

Introduction to Machine Learning*

February 20, 2025

1 What is machine learning?

Machine learning is a particular application of artificial intelligence (AI) that provides machines with the ability to automatically learn and improve from experience – but without being explicitly programmed to do so. This is the essence of machine learning. It focuses on building software that can access data and use it to train itself in order to provide better results, without software developers having to train it themselves. The primary goal of machine learning is to allow machines such as computers to learn automatically without any human intervention or assistance and then adjust their actions accordingly based on the insights they have uncovered.

As humans, we learn through past experiences. We use our senses to obtain these “experiences” and use them later to survive. Machines learn through commands provided by humans. These sets of rules are known as algorithms. Algorithms are sets of rules that a computer is able to follow. Think about how you learned to do long division — maybe you learned to divide the denominator into the first digits of the numerator, subtract the subtotal, and continue with the next digits until you were left with a remainder. Well, that’s an algorithm, and it’s the sort of thing we can program into a computer, which can perform these sorts of calculations much, much faster than we can.

In machine learning, our goal is either prediction or clustering. Prediction is a process where, from a set of input variables, we estimate the value of an output variable. This technique is used for data that has a precise mapping between input and output, referred to as labeled data. This is known as supervised learning. For example, using a set of characteristics of a house, we can predict its sale price. In clustering, we’re not trying to predict a variable; instead, we want to discover hidden patterns within our data that will let us identify groups, or clusters, within that data. For example, clustering is often used in marketing to group users according to multiple characteristics, such as location, purchasing behavior, age, and gender. It can also be used in scientific research, to find population clusters within DNA data.

Artificial intelligence VS. Machine learning VS. Deep learning.

- Artificial intelligence. The goals of artificial intelligence research include reasoning,

*References

- Wikipedia: Machine learning
- Veronika Gladchuk. The History of Machine Learning: How Did It All Start? July 6, 2020
- Christopher Bishop. Pattern Recognition and Machine Learning. 2006
- Andrew Ng. Machine Learning (Stanford course)




Figure 1: Machine learning as subfield of AI

knowledge representation, planning, learning, natural language processing, perception, and the ability to move and manipulate objects.

- **Machine learning:** Machine learning, reorganized as a separate field¹, started to flourish in the 1990s. The field changed its goal from achieving artificial intelligence to tackling solvable problems of a practical nature. ML learns and predicts based on passive observations, whereas AI implies an agent interacting with the environment to learn and take actions that maximize its chance of successfully achieving its goals.
- **Deep learning:** faster computers, algorithmic improvements, and access to large amounts of data enabled advances in machine learning and perception; data-hungry deep learning methods started to dominate accuracy benchmarks around 2012.

2 Successful application areas of machine learning

So far, machine learning has been very successful in many areas, for example

- **Data analysis.** Organizations generate more data than ever before, and machine learning algorithms come to the rescue for making sense of that data. Machine learning can speed up the process of uncovering the most valuable information from data sets by doing the heavy lifting in the time-consuming process of reviewing all the data. Machine learning-based tools assist managers in decision-making processes and help teams in departments such as sales, marketing, or production to crunch the numbers faster.
- **Personalization.** Customers expect to receive personalized experiences from brands today. Personalization, as a driver of customer loyalty, has become especially important when delivered via online and mobile apps. Machine learning can help companies to achieve that by offering customers with personalized product recommendations and channels that they use most.

¹Many sources continue to assert that machine learning remains a subfield of AI. Others have the view that not all machine learning is part of AI, but only an 'intelligent subset' of machine learning should be considered AI.

- **Fraud detection.** As more and more consumers turn to online channels for shopping, cybercriminals gain many opportunities to commit fraud. Organizations employ many types of online security measures, where machine learning holds the greatest promise. For example, they use machine learning tools to identify fraudulent transactions such as money laundering and separate them from legitimate ones. Machine learning algorithms help to examine specific features in the data set and build a model that offers a strong basis for reviewing every single transaction for signs that it could be fraudulent. That way, organizations can stop the process before the transaction is completed and avoid bigger problems.
- **Dynamic pricing.** The travel and retail industries see many opportunities for changing pricing based on fluctuating demand. However, incorporating dynamic pricing can be challenging across large enterprises with multiple locations or customer segments. This is where machine learning helps as well. For example, Airbnb and Uber use machine learning to create dynamic prices for each user on the go. Moreover, machine learning helps to minimize waste time and optimize the ride-sharing aspect of Uber. For instance, the app can temporarily change pricing in a given area to gain a higher revenue stream or reduce rates when the demand is much lower.
- **Natural language processing (NLP).** Tasks like tech support, help desks, customer service, and many others can now be solved thanks to machine learning algorithms and their capability for natural language processing. Computers can take over human agents because NLP offers automated translation between computer and human languages. Machine learning-powered tools like chatbots and virtual assistants focus on context, jargon, meaning, and many other subtle nuances in the human language to sound more human.

