What is machine learning?

Machine learning is the science of training machines to analyze and learn from data the way humans do. Computer systems use machine learning algorithms to process large quantities of historical data and identify data patterns. They can then predict outcomes more accurately from a given input data set. For example, Data scientists could train a medical application to diagnose cancer from x-ray images by storing millions of scanned images and their corresponding diagnoses.

Is machine learning artificial intelligence?

No. While the terms Machine learning and Artificial Intelligence(AI) may be used interchangeably, they are not the same. Artificial Intelligence is an umbrella term for different strategies and techniques used to make machines more human-like. AI includes everything from smart assistants like Alexa to auto chatbots. Machine learning is one among many other branches of Artificial Intelligence. While machine learning is AI, all AI activities cannot be called machine learning.

Are machine learning and data science the same?

No. Data science is a field of study that uses a scientific approach to extract meaning and insights from data. Data scientists use a range of tools for data analysis, and machine learning is one such tool. Data scientists know the bigger picture around the data – the business model, domain, data collection, etc. Machine learning is a computational process that only deals with raw data.

Why is machine learning important?

Data is the critical driving force behind business decision-making. Traditionally, companies have been using data from various business aspects like customer feedback, employees, and finance to improve business operations. Machine learning research automates and optimizes this process, helping businesses deliver greater value in a much shorter

timeframe. Machine learning helps businesses by driving business growth, unlocking new revenue streams, and solving challenging problems.

Where is machine learning used?

Let's look at machine learning applications in some key industries:

Manufacturing

Machine learning can support predictive maintenance, quality control, and innovative research in the manufacturing sector. Machine learning technology also helped companies improve logistical solutions, including assets, supply chain, and inventory management. For example, manufacturing giant 3M uses AWS Machine Learning to innovate sandpaper. Machine learning algorithms enable 3M researchers to analyze how slight changes in shape, size, and orientation improve abrasiveness and durability. In turn, those suggestions inform the manufacturing process.

Healthcare and life sciences

The proliferation of wearable sensors and devices has generated a significant volume of health data. Machine learning programs can analyze this information and support doctors in real-time diagnosis and treatment. Machine learning researchers are significantly impacting human health by developing solutions that detect cancerous tumors and diagnose eye diseases. For example, Cambia health solutions used AWS Machine Learning to support healthcare start-ups. They were able to automate and customize treatment for pregnant women.

Financial services

Financial machine learning projects improve risk analytics and regulation. Machine learning technology can allow investors to identify new opportunities by analyzing stock market movements, evaluating hedge funds, or calibrating financial portfolios. In addition, it can help identify high-risk loan clients and mitigate signs of fraud. Financial software leader Intuit uses AWS Machine Learning system, Amazon Textract, to create more personalized financial management and help end users improve their financial health.

Retail

Retail can use machine learning to improve customer service, stock management, upselling and cross-channel marketing. For example, Amazon fulfillment (AFT) cut infrastructure costs by 40 percent using a machine learning model to identify misplaced inventory. This helped them deliver on Amazon's promise that an item will be readily available to a customer and arrive on time, despite processing millions of global shipments annually.

Media and Entertainment

Entertainment companies turn to machine learning to better understand their target audiences and deliver immersive, personalized, and on-demand content. Machine learning algorithms are deployed to help design trailers and other advertisements, provide consumers with personalized content recommendations, and even streamline production. For example, Disney is using AWS Deep Learning to archive their media library. AWS machine learning tools automatically tag, describe, and sort media content, enabling Disney writers and animators to quickly search for and familiarize themselves with Disney characters.

How does machine learning work?

The central idea behind machine learning is an existing mathematical relationship between any input and output data combination. The machine learning model does not know this relationship in advance, but it can guess if given sufficient data sets. Thus every machine learning algorithm is built around a modifiable math function. An overview of the underlying principle can be understood as follows:

1) We 'train' the algorithm by giving it the following input/output (i,o) combinations – (2,10), (5,19), and (9,31)

2) The algorithm computes the relationship between input and output to be: o=3*i+4

3) We then give it input 7 and ask it to predict the output. It can automatically determine the output as 25.

While this is a basic understanding, machine learning focuses on the principle that all complex data points can be mathematically linked by computer systems as long as they have sufficient data and computing power to process that data. Thus, the accuracy of the output is directly co-relational to the magnitude of the input given.

What are the types of machine learning algorithms?

Depending on the expected output and the input type, you can categorize algorithms by four distinct learning styles.

- Supervised machine learning
- Unsupervised machine learning
- Semi-supervised learning
- Reinforcement machine learning

Supervised machine learning

Data scientists supply algorithms with labeled and defined training data to assess for correlations. The sample data specifies both the input and the output of the algorithm. For example, images of handwritten figures are annotated to indicate which number they correspond to. Given sufficient examples, a supervised-learning system would recognize the clusters of pixels and shapes associated with each number. It would eventually recognize handwritten numbers, reliably distinguishing between the numbers 9 and 4 or 6 and 8.

The strengths of supervised learning are simplicity and ease of design. It is useful when predicting a possible limited set of outcomes, dividing data into categories, or combining results from two other machine learning algorithms. However, labeling millions of unlabeled data sets is challenging.

What is data labeling?

Data labeling is the process of categorizing input data with its corresponding defined output values. Labeled training data is required for supervised learning. For example, millions of apple and banana images would need to be tagged with the words “apple” or “banana.” Then machine learning applications can use this training data to guess the name of the fruit when given a fruit image.

However, labeling millions of new data can be a time-consuming and challenging task. This limitation of supervised learning algorithms can be overcome to some extent by using crowd-working services such as Amazon Mechanical Turk. These services provide access to a large pool of low-cost labor spread across the globe, making data acquisition less challenging.

