


Neuroscientist explains differences between AI and human learning

June 6 2022, by James Fodor



Credit: AI-generated image ([disclaimer](#))

Recently developed artificial intelligence (AI) models are capable of many impressive feats, including recognizing images and producing human-like language. But just because AI can perform human-like behaviors doesn't mean it can think or understand like humans.

As a researcher studying how humans understand and reason about the world, I think it's important to emphasize the way AI systems "think" and learn is fundamentally different to how humans do—and we have a long way to go before AI can truly think like us.

A widespread misconception

Developments in AI have produced systems that can perform very human-like behaviors. The language model [GPT-3](#) can produce text that's often indistinguishable from [human speech](#). Another model, [PaLM](#), can produce explanations for jokes it has never [seen before](#).

Most recently, a general-purpose AI known as Gato has been developed which can [perform hundreds of tasks](#), including captioning images, answering questions, playing Atari video games, and even controlling a robot arm to stack blocks. And DALL-E is a system which has been trained to produce modified images and artwork from a text description.

These breakthroughs have led to some bold claims about the capability of such AI, and what it can tell us about [human intelligence](#).

For example Nando de Freitas, a researcher at Google's AI company DeepMind, argues scaling up existing models will be enough to produce human-level [artificial intelligence](#). Others have [echoed](#) this view.

In all the excitement, it's easy to assume human-like behavior means human-like understanding. But there are several key differences between how AI and humans think and learn.

Neural nets vs the human brain

Most recent AI is built from [artificial neural networks](#), or "neural nets"

for short. The term "neural" is used because these networks are inspired by the human brain, in which billions of cells called neurons form complex webs of connections with one another, processing information as they fire signals back and forth.

Neural nets are a highly simplified version of the biology. A real neuron is replaced with a simple node, and the strength of the connection between nodes is represented by a single number called a "weight."

With enough connected nodes stacked into enough layers, neural nets can be trained to recognize patterns and even "[generalize](#)" to stimuli that are similar (but not identical) to what they've seen before. Simply, generalization refers to an AI system's ability to take what it has learnt from certain data and apply it to new data.

Being able to identify features, recognize patterns, and generalize from results lies at the heart of the success of neural nets—and mimics techniques humans use for such tasks. Yet there are important differences.

Neural nets are typically trained by "[supervised learning](#)". So they're presented with many examples of an input and the desired output, and then gradually the connection weights are adjusted until the network "learns" to produce the desired output.

To learn a language task, a neural net may be presented with a sentence one word at a time, and will slowly learn to predict the next word in the sequence.

This is very different from how humans typically learn. Most human learning is "unsupervised," which means we're not explicitly told what the "right" response is for a given stimulus. We have to work this out ourselves.

For instance, children aren't given instructions on how to speak, but learn this through a [complex process](#) of exposure to adult speech, imitation, and feedback.

Another difference is the sheer scale of data used to train AI. The GPT-3 model was trained on [400 billion words](#), mostly taken from the internet. At a rate of 150 words per minute, it would take a human nearly 4,000 years to read this much text.

Such calculations show humans can't possibly learn the same way AI does. We have to make more efficient use of smaller amounts of data.

Neural nets can learn in ways we can't

An even more fundamental difference concerns the way neural nets learn. In order to match up a stimulus with a desired response, neural nets use an algorithm called "backpropagation" to pass errors backward through the network, allowing the weights to be adjusted in just the right way.

However, it's widely recognized by neuroscientists that [backpropagation can't be implemented](#) in the brain, as it would require [external signals](#) that just don't exist.

Some researchers have proposed [variations](#) of backpropagation could be used by the brain, but so far there is no evidence human brains can use such learning methods.

Instead, humans learn by making [structured mental concepts](#), in which many different properties and associations are linked together. For instance, our concept of "banana" includes its shape, the color yellow, knowledge of it being a fruit, how to hold it, and so forth.

As far as we know, AI systems do not form conceptual knowledge like this. They rely entirely on extracting complex statistical associations from their training data, and then applying these to similar contexts.

Efforts are underway to build AI that [combines different types of input](#) (such as images and text)—but it remains to be seen if this will be sufficient for these models to learn the same types of rich mental representations humans use to understand the world.

There's still much we don't know about how humans learn, understand and reason. However, what we do know indicates humans perform these tasks very differently to AI systems.

As such, [many researchers believe](#) we'll need new approaches, and more fundamental insight into how the [human brain](#) works, before we can build machines that truly think and learn like humans.

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

Provided by The Conversation

Citation: Neuroscientist explains differences between AI and human learning (2022, June 6) retrieved 3 May 2024 from

<https://techxplore.com/news/2022-06-neuroscientist-differences-ai-human.html>

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.