

CS50 AI and Machine Learning

Overview

From self-driving cars to virtual assistants like Siri and Alexa, artificial intelligence (AI) is becoming more and more common in our everyday lives. The field of AI is centered around the development of **machine intelligence**, or the ability of machines to think and react like humans. But why would we want to program computers and robots to think like us? Well, it turns out that sometimes they can make decisions and answer questions even better than we can. AI has already had a major impact on many parts of our world, including banking, technology, marketing, entertainment, and healthcare, and will likely soon change many more!

Key Terms

- machine intelligence
- Turing test
- machine learning
- training

Turing Test

Computer science pioneer and mathematician, Alan Turing, also played an important role in the field of AI. In his 1950 paper, "Computing Machinery and Intelligence," Turing asked the monumental and ever-relevant question, "Can machines think?" This sparked many conversations that influenced the field. Furthermore, he proposed an "imitation game," now known as the **Turing test**, to determine whether a machine exhibits human-like intelligent behavior to the point where the two are indistinguishable. It works by having a "judge" guess, based only on their answers to the same set of questions, which is the computer and which the human.

The test's efficacy has been widely debated. In recent years, many scientists have claimed to have passed the test with their programs, while still others hold that no programs have been able to do this so far. This is a pretty incredible feat since the first artificial intelligence program, the Logic Theorist, was presented at a conference in 1956!

Machine Learning

One method for achieving artificial intelligence is **machine learning**, a term which was coined in 1959 and is generally defined as the techniques we use to have programs learn from data. There are many applications for these programs, including computer vision and natural language processing. And with the rise of big data, this method is becoming ever more relevant.

Typically, the first step in machine learning is collecting and validating data. Then, a model is set up and trained with part of the data. This **training**, or learning, can be supervised, which means the data set is labeled, or unsupervised, where the data is unlabeled. After training, the model can be further tested and improved, and its accuracy can be measured. Finally, we can use our model to make judgements or predictions about other data.

Deep learning is a subset of machine learning which uses artificial neural networks (ANNs). This system of connected nodes is loosely based on the biological neural networks of animal brains. By using ANNs, some scientists think we can program machines to think more like us.

Future of AI

The AI that exists today is known as weak AI, or that which is designed to complete a specific task like classifying images or driving. The term strong AI refers to the point at which machines can think for themselves. We don't know when, if ever, we will reach the point where we have superintelligent machines. Still, the media is full of depictions of strong AI gone terrible.

In the meantime, we need to grapple with the issues emerging from the rise of AI. Consider, for instance, that we tell a self-driving car to get from point A to point B. It successfully completes this task but it may not be able to make ethical decisions like if it had to pick between hitting three pedestrians or one pedestrian (commonly known as the trolley problem). The machine did what we asked it to do, but not without serious consequences. Who should be held accountable?

Some scholars worry about superintelligence with goals that are not aligned with ours. But will artificially intelligent computers take over the world? We'll soon find out!