3 The history of machine learning

When thinking of AI, machine learning and smart computers, we tend to imagine something very contemporary, something that has appeared only recently. Many would indeed be surprised to know that machine learning history has started back in the 40s with a very important book on human cognition, and it has been accelerating only recently due to the development of new algorithms and methods but also due to the wide availability of the technology itself. It is impossible to pinpoint when machine learning was invented or who invented it, rather, it is a combination of many individuals' work, who contributed with separate inventions, algorithms or frameworks.

The following is only a partial list of important events in the history of AI and machine learning.

- 1943 – the first mathematical model of neural networks presented in the scientific paper "A logical calculus of the ideas immanent in nervous activity" by Walter Pitts and Warren McCulloch.
- 1950 – this is the year when Alan Turing, one of the most brilliant and influential British mathematicians and computer scientists, created the Turing test. The test was designed to determine whether a computer has human-like intelligence. In order to pass the test, the computer needs to be able to convince a human to

believe that it's another human. Apart from a computer program simulating a 13-year-old Ukrainian boy who is said to have passed the Turing test, there were no other successful attempts so far.

- 1952 – Arthur Samuel, a pioneer in machine learning, created a program for playing championship-level computer checkers. Instead of researching each and every possible path, the game used alpha-beta pruning that measured chances of winning. Additionally, Samuel utilized a minimax algorithm (which is still widely used for games today) of finding the optimal move, assuming that the opponent is also playing optimally. He also designed mechanisms for his program to continuously improve, for instance, by remembering previous checker moves and comparing them with chances of winning. Arthur Samuel is the first person to come up with and popularize the term "machine learning".
- 1956 – The Dartmouth Workshop is sometimes referred to as "the birthplace of artificial intelligence". During a two-month period, a group of prominent scientists in the fields of math, engineering, computer and cognitive sciences have gathered to establish and brainstorm the fields of AI and ML research.
- 1957 – this year witnessed the design of the very first neural network for computers called the perceptron by Frank Rosenblatt. It successfully stimulated the thought processes of the human brain. This is where today's neural networks originate from.
- 1965 – Ukrainian-born soviet scientists Alexey (Oleksii) Ivakhnenko and Valentin Lapa have developed hierarchical representation of neural network that uses polynomial activation function and are trained using Group Method of Data Handling (GMDH). It is considered as the first ever multi-layer perceptron and Ivakhnenko is often considered as the father of deep learning.
- 1967 – The nearest neighbor algorithm was written for the first time this year. It allows computers to start using basic pattern recognition. This algorithm can be used to map a route for a traveling salesman that starts in a random city and ensures that the salesman passes by all the required cities in the shortest time. Today, the nearest neighbor algorithm called KNN is mostly used to classify a data point on the basis of how their neighbors are classified. KNN is used in retail applications that recognize patterns in credit card usage or for theft prevention when implemented in CCTV image recognition in retail stores. (Thomas Cover and Peter E. Hart. Nearest Neighbor Pattern Classification. IEEE Transactions on Information Theory, Volume: 13, Issue: 1, January 1967, pages 22-27).
- 1979 – Japanese computer scientist Kunihiko Fukushima publishes his work on neocognitron, a hierarchical multilayered network which is used to detect patterns and inspires convolutional neural networks — systems used nowadays for analyzing images.
- 1981 – Gerald Dejong introduced the concept of explanation-based learning (EBL). In this type of learning, the computer analyzes training data and generates a general rule that it can follow by discarding the data that doesn't seem to be important.

