DataCamp

The difference between AI and machine learning

by Dr. Hugo Bowne-Anderson

The business world is overloaded with buzz terms like artificial intelligence, machine learning, AI transformation, deep learning, and data science. We know that these fields, technologies, and tools are changing the competitive landscape across verticals and are soon to become more table stakes and foundational than disruptive. However, people can know they're important while not even knowing what they really mean. I'd forgive you for your confusion because these words are overloaded and they're not even used consistently. My goal here is to dispel any confusion by demystifying these terms: What are artificial intelligence, machine learning, and data science? Where do they intersect? Where do they diverge? Read on to find out.

TL;DR

Artificial Intelligence (AI) is a "a huge set of tools for making computers behave intelligently" and in an automated fashion. This includes voice assistants, recommendation systems, and self-driving cars.

Machine Learning (ML) is the "field of study that gives computers the ability to learn without being explicitly programmed." The lion's share of ML involves computers learning patterns from existing data and applying it to new data in the form of making predictions, such as predicting whether an email is spam or not, whether a customer will churn or not, and diagnosing a particular piece of medical imaging.

Data Science (DS) is about making discoveries and creating insights from data and communicating these insights and discoveries to non-technical stakeholders.

How are these related?

ML feeds into both AI and DS
If the output of your machine learning model is fed into a computational system that performs an action in an automated fashion, such as recommending a movie, decelerating a self-driving car, or serving search results, it can be viewed as a component in your AI system.

If the output of your machine learning model is fed into a human decision making process, it can be considered data science work (for example, when predicting a customer may churn results in a human deciding to incentivize the customer to stay, this insight or discovery informs a data science decision).

Much of ML and AI rely on high quality data, meaning the most impactful and effective AI strategies will stand on the shoulders of robust data science capabilities.

Artificial Intelligence

In his Coursera course AI for Everyone, Andrew Ng, co-founder of Google Brain and former Chief Scientist at Baidu, defines artificial intelligence as "a huge set of tools for making computers behave intelligently." This definition, although good, does cast a wide net and it's worth making clear that "behaving intelligently" means by providing several examples:

- Voice assistants, such as Siri
- Recommendation systems, such as Netflix
- Self-driving cars
- Drones that fly over fields and capture footage used to optimize crop yield
- Google Search
- Surfacing algorithms, such as those employed by Twitter and Facebook, that decide what content to show you in your feed

It's important to recognize that here, AI means that actions and decisions are automated. It's also key to note that all of these are examples of Artificial Narrow Intelligence (ANI), that is, algorithms that can do one thing well. This is not to be confused with Artificial General Intelligence (AGI), which is a hypothetical, futuristic AI that can do anything a human is capable of. Nor is it a Superintelligent AI, a hypothetical software agent whose intelligence surpasses that of humans. Both AGI and Superintelligent AI are a long way off, if at all possible, and serve as distractions for real, present, and necessary conversations around the capabilities and limitations of AI, as we know it today, resulting in headlines such as An AI god will emerge by 2042 and write its own bible. WIll you worship it? This is clearly absurd and distracts from all the current examples of AI that allow computers to perform tasks that mimic aspects of human intelligence, such as recognizing stop signs and people in images and videos (self-driving cars), holding basic conversations, retrieving information, and performing tasks (voice assistants), and ranking text documents based on their relevance to a particular query (Google Search).

Now if AI is a huge set of tools, what tools are we talking about? One tool of central importance to modern AI is machine learning (ML), which we'll now explore.

Machine Learning

ML powers recommendation systems, content discovery, search engines, email spam filters, and matching problems across the board in tech. In healthcare, it's being leveraged for drug discovery and high throughput diagnostic imaging diagnosis. In finance, ML is now foundational for fraud detection, process automation, algorithmic trading, and robo-advisory. In retail, Walmart is at the forefront of using ML to reinvent supply chain management. The list goes on. So what actually is machine learning?

Machine learning was a term popularized in 1959 by Arthur Samuel, a pioneer in artificial intelligence and computer gaming. Samuel defined machine learning as the **"field of study that gives computers the ability to learn without being explicitly programmed."**

The majority of ML involves computers learning patterns from existing data and then applying it to

- Predicting whether a given credit card transaction is fraudulent or not, given transaction details.
- Predicting whether an email is spam or not, given the email sender, subject, and body.
- Predicting the diagnosis of a particular piece of medical imaging.
- Predicting the present and future location of pedestrians, cars, and other stationary/ moving objects in a video feed (such as those used by self-driving cars).

These prediction and classification problems currently form the majority of ML, and are referred to as supervised learning. This terminology results from the fact that the label you're trying to predict, e.g. spam or not, are said to supervise the learning process.

It's essential to note that the power of modern ML rests firmly on having good quality data for your algorithm to learn from, or be "trained on," and that such "training data," as it's commonly called, needs to be labeled. In the spam classification example, you'll need many examples of emails labeled with whether they were spam or not; in the diagnostic imaging example, you'll require at least thousands of images labeled with their diagnosis.

