

Data versus lore: An introduction to the technology behind Artificial Intelligence

By MANDY HATHAWAY

Because of its simplicity, this article is not meant to be the whole of the learning experience and educational links and videos will be provided throughout.

What is Artificial Intelligence?

Television, movies, books, everywhere you turn these days the discussion and stories often center around Artificial Intelligence (AI). Because of the media attention and our vast imaginations, understanding what is real and what is science-fiction can be a challenge.

What is AI really like and what is fantasy, at least for now? This short series will explore that and much more. If you've ever wondered what AI is and how it works, then this article is for you.

Let's dive in and explore the different kinds of AI before taking a brief look at how these systems work.

Types of Artificial Intelligence

Narrow Intelligence: Current systems are something called Artificial Narrow Intelligence, or weak AI. These systems can outperform humans, but only in very particular ways such as speech recognition or data analysis. These systems can be used, sometimes in conjunction with one another, to do complex things. However, the systems are not thinking, they are following predesigned formulas and algorithms to achieve user-defined goals. And while these goals can vary widely, from Amazon and Netflix's personalized recommendations to medical imaging analysis, all of these systems are confined by their preprogrammed parameters. Narrow AI would include all current AI systems.

General Intelligence: AGI is what you see in the movies and read about in books. This term includes those robots that think, love, fear and dream of a better future. These are something called strong AI. These systems would be complex, they would be self-aware or have what is often called consciousness. As a result, they would have the ability to set their own goals, think and use reason to make their own decisions. These systems would be as good or better than humans at a number of things, but AGI has not been developed yet.

Super Intelligence: The next level of AI systems, ones that have far surpassed human abilities are often called Artificial Super Intelligence. These are the sci-fi machines that we imagine will take over the world or plot the destruction of humankind. These systems are similar to AGI but they are orders of magnitude more intelligent, they outperform humans in every way. These systems do not yet exist outside of the imagination and are the systems that are still the farthest away from being developed.

How do current AI systems work?

Machine Learning

Most current AI systems are based on something called machine learning. A classic way to conceptualize this is to think of dogs and cats. When you see either one, you generally know right away which animal it is. Though both are similar in size and shape, your lifelong exposure to these animals, or at least to images of these animals, allows you to easily recognize which animal you are viewing. That is the goal of machine learning, through a large volume of exposure systems, designers are teaching the AI to recognize patterns and statistical relationships.

Systems are fed information that is particular to what they are designed to do. Medical diagnostic software is trained on medical histories, diagnosis, and scan or test results. A chatbot or voice recognition system, such as those in Siri, Alexa and your cable company's phone system, are trained on huge databases of typed conversations or spoken audio respectively. No matter the job, machine learning employs data and statistics to formulate its output.

Neural Networks

One of the most common ways to design an Artificial Intelligence system is by using a neural network. These systems are thus named because the nodes connect to one another in a similar way to the connections between biological neurons and transmit small bits of data to make complex computations.

During the system's training, machine learning is used. Data is provided at the input layer, including all the variables and data that the system will be given to learn from. Information about the desired output is also often included. For example, images of dogs would be labeled 'dogs,' images of cats would be labeled 'cats.' The system runs the information from the input layer through its hidden layers and looks for statistical relationships that lead to the output. The system uses these relationships to create a custom algorithm that will be used to process similar data in the future.

The hidden layers do very different things depending upon the kind of data they are being trained with. Image analysis incorporates a value for each pixel of an image and uses statistics to predict what those pixels represent. Chatbots and their cousins, the virtual assistants, are trained using possible behaviors or responses and a database of possible voice or text inputs. The hidden layers identify aspects or sounds from the speech or text which they can match to their list of possible behaviors or responses.

The trainer then reviews the system's output and offers additional input data and corrections. This process is repeated until the output is reliable enough to satisfy the program designer. This can often mean that the output is exponentially more reliable than similar work done by humans.

Now, let's return to the previous analogy about identifying cats and dogs. When you see a small furry creature, your brain does a few things automatically.

First, you take in the available data without even thinking about it. You observe the animal's size, shape and location; its approximate size compared to its surroundings; and its body and face type. All of these things are valuable data for you to process and compare. This is much like the input layer of a neural network.

Next, the brain takes all of that data and compares it to other information and experiences that you have stored. There is more than one way, and no correct versus incorrect way, for your brain to do this. Perhaps the animal looks like your aunt's cat, perhaps it reminds you of a dog breed you have seen many times, or perhaps, while you can't see it well, the size and the fact that it is walking on top of a fence reminds you of the way you have seen cats, but never dogs, move.

What happens in the hidden layers of your brain or a neural network is only as correct or incorrect as the final output. In life, when you realize you have made a mistake, you learn from that by adding new data which will be used to do later comparisons. When a neural network gets incorrect results in training, the programmer adjusts the parameters or weights and the system adjusts its internal algorithm until it starts getting the answers within an acceptable threshold of error.

Intelligent does not mean thinking

It is crucial to remember that current systems do not have an understanding of what they are doing outside of their specific purpose. Some systems, such as IBM's Watson and Google's DeepMind, outperform doctors at cancer diagnosis through analysis of people's medical scans. These are very impressive results and demonstrate potentially life-changing systems.

However, in spite of his diagnostic prowess, Watson doesn't actually know what cancer is or have any conceptualization of what it means to be either alive or dead. It doesn't understand how cancer grows, doesn't know what caused it; Watson doesn't even understand the value of what it is doing, all it knows is how to expertly identify cancer in diagnostic imaging. It is possible that Watson could be taught to do some other things, but only if it was trained on additional data that allowed it to learn.

This is why these systems are called narrow artificial intelligence. They can be very good at what they do, but they cannot do or accomplish more than they are trained or programmed to achieve.

One final note, for those of you who were surprised to not see Sophia, the AI granted citizenship in Saudi Arabia in 2017, mentioned here. Sophia is not AGI, she is a narrow AI who appears to be a complex chatbot with a robotic case. This means that she isn't expressing her own thoughts or feelings, rather she has learned to mimic human conversation after being trained on a database of the same. While Saudi Arabia granted her full citizenship, this publicity stunt does not signal any real breakthrough or turning point in AI development.

Don't forget to look for the next article in this series, where we will explore current AI applications.