- 1985 – Terry Sejnowski invented the NetTalk program that could learn to pronounce words just like a baby does during the process of language acquisition. The artificial neural network aimed to reconstruct a simplified model that would show the complexity of learning human-level cognitive tasks.
- 1990 – Paper “The Strength of Weak Learnability” by Robert Schapire and Yoav Freund introduce boosting for machine learning. Boosting is an algorithm which aims to enhance predicting power of an AI model. Instead of using a single strong model, it generates many weak models and converts them into strong ones by combining their predictions (usually, using averages or voting).
- 1995 – Random decision forests are introduced in a paper published by Tin Kam Ho. This algorithm creates and merges multiple AI decisions into a ”forest”. When relying on multiple different decision trees, the model significantly improves in its accuracy and decision-making.
- 1997 – IBM chess computer, Deep Blue, beats world champion Garry Kasparov in chess. At the time this achievement was seen as a proof of machines catching up to human intelligence.
- 2000 – First mention of the term ”deep learning” by a Ukrainian-born neural networks researcher Igor Aizenberg in the context of Boolean threshold neurons.
- 2006 – this is the year when the term “deep learning” was coined by Geoffrey Hinton. He used the term to explain a brand-new type of algorithms that allow computers to see and distinguish objects or text in images or videos.
- 2009 – A massive visual database of labeled images ImageNet is launched by Fei-Fei Li. Li wanted to expand on the data available for training algorithms, since she believed that AI and ML must have good training data that reflects the real world in order to be truly practical and useful. The Economist described the creation of this database as an exceptional event for popularizing AI throughout the whole tech community, marking the new era of deep learning history.
- 2010 – this year saw the introduction of Microsoft Kinect that could track even 20 human features at the rate of 30 times per second. Microsoft Kinect allowed users to interact with machines via gestures and movements.
- 2011 – this was an interesting year for machine learning. For starters, IBM’s Watson managed to beat human competitors at Jeopardy. Moreover, Google developed Google Brain equipped with a deep neural network that could learn to discover and categorize objects (in particular, cats).
- 2012 – Google X lab developed a machine learning algorithm able to autonomously browse YouTube videos and identify those that contained cats.
- 2014 – A group of prominent scientists (Goodfellow, Pouget-Abadie, Mirza, Xu, Warde-Farley, Ozair, Courville, Bengio) develop Generative adversarial networks (GAN) frameworks that teach AI how to generate new data based on the training set.

- 2014 – Facebook research team develops DeepFace, a deep learning facial recognition system — nine-layer neural network trained on 4 million images of Facebook users. This AI is able to spot human faces in images with the same accuracy as humans do (approximately 97.35%).
- 2015 – this is the year when Amazon launched its own machine learning platform, making machine learning more accessible and bringing it to the forefront of software development. Moreover, Microsoft created the Distributed Machine Learning Toolkit, which enables developers to efficiently distribute machine learning problems across multiple machines. During the same year, however, more than three thousand AI and robotics researchers endorsed by figures like Elon Musk, Stephen Hawking, and Steve Wozniak signed an open letter warning about the dangers of autonomous weapons that could select targets without any human intervention.
- 2016 – this was the year when Google’s artificial intelligence algorithms managed to beat a professional player at the Chinese board game Go. Go is one of the oldest and hardest abstract strategy games, which was previously thought to be a near-impossible game to teach a computer. The AlphaGo algorithm developed by Google won five out of five games in the competition, bringing AI to the front page.
- 2020 – Open AI announced a groundbreaking natural language processing algorithm GPT-3 with a remarkable ability to generate human-like text when given a prompt. Today, GPT-3 is considered the largest and most advanced language model in the world, using 175 billion parameters and Microsoft Azure’s AI supercomputer for training.
- 2021 – DeepMind (a subsidiary of Alphabet) used its neural network to tackle one of biology’s grand challenges, the protein-folding problem. Its neural net, known as AlphaFold, was able to predict the 3D structures of proteins based on their amino acid sequences with unprecedented accuracy. DeepMind’s breakthrough demonstrates that deep learning has the potential to dramatically accelerate scientific discovery.

4 The future of machine learning

Improvements in unsupervised learning algorithms. In the future, we’ll see more effort dedicated to improving unsupervised machine learning algorithms to help to make predictions from unlabeled data sets. This function is going to become increasingly important as it allows algorithms to discover interesting hidden patterns or groupings within data sets and help businesses understand their market or customers better.

The rise of quantum computing. One of the major applications of machine learning trends lies in quantum computing that could transform the future of this field. Quantum computers lead to faster processing of data, enhancing the algorithm’s ability to analyze and draw meaningful insights from data sets.

Focus on cognitive services. Software applications will become more interactive and intelligent thanks to cognitive services driven by machine learning. Features such as visual recognition, speech detection, and speech understanding will be easier to implement. We’re going to see more intelligent applications using cognitive services appear on the market.

5 The classification of machine learning algorithms

We can generally divide machine learning algorithms into those based on supervised and unsupervised learning.

5.1 Unsupervised learning

Unsupervised machine learning algorithms are used when data fed to the algorithm is not labeled or classified in any way. Unsupervised learning is based on systems that analyze patterns in unlabeled data. The system doesn't arrive at the output that can be just right or wrong, but rather it explores the data and draws inferences from data sets.