Machine learning recycles intelligence

What your ML algorithms then do is pick up patterns in your “training data” and generalize those patterns to unlabeled data, where you don’t know the outcome that you’re trying to predict. It’s for this reason that mathematician and Stanford Professor David Donoho prefers the term “recycled intelligence” to the term “artificial intelligence” for machine learning, as no new intelligence is created, but human intelligence, as captured by humans with domain expertise hand-labelling datasets, is recycled and re-applied to new data. There is a huge and hidden supply chain behind the worlds of machine learning and artificial intelligence: people leverage services such as Amazon Mechanical Turk to crowd-source labeled data. And Scale AI, a start-up that works with tens of thousands of contractors worldwide to hand-label data, recently raised $100 million, which speaks to the scale of the challenge (the irony of a company that procures hand-labeled data being called Scale AI is not lost on me).

Predicting behavior with reinforcement learning

It’s worth mentioning that, although the vast majority of AI and ML rely firmly on labeled data and recycling intelligence contained therein, there is a growing sub-field of ML called reinforcement learning (RL) that relies far less, if at all, on pre-existing training data. In RL, which draws on behavioral psychology, software agents are placed in constrained environments and given “rewards” and “punishments” based on their activity. If playing games sounds like a relevant application of RL to you, you’re spot on: RL was how AlphaGo Zero became the world Go champion in 2017, beating AlphaGo, which was trained on human data. RL also has serious applications in the self-driving car world and is a different paradigm that we’ll definitely see more of. More recently, in 2019, Pluribus beat the best professional players in six-player no-limit Texas Hold’em poker.

While we’re talking about ML, let’s also demystify deep learning, which is a specific form of ML that has been receiving a great deal of attention for some time now, and for good reason.

Deep Learning

Deep learning is a form of machine learning that uses algorithms called neural networks, which are loosely inspired by biological neural networks in human brains. Be clear that this is the extent to which the analogy holds and we should not anthropomorphize deep learning systems to consider them human or conscious in any way. The majority of the applications of deep learning occur in the supervised learning world in the form of image classification (facial recognition, self-driving cars, drone footage utilized to estimate crop yield in AgTech) and natural language processing (Google translation, document classification, sentiment analysis), although there are other applications in time-series prediction, such as financial prediction problems.

It's important to emphasize that deep learning systems are rarely good at more than one task: an algorithm that is built for facial recognition will not be any good at classifying legal documents. This is to say that, although you may like to call deep learning a form of AI, it is so in the sense of narrow artificial intelligence, not artificial general intelligence, the realm of hypothesized computational systems that are as intelligent as humans across the board.

Is ML a form of AI?

So is ML a form of AI? Colloquially, yes, machine learning is regarded as a form of artificial intelligence, but if we're thinking of AI as a set of tools for making computers behave «intelligently», then ML becomes one of these tools. For example, the Google spam filter is an example of an AI, and the ML algorithm that classifies a given email as spam or not is one component of this AI (another component is the software that pushes emails classified as spam by the ML algorithm into your spam folder).

Now that we've got a handle on AI, ML, RL, and DL, let's see what data science is all about.

Data Science

In their seminal 2012 Harvard Business Review article [Data Scientist: The Sexiest Job of the 21st Century], Thomas Davenport and DJ Patil state unequivocally that “more than anything, what data scientists do is make discoveries while swimming in data.” Data science is about creating insights from data, often in a business setting. How do data scientists do this, though? [Wikipedia] provides a bit more clarity in stating that “data science is a multidisciplinary field that uses scientific methods, processes, algorithms, and systems to extract knowledge and insights from structured and unstructured data.” How does this play out in practice? There are so many tools and techniques in a modern data scientist’s toolbox that it’s helpful to partition the space. One way to slice the data science space is into the following categories:

- Descriptive analytics
- Predictive analytics
- Prescriptive analytics

Descriptive analytics is essentially about getting the right pre-existing data in front of the right people in the form of dashboards, reports, or emails (for example). This can include both past and real-time time data about revenue, customer engagement, churn, and company and employee performance.

Predictive analytics is synonymous with machine learning and is the realm of predicting the future, such as whether a customer will churn or not, and more general classification tasks: Is an email spam or not? Is a tumor in a diagnostic image benign or malignant?

Prescriptive analytics is the realm of decision science and how to make decisions based on data. If your machine learning model, for example, tells you that a particular customer will churn, this doesn’t tell you what to do about it—prescriptive analytics is concerned with finding frictionless interfaces between the data and decision functions in any given organization. Exciting spaces to watch are data translation, advances in reinforcement learning (which bleeds into machine learning, as we’ve seen), and the work of Cassie Kozyrkov, Chief Decision Scientist at Google Cloud, with whom I discussed decision science on DataFramed, the DataCamp podcast.

So wait, is machine learning part of data science or part of AI? And what is the relation between data science and AI? As discussed, these terms are used in a variety of inconsistent ways, but a good rule of thumb is that:

- If the output of your machine learning model is fed into a human decision making process, it can be considered data science work (for example, if predicting a customer may churn results in a human deciding to incentivize the customer to stay, this can be considered insights or discoveries made from the data).

- If the output of your machine learning model is fed into a computational system that performs an action in an automated fashion, such as recommending a movie, decelerating a self-driving car, or serving search results, it can be viewed as a component in your AI system.

The main distinction between AI and data science we see emerge here is that, although many of the tools, techniques, infrastructures, and processes are the same, data science is often fed into human decision-making processes while AI is concerned with automation. However, remember that much of ML and AI relies on high quality data. What this means is that the most impactful and effective AI strategies will stand on the shoulders of robust data science capabilities.