Unsupervised machine learning

Unsupervised learning algorithms train on unlabeled data. They scan through new data, trying to establish meaningful connections between the inputs and predetermined outputs. They can spot patterns and categorize data. For example, unsupervised algorithms could group news articles from different news sites into common categories like sports, crime, etc. They can use natural language processing to comprehend meaning and emotion in the article. Or unsupervised learning in retail could find patterns in customer purchases. It could give data analysis results like the customer is most likely to purchase bread if also buying butter.

Unsupervised learning is useful for pattern recognition, anomaly detection, and automatically grouping data into categories. As the training data does not require labeling, set up is easy. These algorithms can be used to automatically clean and process data for further modeling as well. The limitation of this method is that it cannot give precise predictions. In addition, it cannot single out specific types of data outcomes independently.

Semi-supervised learning

As the name suggests, this method combines supervised and unsupervised learning. The technique relies upon using a small amount of labeled data and a large amount of unlabelled data to train systems. First, the labeled data is used to train the machine-learning algorithm partially. After that, the partially trained algorithm itself labels the unlabelled data. This process is called pseudo-labeling. The model is then re-trained on the resulting data mix without being explicitly programmed.

The advantage of this method is that you do not require large amounts of labeled data. It is handy when working with data like long documents that would be too time-consuming for people to read and label.

Reinforcement learning

Reinforcement learning is a method with reward values attached to the different steps that the algorithm is supposed to go through. So the model's goal is to accumulate as many reward points as possible and eventually reach an end goal. Most of the practical application of reinforcement learning in the past decade has been in the realm of video games. Cutting edge reinforcement learning algorithms have achieved impressive results in classic and modern games, often beating out their human counterparts significantly.

While this method works best in uncertain and complex data environments, it is rarely implemented in business contexts. It is not efficient for well-defined tasks, and developer bias can affect the outcomes. As the data scientist designs the rewards, she can influence the results.

Are machine learning models deterministic?

If a system's output is predictable, then it is said to be deterministic. Since most software applications always respond predictably to the user's action, you can say, "If the user does this, he gets that." However, machine learning algorithms learn through observation along with experiences. Therefore, they are probabilistic in nature. The statement now changes to – "If the user does this, there is an X% chance of that happening."

In machine learning, determinism is a strategy used while applying the learning methods described above. Any of the supervised, unsupervised and other training methods can be made deterministic depending on what outcomes are wanted by the business. The research question, data retrieval, structure, and storage decisions determine if a deterministic or non-deterministic strategy is adopted.

Deterministic Vs. Probabilistic approach

The deterministic approach focuses on the accuracy and the amount of data collected. Efficiency is prioritized over uncertainty. On the other hand, the non-deterministic (or probabilistic) process is designed to manage the chance factor. Built-in tools are integrated into machine learning algorithms to help quantify, identify and measure uncertainty during the learning and observation.

What is deep learning?

Deep learning is a type of machine learning technique that is modeled on the human brain. Deep learning algorithms analyze data with a logic structure similar to that used by humans. Deep learning uses intelligent systems called artificial neural networks to process information in layers. Data flows from the input layer through multiple "deep" hidden neural network layers before coming to the output layer. The additional hidden layers support learning that's far more capable than that of standard machine learning models.

What is an artificial neural network?

The deep learning layers are made of artificial neural network (ANN) nodes that operate like human brain neurons. Nodes can be a combination of hardware and software. Each layer in a deep learning algorithm is made up of ANN nodes. Each node, or artificial neuron, connects to another and has an associated value number and threshold number. A node

sends its value number as an input to the next layer node when activated. It is activated only if its output is above the specified threshold value. Otherwise, no data is passed along.

What is computer vision?

Computer vision is a real-world application of deep learning. Just like artificial intelligence enables computers to think, computer vision enables them to see, observe and respond. Self-driving cars use computer vision to "read" road signs. A car's camera captures a photo of the sign. This photo is sent to the deep learning algorithm in the car. The first hidden layer detects edges, the next differentiate colors, while the third layer detects the details of the alphabet on the sign. The algorithm predicts that the sign reads STOP, and the car responds by triggering the brake mechanism.

Are machine learning and deep learning the same?

Deep learning is a subset of machine learning. Deep learning algorithms can be regarded both as a sophisticated and mathematically complex evolution of machine learning algorithms.

What are the advantages and disadvantages of machine learning?

Let us look at some things that machine learning can and cannot do.

Advantages of machine learning models:

· They can identify data trends and patterns that human beings might miss.

· They can work without human intervention after setup. For example, machine learning in cybersecurity software can continuously monitor and identify irregularities in network traffic without administrator input.

· Results can become more accurate over time

· They can handle a variety of data formats in dynamic and complex data environments.

Disadvantages of machine learning models:

· Initial training is a costly and time-consuming process. It may be hard to implement if sufficient data is not available.

· It is a compute-intensive process requiring heavy initial investment if the hardware is set up in-house.

· It can be challenging to interpret results correctly and eliminate uncertainty without expert help.

How can Amazon machine learning help?

AWS is putting machine learning in the hands of every developer, data scientist, and AI practitioner. Amazon machine learning services provide high-performing, cost-effective and scalable infrastructure to meet business needs.

· Just starting? – Use Amazon SageMaker Studio to build, train and deploy Machine Learning Models at scale.

· Have an existing data archive? – Use Amazon SageMaker Ground Truth for built-in data labeling workflows that support video, images, and text.

· Have existing Machine Learning systems? – Use Amazon SageMaker Clarify for detecting bias and Amazon SageMaker Debugger to monitor and optimize performance.

· Want to implement deep learning? – Use Amazon SageMaker Distributed Training to train large deep learning models automatically.

Sign up for a free account to begin your machine learning journey today!