MACHINE LEARNING, DATA SCIENCE, ARTIFICIAL INTELLIGENCE, DEEP LEARNING, AND STATISTICS

Machine learning. Data science. Artificial intelligence. Deep learning. Statistics. Most organizations, companies and individuals today are using these technologies – whether they know it or not. If your work involves computers, you're likely familiar with at least some of them – but the terms can be confusing, and their use sometimes conflicting.

The 21st century is the era of big data. Big data refers to data sets that are so large and complex that previous applications of data processing aren't adequate. Researchers and companies are harnessing and experimenting with various methods of extracting value from big data. The global connected world offers infinite ways to generate, collect, and store data for analysis. Never before have we had access to this much data, and we are only now beginning to find ways to unleash the immense amount of meaning and information contained within.

The relatively recent concepts of data science, machine learning, and deep learning offer a new set of techniques and methods, but also find their way into hype and branding. Companies may adopt these terms without necessarily using their processes for a "cutting-edge" appeal to customers. In this article, we'll explore the differences between these terms, whether they're new or a return of the old, and whether they're just different names for the same thing.

Statistics and artificial intelligence

Let's begin with statistics, as these fields have been around for decades, even centuries, before computers were invented. The study of statistics and the application of statistical modeling are a subfield of mathematics. Both the theories and applications are aimed at identifying and formalizing relationships in data variables, based on mathematical equations. Statistical modeling relies on tools like samples, populations, and hypotheses.

In the latter part of 20th century, as access to computers became more widely available and computational power commoditized, people began to do statistics in computational applications. This allowed for treatment of larger and different data sets as well as the application of statistical methods that were untenable without computing power.

Artificial Intelligence is ultimately an evolution of this first encounter between math and computer science. Statistical modeling started as a purely mathematical or scientific exercise, but when it became computational, the door opened to using statistics to solve 'human' problems. In the post-war, due to enthusiastic optimism around the promise of computing as well as the belief that human thought processes were essentially computational, the idea that we could build an 'artificial' human intelligence gained currency.

In the 1960s, the field of artificial intelligence was formalized into a subset of computer science. New technology and a more expansive understanding of how humans' minds work changed artificial intelligence, from the original computational statistics paradigm to the modern idea that machines could mimic actual human capabilities, such as decision making and performing more "human" tasks.

Modern artificial intelligence is often broken into two areas: general artificial intelligence and applied artificial intelligence. Applied artificial intelligence is at play when we consider systems like driverless cars or machines that can smartly trade stocks. Much less common in practice is general artificial intelligence, the concept that a system could, in theory, handle any task, such as:

- Planning
- Getting around
- Recognizing objects and sounds
- Speaking and translating
- Performing social or business transactions
- Working creatively

The concept of artificial intelligence grows and shifts as technology advances, and likely will do so for the foreseeable future. Currently the only solid criterion for success or failure is how it can accomplish applied tasks.

Machine learning

By 1959, the idea of artificial intelligence had gained solid traction in computer science. Arthur Samuel, a leader and expert in the field, imagined that instead of engineers "teaching" or programming computers to have what they need to carry out tasks, that perhaps computers could teach themselves – learn something without being explicitly programmed to do so. Samuel called this "machine learning".

Machine learning is a form of applied artificial intelligence, based on the theory that systems that can change actions and responses as they are exposed to more data will be more efficient, scalable and adaptable for certain applications compared to those explicitly programmed (by humans). There are certainly many current applications proving this point: navigation apps and recommendation engines (shopping, shows, etc.) being two of the obvious examples.

Machine learning is typically categorized as either 'supervised' or 'unsupervised'. Supervised learning involves the machine to infer functions from known inputs to known outputs. Unsupervised MACHINE LEARNING works with the inputs only, transforming or finding patterns in the data itself without a known or expected output. For a more detailed discussion, see my blog about the differences between supervised and unsupervised machine learning.

Machine learning is a task-oriented application of statistical transformations. Accomplishing the task will require a process or set of steps, rules, etc. The process or set of rules to be followed in calculations or problem-solving operations is called an algorithm. When designing a learning machine, the engineer programs a set of algorithms through which the machine will process data.