5.2 Supervised learning

Supervised machine learning algorithms are able to apply what they have learned in the past to new data with the help of labeled examples. They do that to predict future events. Such algorithms begin with the analysis of the known training data set and then produce insights to make predictions about the output values. The system can provide targets for any new input after being exposed to sufficient training. The algorithm can also compare its results with the correct and intended output to find errors and modify the model accordingly.

5.3 Semi-supervised learning

Semi-supervised machine learning algorithms fall somewhere in between these two sides. They use both labeled and unlabeled data for training. Typically, it's a small number of labeled data and a very large amount of unlabeled data. Systems that use this method can considerably improve their learning accuracy. Most of the time, semi-supervised learning systems are chosen when the acquired labeled data calls for skilled and relevant resources in order to learn from it. Otherwise, acquired unlabeled data doesn't really mean that we need any additional resources.

5.4 Reinforcement learning

Reinforcement machine learning algorithms represent another type of learning method. This type of algorithm interacts with the environment by producing actions and then discovering errors or rewards. Delayed reward or trial and error searches are the key characteristics of reinforcement learning. The method allows machines to automatically determine the best behavior within a specific context to maximize their performance and get the best reward. They require to learn which actions work best – the so-called reinforcement signal.

6 Limitation and danger of using machine learning

Machine learning has been successful in many applications. However, it is important to understand that machine learning is not the answer to all problems. There are times when using machine learning is just unnecessary, does not make sense, and other times when its implementation can get you into difficulties.

6.1 Ethics

It is easy to understand why machine learning has had such a profound impact on the world, what is less clear is exactly what its capabilities are, and perhaps more importantly, what its limitations are. Yuval Noah Harari famously coined the term ‘dataism’, which refers to a putative new stage of civilization we are entering in which we trust algorithms and data more than our own judgment and logic.

The idea of trusting data and algorithms more than our own judgment has its pros and cons. These algorithms allow us to automate processes by making informed judgments using available data. Sometimes, however, this means replacing someone’s job with an algorithm, which comes with ethical ramifications. Additionally, who do we blame if something goes wrong? The most commonly discussed case currently is self-driving cars — how do we choose how the vehicle should react in the event of a fatal collision? In the future will we have to select which ethical framework we want our self-driving car to follow when we are purchasing the vehicle? If my self-driving car kills someone on the road, whose fault is it? Clearly, machine learning cannot tell us anything about what normative values we should accept, i.e. how we should act in the world in a given situation.

6.2 Deterministic problems

Machine learning is stochastic, not deterministic. Below are two examples:

Example 1: Weather forecast. Running computer models that simulate global weather is very computationally expensive. In fact, it is so computationally expensive, that a research-level simulation can take weeks even when running on a supercomputer. Running weather prediction models is fine, but now that we have machine learning, can we just use this instead to obtain our weather forecasts? Can we leverage data from satellites, weather stations, and use an elementary predictive algorithm to discern whether it is going to rain tomorrow? The answer is, surprisingly, yes. If we have knowledge of the air pressures around a certain region, the levels of moisture in the air, wind speeds, and information about neighboring points and their own variables, it becomes possible to train, for example, a neural network. Using a neural network with a thousand inputs to determine whether it will rain tomorrow in a region is possible. However, utilizing a neural network misses the entire physics of the weather system, and you can never assert that a result is 100% correct.

Example 2: Medical care. The most obvious risk is that AI systems will sometimes be wrong, and that patient injury or other health-care problems may result. If an AI system recommends the wrong drug for a patient, fails to notice a tumor on a radiological scan, or allocates a hospital bed to one patient over another because it predicted wrongly which patient would benefit more, the patient could be injured. Of course, many injuries occur due to medical error in the health-care system today, even without the involvement of AI. AI errors are potentially different for at least two reasons. First, patients and providers may react differently to injuries resulting from software than from human error. Second, if AI systems become widespread, an underlying problem in one AI system might result in injuries to thousands of patients—rather than the limited number of patients injured by any single provider’s error.

6.3 Data and data quality

This is the most obvious limitation. If you feed a model poorly, then it will only give you poor results. This can manifest itself in two ways: lack of data, and lack of good data.

Lack of data. Many machine learning algorithms require large amounts of data before they begin to give useful results. A good example of this is a neural network that require copious amounts of training data. The larger the architecture, the more data is needed to produce viable results. Reusing data is a bad idea, and data augmentation is useful to some extent, but having more data is always the preferred solution.

