

Why AI can't ever reach its full potential without a physical body

October 6 2020, by Mark Lee



Credit: Alexander James Spence, Author provided

Artificial intelligence seems to be making enormous advances. It has become the key technology behind self-driving cars, [automatic translation systems](#), speech and textual analysis, image processing and all kinds of diagnosis and recognition systems. In many cases, AI can [surpass the best](#) human performance levels at specific tasks.

We are witnessing the emergence of a new commercial industry with

intense activity, massive financial investment, and tremendous potential. It would seem that there are no areas that are beyond improvement by AI—no tasks that cannot be automated, no problems that can't at least be helped by an AI application. But is this strictly true?

Theoretical studies of computation have shown there are some things that are not computable. Alan Turing, the brilliant mathematician and code breaker, proved that some computations might never finish (while others would take years or even centuries).

For example, we can easily compute a few moves ahead in a game of chess, but to examine all the moves to the end of a typical 80-move chess game is completely impractical. Even using one of the world's fastest supercomputers, running at over one hundred thousand trillion operations per second, it would take over a year to get just a tiny portion of the chess space explored. This is also known as the [scaling-up problem](#).

Early AI research often produced good results on small numbers of combinations of a problem (like noughts and crosses, known as toy problems) but would not scale up to larger ones like chess (real-life problems). Fortunately, modern AI has developed alternative ways of dealing with such problems. These can beat the world's best human players, not by looking at all possible moves ahead, but by looking a lot further than the human mind can manage. It does this by using methods involving approximations, probability estimates, large neural networks and other machine-learning techniques.

But these are really problems of computer science, not artificial intelligence. Are there any fundamental limitations on AI performing intelligently? A serious issue becomes clear when we consider [human-computer interaction](#). It is widely expected that future AI systems will communicate with and assist humans in friendly, fully interactive, social

exchanges.

Theory of mind

Of course, we already have primitive versions of such systems. But audio-command systems and call-center-style script-processing just pretend to be conversations. What is needed are proper social interactions, involving free-flowing conversations over the long term during which AI systems remember the person and their past conversations. AI will have to understand intentions and beliefs and the meaning of what people are saying.

This requires what is known in psychology as a [theory of mind](#) – an understanding that the person you are engaged with has a way of thinking, and roughly sees the world in the same way as you do. So when someone talks about their experiences, you can identify and appreciate what they describe and how it relates to yourself, giving meaning to their comments.

We also observe the person's actions and infer their intentions and preferences from gestures and signals. So when Sally says, "I think that John likes Zoe but thinks that Zoe finds him unsuitable," we know that Sally has a first-order model of herself (her own thoughts), a second-order model of John's thoughts, and a third-order model of what John thinks Zoe thinks. Notice that we need to have similar experiences of life to understand this.

Physical learning

It is clear that all this social interaction only makes sense to the parties involved if they have a "sense of self" and can similarly maintain a model of the self of the other agent. In order to understand someone

else, it is necessary to know oneself. An AI "self model" should include a subjective perspective, involving how its body operates (for example, its visual viewpoint depends upon the physical location of its eyes), a detailed map of its own space, and a repertoire of well understood skills and actions.

That means a physical body is required in order to ground the sense of self in concrete data and experience. When an action by one agent is observed by another, it can be mutually understood through the shared components of experience. This means social AI will need to be realized in robots with bodies. How could a software box have a subjective viewpoint of, and in, the physical world, the world that humans inhabit? Our conversational systems must be not just embedded but embodied.

A designer can't effectively build a software sense-of-self for a robot. If a subjective viewpoint were designed in from the outset, it would be the designer's own viewpoint, and it would also need to learn and cope with experiences unknown to the designer. So what we need to design is a framework that supports the learning of a subjective viewpoint.

Fortunately, there is a way out of these difficulties. Humans face exactly the same problems but they don't solve them all at once. The first years of infancy display [incredible developmental progress](#), during which we learn how to control our bodies and how to perceive and experience objects, agents and environments. We also learn how to act and the consequences of acts and interactions.

Research in the new field of developmental robotics is now exploring how robots can learn from scratch, like infants. The first stages involve discovering the properties of passive objects and the "physics" of the robot's world. Later on, robots note and copy [interactions with agents](#) (carers), followed by gradually more complex modeling of the self in context. In [my new book](#), I explore the experiments in this field.

So while disembodied AI definitely has a fundamental limitation, future research with robot bodies may one day help create lasting, empathetic, social interactions between AI and humans.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: Why AI can't ever reach its full potential without a physical body (2020, October 6) retrieved 30 April 2024 from <https://techxplore.com/news/2020-10-ai-full-potential-physical-body.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--