As the machine learns – gets feedback – it typically will not change the employed statistical transformations but rather alter the algorithm. For example, if the machine is trained to factor two criteria in evaluating data and it learns that a third criteria has high correlation to the other two and refines the accuracy of calculation, it could add that third criteria to the analysis. This would be a change to the steps (algorithm), but not the underlying math.

Ultimately, machine learning is a way to "teach" computers to be adaptable to changes in data. We now have essentially infinite amounts of digital data being created constantly. The volume and diversity of that data increases rapidly and exponentially. Machines analysis has the advantages of speed, accuracy and lack of bias over human analysis, which is why machine learning is critical and has hit a tipping point.

Deep learning

Deep learning goes even further than machine learning as applied ARTIFICIAL INTELLIGENCE – it could be considered the cutting edge, says industry expert Bernard Marr. Machine learning trains and works on large sets of finite data, e.g. all the cars made in the 2000s. Machine learning does a good job of learning from the 'known but new' but does not do well with the 'unknown and new'.

Where machine learning learns from input data to produce a desired output, deep learning is designed to learn from input data and apply to *other* data. A paradigmatic case of deep learning is image identification. Suppose you want a machine to look at an image and determine what it represents to the human eye. A face, flower, landscape, truck, building, etc. To do this, the machine would have to learn from thousands or millions of images and then apply that knowledge to each specific new image you want it to identify.

Machine learning is not sufficient for this task because machine learning can only produce an output from a data set - whether according to a known algorithm or based on the inherent structure of the data. You might be able to use machine learning to determine whether an image was of an "X" – a flower, say – and it would learn and get more accurate. But that output is binary (yes/no) and is dependent on the algorithm, not the data. In the image recognition case, the outcome is not binary and not dependent on the algorithm.

This is because deep learning uses neural networks. Neural networks require their own deeper dive

in another post but for our purposes here, we just need to understand that neural networks don't calculate like typical machines. Rather than following an algorithm, neural networks are designed to make many 'micro' calculations about data. Which calculations and in what order is determined by the data, not an algorithm. Neural networks also support weighting data for 'confidence'. This results in a system that is probabilistic, vs. deterministic, and can handle tasks that we think of as requiring more 'human-like' judgement.

Deep learning neural networks are large and complex, requiring many layers and distributions of micro calculations. The machine still trains on data, but it can perform more nuanced actions than machine learning. Deep learning is appropriate for machine classification tasks like facial, image, or handwriting recognition.

Here are interesting examples of current, real-world technology using machine learning and deep learning:

- Driver-less cars use sensors and onboard analytics that better recognize obstacles, so they can quickly and more accurately react appropriate.
- Software applications are able to recolor black and white images by recognizing objects and predicting the colors that humans see.
- Machines are able to predict the outcome of legal proceedings when basic case facts are input into the computer.

Data science

Statistics is a field of mathematics. Artificial intelligence, deep learning and machine learning all fit within the realm of computer science. Data science is a separate thing altogether.

Formally defined, data science is an interdisciplinary approach to data mining, which combines statistics, many fields of computer science, and scientific methods and processes in order to mine data in automated ways, without human interaction. Modern data science is increasingly concerned with big data.

Data science has many tools, techniques, and algorithms culled from these fields, plus others – in order to handle big data. The goal of data science, somewhat similar to machine learning, is to make accurate predictions and to automate and perform transactions in real time, such as purchasing internet traffic or automatically generating content.

Data science relies less on math and coding and more on data and building new systems to process the data. Relying on the fields of data integration, distributed architecture, automated machine learning, data visualization, data engineering, and automated data-driven decisions, data science can cover an entire spectrum of data processing, not only the algorithms or statistics related to data.

Terminology branding

These terms are sometimes used interchangeably, and sometimes even incorrectly. A company with a new technology to sell may talk about their innovative data science techniques, when really, they may be using nothing close to it. In this way, companies are simply aligning themselves with what the concepts stand for: innovation, forward-thinking, and newfound uses for technology and our data. This isn't inherently bad, it's simply a caution that because a company claims use of these tools in its product design doesn't mean it does. Caveat emptor.