Lack of good data. Let's imagine you think you can cheat by generating ten thousand fake data points to put in your neural network. What happens when you put it in? It will train itself, and then when you come to test it on an unseen data set, it will not perform well. You had the data but the quality of the data was not up to scratch. In the same way that having a lack of good features can cause your algorithm to perform poorly, having a lack of good ground truth data can also limit the capabilities of your model.

Data biases. The infallibility of an AI solution is based on the quality of its inputs. For example, facial recognition has had a large impact on social media, human resources, law enforcement, and other applications. But biases in the data sets provided by facial recognition applications can lead to inexact outcomes. If the training data is not neutral the outcomes will inherently amplify the discrimination and bias that lies in the data set. The most ideal way to mitigate such risks is by collecting data from multiple random sources. A heterogeneous dataset limits the exposure to bias and results in higher quality ML solutions.

6.4 Interpretability

Interpretability is one of the primary problems with machine learning. An AI consultancy firm trying to pitch to a firm that only uses traditional statistical methods can be stopped dead if they do not see the model as interpretable. If you cannot convince your client that you understand how the algorithm came to the decision it did, how likely are they to trust you and your expertise? As bluntly stated in *Business Data Mining — a machine learning perspective*²: *A business manager is more likely to accept the [machine learning method] recommendations if the results are explained in business terms.* These models as such can be rendered powerless unless they can be interpreted, and the process of human interpretation follows rules that go well beyond technical prowess. For this reason, interpretability is a paramount quality that machine learning methods should aim to achieve if they are to be applied in practice. The blossoming *-omics* sciences (genomics, proteomics, metabolomics and the like), in particular, have become the main target for machine learning researchers precisely because of their dependence on large and non-trivial databases. However, they suffer from the lack of interpretability of their methods, despite their apparent success.

²Indranil Bose^{1a} and Radha Mahapatra. Business data mining — a machine learning perspective. Information & Management. 39(3), 211-225, 20 December 2001

6.5 Susceptibility to adversarial attack or manipulation

Computer scientists regularly test machine learning systems with so-called "adversarial examples" crafted to make the systems misclassify them in order to find out the possible limitations of current deep learning methods. However, the existence of adversarial examples also indicates that machine learning algorithms are susceptible to risk of adversarial attack. An adversarial attack describes an otherwise-effective model that is susceptible to manipulation by inputs explicitly designed to fool them. For example, a team at Harvard Medical School and MIT showed that it's pretty easy to fool an AI system analyzing medical images³. In their tests, pixels within images are modified in a way that might seem like a minimal amount of noise to humans, but could trick these systems into classifying these pictures incorrectly. The scientists note their attacks could make deep learning systems misclassify images up to 100 percent of the time, and that modified images were imperceptible from real ones to the human eye. Researchers also shown that such attacks could work on any image, and could even be incorporated directly into the image-capture process. If someone has access to the underlying data, then they could commit many different kinds of fraud, not just using adversarial attacks (that it would be very difficult to detect that the attack has occurred). This is why AI safety or security is another hot topic in AI research.

7 Common terms in machine learning

You have heard some of the following terminology in the first lecture of the course. Most of these terminology will be further explained in future lectures.

- **Algorithm:** a set of rules used to make a calculation or solve a problem.
- **Feature:** also known as an independent variable or a predictor variable. A feature is an observable quantity, recorded and used by a prediction model. You can also engineer features by combining them or adding new information to them.
- **Model:** a mathematical representation of a real world process; a predictive model forecasts a future outcome based on past behaviors.
- **Regression:** a prediction method whose output is a continuous real number, that is, a value that represents a quantity. For example: predicting the temperature of an engine or the revenue of a company.
- **Classification:** a prediction method that assigns each data point to a predefined category (a discrete value), e.g., a type of object (car, human, tree).
- **Training set:** a dataset used to find potentially predictive relationships that will be used to create a model.
- **Test set:** a dataset, separate from the training set but with the same structure, used to measure and benchmark the performance of various models.

³Charles Choi. Medical Imaging AI Software Is Vulnerable to Covert Attacks. IEEE Spectrum. 04 Jun 2018

- Training/Learning: the process of creating a model from the training data. The data is fed into the training algorithm, which learns a representation for the problem, and produces a model.
- Target: in statistics, it is called the dependent variable; it is the output of the model or the variable you wish to predict.
- Over fitting: a situation in which a model that is too complex for the data has been trained to predict the target. This leads to an overly specialized model, which makes predictions not reflect the reality of the underlying relationship between the features and target.