Researchers discover machines can learn by simply observing

August 30 2016

Credit: Piotr Siedlecki/Public Domain

It is now possible for machines to learn how natural or artificial systems work by simply observing them, without being told what to look for, according to researchers at the University of Sheffield.

This could mean advances in the world of technology with machines able to predict, among other things, human behaviour.

The discovery takes inspiration from the work of pioneering computer scientist Alan Turing, who proposed a test, which a machine could pass if it behaved indistinguishably from a human. In this test, an interrogator exchanges messages with two players in a different room: one human, the other a machine.

The interrogator has to find out which of the two players is human. If they consistently fail to do so - meaning that they are no more successful than if they had chosen one player at random - the machine has passed the test, and is considered to have human-level intelligence.

Dr Roderich Gross from the Department of Automatic Control and Systems Engineering at the University of Sheffield, said: "Our study uses the Turing test to reveal how a given system - not necessarily a human - works. In our case, we put a swarm of robots under surveillance and wanted to find out which rules caused their movements. To do so, we put a second swarm - made of learning robots - under surveillance too. The movements of all the robots were recorded, and the motion data shown to interrogators."

He added: "Unlike in the original Turing test, however, our interrogators are not human but rather computer programs that learn by themselves. Their task is to distinguish between robots from either swarm. They are rewarded for correctly categorising the motion data from the original swarm as genuine, and those from the other swarm as counterfeit. The learning robots that succeed in fooling an interrogator - making it believe their motion data were genuine - receive a reward."

Dr Gross explained the advantage of the approach, called 'Turing Learning', is that humans no longer need to tell machines what to look for.

"Imagine you want a robot to paint like Picasso. Conventional machine

[learning](#) algorithms would rate the robot's paintings for how closely they resembled a Picasso. But someone would have to tell the algorithms what is considered similar to a Picasso to begin with. Turing Learning does not require such prior knowledge. It would simply reward the robot if it painted something that was considered genuine by the interrogators. Turing Learning would simultaneously learn how to interrogate and how to paint."

Dr Gross said he believed Turing Learning could lead to advances in science and technology.

"Scientists could use it to discover the rules governing natural or artificial systems, especially where behaviour cannot be easily characterised using similarity metrics," he said.

"Computer games, for example, could gain in realism as virtual players could observe and assume characteristic traits of their human counterparts. They would not simply copy the observed behaviour, but rather reveal what makes human players distinctive from the rest."

The discovery could also be used to create algorithms that detect abnormalities in behaviour. This could prove useful for the health monitoring of livestock and for the preventive maintenance of machines, cars and airplanes.

Turing Learning could also be used in security applications, such as for lie detection or online identity verification.

So far, Dr Gross and his team have tested Turing Learning in robot swarms but the next step is to reveal the workings of some animal collectives such as schools of fish or colonies of bees. This could lead to a better understanding of what factors influence the behaviour of these animals, and eventually inform policy for their protection.