Science Forward: Artificial Intelligence

A production of CUNY Advance and Macaulay Honors College
http://cuny.is/scienceforward

Flora Lichtman: [0:00] Humans have a longstanding fascination with seeing our creations come to life. We've been at it for millennia. Remember the ancient myth of Pygmalion, the sculptor so enchanted with the woman he chisels that goddess Aphrodite brings the sculpture to life?

[0:15] We never seem to tire of the concept. It might be because these stories raise such fundamental questions about what it means to be alive and what it takes to think and to feel.

(Movie Voice): [0:26] It's alive!

Flora Lichtman: [0:28] But these questions aren't reserved for fiction anymore. We're facing them in real life through the field of artificial intelligence.

[0:36] In 1770 an Austrian inventor premiered his so-called mechanical Turk, a supposedly intelligent automaton dressed in a European conception of Turkish clothing. The automaton played chess, and for decades the machine delighted audience throughout Europe and the Americas, beating the majority of its challengers, including Benjamin Franklin and Napoleon Bonaparte.

[0:59] The secret was simple. The intelligence was inside the machine. Hiding inside the cabinet, which people could inspect, was a person. A living, breathing chess master who operated the automaton and defeated the opponents.

[1:13] In the last 250 years we've made a lot of progress. Robots can amaze without a person hiding inside. But the people creating artificial intelligence still face deep questions about what intelligence means and how we know and created it.

[1:28] In this video we're exploring what artificial intelligence, or AI, is. We're talking about the differences between how computers and people process information, and we will see some AI in action.

[1:40] [music playing]

Suzanne Tamang: [1:49] AI is the demonstration of humanlike intelligence in machines. It was first coined in the '50s, and people continued to debate what artificial intelligence is and is not, both among the general public and among computer scientists.

Flora Lichtman: [2:08] Now machines really can play chess.

Andrew Rosenberg: [2:11] Chess for a long time was considered a pinnacle and marker of intelligence because it's highly reasoned. Only people can do it. To have a system that in this very constrained rules-based world can best human intelligence and can demonstrate what we considered to be a pinnacle of intelligence forced us in a lot of ways to rethink what intelligence is, and then what artificial intelligence is within that.

Susan Epstein: [2:42] Being smart is knowledge, but it's also knowing how to use knowledge. That combination is what we look for.

Flora Lichtman: [2:50] One way that scientists have attempted to define artificial intelligence is through the Turing test, named for its creator pioneering British computer scientist Alan Turing, who began his work in the 1940s.

Massimo Pigliucci: [3:03] The basic idea is this. I am talking to you right now, and the reason I think, according to Turing, the reason I think that I'm talking to a human being and not a computer is because I get into a conversation and you give me answers that are meaningful in response to my questions, and vice-versa.

[3:21] The basic idea therefore is that if we were having this conversation not face-to-face but using a computer terminal, and I didn't know whether on the other side of the terminal there was a human being or a computer program, if the conversation keeps going and I still can't tell the difference, and I find out that it is in fact a computer software, a piece of software that is on the other side of my terminal, then I would have to agree that that piece of software is intelligent, that that computer is intelligent.

[3:50] I think this is flawed for a variety of reasons. Number one, because you can actually fool human beings very easily.

Flora Lichtman: [3:57] People disagree about using the Turing test as a marker of intelligence. Susan Epstein provides a different framework.

Susan Epstein: [4:04] The way I teach about artificial intelligence, I talk about something we call an agent. An agent is simply a device or an object or a being that does three things over and over again. It senses its world, perceives it. It makes some sort of decision based on those perceptions. Then it takes an action based on that decision. Then it starts all over. Sense, decide, act. Sense, decide, act. Over and over again.

[4:37] That loop can be implemented a lot of different ways. People are really good at this because we have a lot of ways to perceive the world. We can see, we can smell, we can taste, we can hear, we can touch. All those abilities are what you would think of for a computer as input devices. They're ways of getting information in.

[5:03] It appears, at least to us, that we seamlessly integrate all this data that's coming in and we just know what's going on. If you speak to a computer, it's quite different from if you feel it text. There is all this noise in a signal, as you probably know.

Flora Lichtman: [5:21] We often hear scientists talk about signal versus noise. This is a data analysis skill that scientists use when interpreting observations and results. It sounds technical, but in fact we do these calculations every day. Think of what it's like to be in a large classroom as you try and focus on your teacher's lesson. All around you there's noise.

