect what is already known about how humans learn and understand the world.

Real progress will come only when researchers get out of our offices and start talking to people in other fields. Together we will come closer to understanding intelligence and figuring out how to replicate it in intelligent machines that can speak, see, and operate in our physical world.

The challenge of creating truly intelligent machines is a very difficult one, but it is also one of the most important challenges we have. **T**

by Will Knight

DEEP DIVE

ARTIFICIAL INTELLIGENCE

How a top Chinese AI model overcame US sanctions

With a new reasoning model that matches the performance of ChatGPT o1, DeepSeek managed to turn restrictions into innovation.

By Caiwei Chen

Everyone in AI is talking about Manus. We put it to the test.

The new general AI agent from China had some system crashes and server overload—but it's highly intuitive and shows real promise for the future of AI helpers.

By Caiwei Chen

The second wave of AI coding is here

A string of startups are racing to build models that can produce better and better software. They claim it's the shortest path to AGI.

By Will Douglas Heaven

OpenAI launches Operator—an agent that can use a computer for you

The announcement confirms one of two rumors that circled the internet this week. The other was about superintelligence.

By Will Douglas Heaven

STAY CONNECTED

Illustration by Rose Wong

**Get the latest updates from
MIT Technology Review**

Discover special offers, top stories, upcoming events, and more.

Enter your email

 →

[Privacy Policy](#)

The latest iteration of a legacy

Founded at the Massachusetts Institute of Technology in 1899, MIT Technology Review is a world-renowned, independent media company whose insight, analysis, reviews, interviews and live events explain the newest technologies and their commercial, social and political impact.

[READ ABOUT OUR HISTORY](#)

Advertise with MIT Technology Review

Elevate your brand to the forefront of conversation around emerging technologies that are radically transforming business. From event sponsorships to custom content to visually arresting video storytelling, advertising with MIT Technology Review creates opportunities for your brand to resonate with an unmatched audience of technology and business elite.

ADVERTISE WITH US

About us

Careers

Custom content

Advertise with us

International Editions

Republishing

MIT Alumni News

Help & FAQ

My subscription

Editorial guidelines

Privacy policy

Terms of Service

Write for us

Contact us



© 2025 MIT Technology Review

