



SDS PODCAST

EPISODE 464

FIVE MINUTE

FRIDAY:

A.I. VS MACHINE

LEARNING VS DEEP

LEARNING



(00:05):

This is Five-Minute Friday, on A.I. versus Machine Learning versus Deep Learning.

(00:19):

One of the most confusing aspects of terminology in the data science world is the distinction between artificial intelligence, machine learning, and deep learning. The media and businesspeople generally use the three terms interchangeably even though they represent different concepts; even academic experts will sometimes use them inaccurately when they're not being careful.

(00:43):

Let's start with artificial intelligence: "A.I." is the buzziest, vaguest, and broadest of the three terms. Taking a stab at a technical definition regardless, a decent one is that AI involves a machine processing information from its surrounding environment and then factoring in that information to achieve some desired outcome. Perhaps given this, some consider the goal of AI to be the achievement of "general intelligence" — intelligence as it is generally referred to with respect to broad reasoning and problem-solving capabilities. If you want to learn more about this idea, you can check out SuperDataScience episode #438, I talk about that in a huge amount of detail.

(01:29):

Back to our main point here. In practice and particularly in the popular press, "AI" is used to describe any cutting-edge machine capability. Presently, these capabilities include voice recognition, describing what's happening in a video, question-answering, driving a car, industrial robots that mimic human actors in the factory, or dominating humans at "intuition-heavy" board games like Go. Those are the kinds of capabilities



that AI has today. But once an AI capability becomes common-place (e.g., recognizing handwritten digits, which was cutting-edge in the 1990s), the “AI” moniker is dropped by the popular press for that particular capability such that the goalposts on the definition of AI are always moving.

(02:23):

Ok, so now machine learning. Machine learning is actually a subset of AI alongside other facets of AI like robotics and approaches such as expert systems that are hard-coded so they don’t learn directly from data. Machine learning, in contrast, is a field of computer science concerned with setting up software in a manner so that the software can recognize patterns in data without the programmer needing to explicitly dictate how the software should carry out all aspects of that recognition. That said, the programmer would typically have some insight into or some hypothesis about how the problem might be solved, and would thereby provide a rough model framework and relevant data such that the learning software is well prepared and well equipped to solve the problem on hand.

(03:18):

Before I can dig into what deep learning is, I first need to introduce the term artificial neural networks (ANNs). Artificial neurons are simple algorithms, really quick ones, that are inspired by biological brain cells, especially in the sense that individual neurons—whether they are biological or artificial—receive input from many other neurons, perform some quick speed computation, and then produce a single output. An artificial neural network, then, is a collection of these artificial neurons arranged so that they send and receive information between each other. Data (e.g., images of cats and dogs) are fed into an ANN, which then processes these data in some way with the goal of producing some desired result (e.g., a guess as to whether the image is a cat or a dog). So an ANN is just an example of machine learning approach.



(04:24):

So now that we know what an artificial neural network is, deep learning is fairly straightforward to define: It's a machine learning approach that involves an ANN composed of at least a few separate layers of artificial neurons. We can call that a deep learning network. More specifically, deep learning networks have a total of five or more layers with the following structure: They have a single input layer that is reserved for the data being fed into the network. And then they have three or more hidden layers that can represent the inputs in increasingly complex, increasingly abstract ways as we add more and more of these hidden layers. And then finally, a single output layer that is reserved for the values (e.g., predictions) that the network outputs.

(05:19):

With each successive layer in the deep learning artificial neural network being able to represent increasingly abstract, nonlinear recombinations of the previous layers, deep learning models with fewer than a dozen layers of artificial neurons are often sufficient for learning to make accurate predictions with a given dataset. That said, there are deep learning networks out there with hundreds or even upwards of a thousand layers that have in occasional circumstances think largely academic been demonstrated to provide a big value.

(05:53):

As rapidly improving accuracy benchmarks and countless competition wins in the past decade have demonstrated, the deep learning approach to modeling data excels at a broad range of machine learning tasks. Indeed, with deep learning driving so much of the contemporary progress in AI capabilities, I think this is why we see the words “deep learning” and “artificial intelligence” used so interchangeably by the popular press — and even by experts who should know better!



(06:22):

Alright, so I hope you enjoyed that quick intro to A.I., machine learning, and deep learning. If you'd like to learn more about those three terms I recommend checking out my book, Deep Learning Illustrated. Much of this podcast content was inspired directly by content from Chapter 4.

(06:39):

A quick announcement that starting with the next episode of this podcast, Episode 465, we will begin releasing guest episodes on Tuesday mornings New York time. Historically, we've released Wednesday evenings, but by releasing 36 hours earlier, we'll be giving you two more morning commutes in your week to enjoy the episode. I haven't been able to imagine any downsides to this change, but I didn't want to catch you off guard when it happens.