[5:40] Now, it doesn't have to be just sound. There are visual distractions too. Any irrelevant information counts as noise. You try to filter out all this extraneous information around you and focus on the signal, your teacher's lecture. It's tough even for humans, and we practice this skill every day of our lives.

[5:57] Now imagine trying to teach a computer how to do this. How do you teach a computer what to pay attention to and what to ignore?

[6:04] In our robotic heroes and villains we like to imagine sophisticated, autonomous machines with lots of skills and their own personalities. But here's what they look like in real life.

[6:15] Flora Lichtman: Go to the kitchen.

Are you asking me to go to room YK95043

Flora Lichtman: [6:23] Yes.

Manuela Veloso: [6:25] The goal was to really have it completely autonomous. Nobody following it, nobody checking. Just letting it go down the corridors by itself and just by itself.

[6:38] [robot noises]

Manuela Veloso: [6:45] That was a major thing for us, because a lot of these experiments have been done in the past, and ours is still being done, but people follow this just to make sure. How do we go from not following the robot to just let it go? What is this difference?

[7:00] It happened that no matter what when we sent it it always got stuck somewhere. It was either a chair that was not on its map or a person that hit the robot, or God knows what. It became a little bit of a frustration to understand that this robot was always never going to be able to be by itself.

[7:19] At one point I realized that this is it. We still are going to have to deploy this thing. I conceived this concept of symbiotic autonomy, in which the robot asks for help.

[7:32] CoBot stands for Collaborative Robot. It became collaborative because it's not really capable of doing the whole thing by itself. It doesn't understand everything you say up front. But it asks, keeps asking, and eventually converges to knowing more and more by asking. Actually they also go to the Internet, so if I say, "Bring me coffee," it's going to say, "What is this coffee thing about?" It goes to the web. So it has these external resources.

[7:57] It talks with other robots in the building. We have four of them at Carnegie Melon. It's just changed the paradigm, this concept of not having to do it all by yourself. It's still autonomous, but it can ask for help. That's why we feel so strongly about this "co" part, CoBot, because it's collaborative. With humans, with the web, and with other robots.

Flora Lichtman: [8:20] As CoBots move around the building doing tasks, delivering things, escorting people, they can also collect data about the environment. Like the strength of the wireless signal in different areas of the office, or temperature in the office throughout the day.

Manuela Veloso 5: [8:33] For a person it's very hard to actually acquire these detail maps. You have this machine that tediously also goes over and over and then creates probabilistic averages, does not take one measurement, and knows it all. This robot has moved four kilometers in the building over at different times, many times, many samples. It goes and it goes and it goes. It brings all that knowledge from tedious repetition and accurate localization in the building that then humans can use to make decisions about a location of resources, about...

Flora Lichtman: [9:11] Where to put another router.

Manuela Veloso: [9:12] Where to put another router.

Flora Lichtman: [9:14] CoBots can understand speech, but how do you get a machine not only to understand words and sentences but also subtext? That is, how the meaning changes based on how the words are spoken.

Man 3: [9:27] Currently speech recognition does a pretty good job of recognizing the words you're saying, but generally ignores a lot of the inflection and intonation which carries a lot of meaning for a human being, but unfortunately machines are currently not making particularly good use of.

Bon Sy: [9:43] Our natural language is inherently ambiguous, such as if we say something like, "I see a boy playing in the park." That sentence itself is inherently ambiguous, that which is not clear. Are we referring to I am in the park, or are we referring the boy is playing in the park?

[10:07] There are interesting problems that we get into. If we try to get the machine to actually exhibit human intelligent behavior, then what we would actually like the machine to be able to do is not just about cranking numbers, solving a specific problem, but actually try to go solve problems that we may not know very well how to solve at all ourselves.

Flora Lichtman: [10:29] Humans are pretty good at sorting through messy data, separating signal from noise, and combining lots of different data streams to get a contextualized picture of the world. Researchers say if we want to replicate some of those skills and machines, one approach is to start by better understanding ourselves.

Woman 5: [10:47] First of all, we actually don't know very well how we do it ourselves.

Woman 6: [10:53] Knowledge representation is a really essential piece of AI. It's how we describe our world, which is so different to a computer, in a way that it can capture it.

Flora Lichtman: [11:04] Hi CoBot. Go to the kitchen.

CoBot: [11:13] You are asking me to go to room YK1943. Is that correct?

Flora Lichtman: [11:19] Yes.

[11:21] [robot rolling away]

CoBot: [11:29] I am done with my task. Please press the done button.

[11:33] [music playing]