

MACHINE LEARNING



Prepared by- Revanwar Gayatri
if5i(1531)

Content

what is machine learning

machine learning methods

how machine learning works

real world use cases

advantages and disadvantages of machine learning

**what is future of machine
learning**

what is Artificial Intelligence

**Why is artificial intelligence
important**

**what are the advantages and disadvantages of
artificial intelligence**

Strong AI vs. weak AI

What are the 4 types of artificial intelligence?

**What are examples of AI technology and how is it used
today?**

VISION

“Become premier center in Information Technology with value based education that will prepare students for ever changing technological challenges of 21st century”.

MISSION

- To train the students in the latest technologies.
- Provide an environment that inculcates ethics and effective soft-skills.
- Develop the skill sets among students that will benefit employer and society.

Program Outcomes (POs)

- Basic and Discipline specific knowledge : Apply knowledge of basic mathematics, science and engineering fundamentals and engineering specialization to solve the engineering problems
- Problem analysis : Identify and analyse well-defined engineering problems using codified standard methods
- Design/ development of solutions : Design solutions for well-defined technical problems and assist with the design of systems components or processes to meet specified needs
- Engineering Tools, Experimentation and Testing : Apply modern engineering tools and appropriate technique to conduct standard tests and measurements
 - Engineering practices for society, sustainability and environment : Apply appropriate technology in context of society, sustainability, environment and ethical practices.
- Project Management : Use engineering management principles individually, as a team member or a leader to manage projects and effectively communicate about well-defined engineering activities
- Life-long learning : Ability to analyse individual needs and engage in updating in the context of technological changes

Program Educational Objectives (PEOs)

- Become competent Information Technology engineer to work as a programmer or an administrator in a team or as an individual.
- Pursue higher studies in relevant field of engineering with a desire for lifelong learning.
 - Become a successful professional with ethical and societal responsibilities.

Program Specific Outcomes (PSOs)

- Information Technology: Use latest technologies for operation and application of information.
 - Information Technology Process: Maintain the information processes using modern information and communication technologies.

what is machine learning

Machine learning is a branch of artificial intelligence (AI) and computer science which focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy.

IBM has a rich history with machine learning. One of its own, Arthur Samuel, is credited for coining the term, “machine learning” with his research (PDF, 481 KB) (link resides outside IBM) around the game of checkers. Robert Nealey, the self-proclaimed checkers master, played the game on an IBM 7094 computer in 1962, and he lost to the computer. Compared to what can be done today, this feat almost seems trivial, but it’s considered a major milestone within the field of artificial intelligence.

Over the next couple of decades, the technological developments around storage and processing power will enable some innovative products that we know and love today, such as Netflix’s recommendation engine or self-driving cars.

Machine learning is an important component of the growing field of data science. Through the use of statistical methods, algorithms are trained to make classifications or predictions, uncovering key insights within data mining projects. These insights subsequently drive decision making within applications and businesses, ideally impacting key growth metrics. As big data continues to expand and grow, the market demand for data scientists will increase, requiring them to assist in the identification of the most relevant business questions and subsequently the data to answer them.

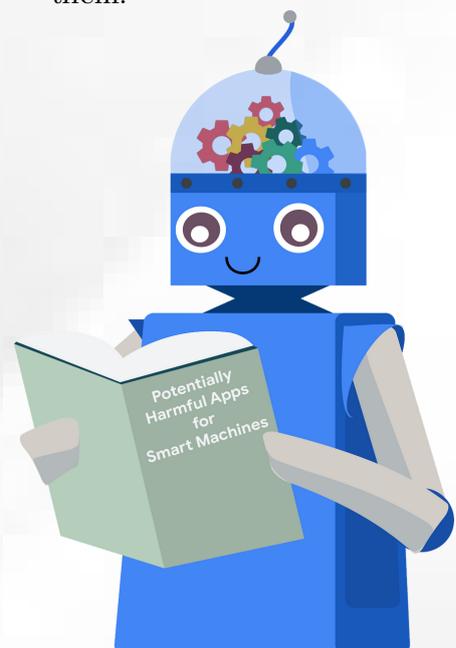
Machine learning is a branch of artificial intelligence (AI) and computer science which focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy.

IBM has a rich history with machine learning. One of its own, Arthur Samuel, is credited for coining the term, “machine learning” with his research (PDF, 481 KB) (link resides outside IBM) around the game of checkers.

Robert Nealey, the self-proclaimed checkers master, played the game on an IBM 7094 computer in 1962, and he lost to the computer. Compared to what can be done today, this feat almost seems trivial, but it’s considered a major milestone within the field of artificial intelligence.

Over the next couple of decades, the technological developments around storage and processing power will enable some innovative products that we know and love today, such as Netflix’s recommendation engine or self-driving cars.

Machine learning is an important component of the growing field of data science. Through the use of statistical methods, algorithms are trained to make classifications or predictions, uncovering key insights within data mining projects. These insights subsequently drive decision making within applications and businesses, ideally impacting key growth metrics. As big data continues to expand and grow, the market demand for data scientists will increase, requiring them to assist in the identification of the most relevant business questions and subsequently the data to answer them.



machine learning methods

Machine learning classifiers fall into three primary categories.

Supervised machine learning

Supervised learning, also known as supervised machine learning, is defined by its use of labeled datasets to train algorithms that to classify data or predict outcomes accurately. As input data is fed into the model, it adjusts its weights until the model has been fitted appropriately. This occurs as part of the cross validation process to ensure that the model avoids overfitting or underfitting. Supervised learning helps organizations solve for a variety of real-world problems at scale, such as classifying spam in a separate folder from your inbox. Some methods used in supervised learning include neural networks, naïve bayes, linear regression, logistic regression, random forest, support vector machine (SVM), and more.

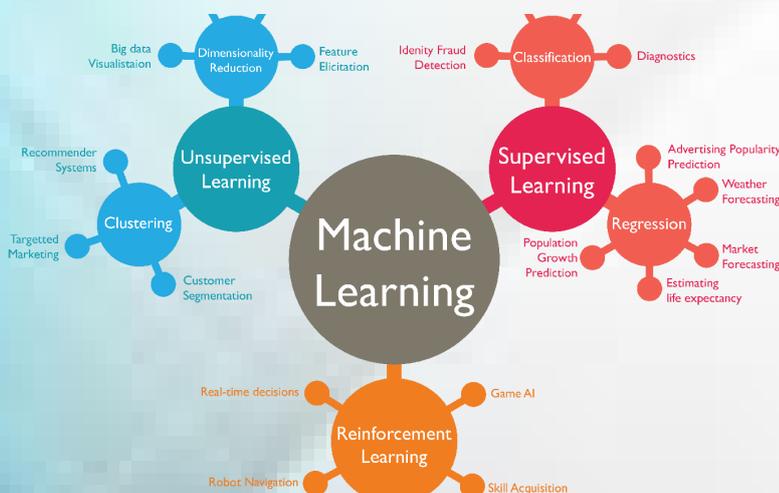
Unsupervised machine learning
Unsupervised learning, also known as unsupervised machine learning, uses machine learning algorithms to analyze and cluster unlabeled datasets. These algorithms discover hidden patterns or data groupings without the need for human intervention. Its ability to discover similarities and differences in information make it the ideal solution for exploratory data analysis, cross-selling strategies, customer segmentation, image and pattern recognition. It's also used to reduce the number of features in a model

through the process of dimensionality reduction; principal component analysis (PCA) and singular value decomposition (SVD) are two common approaches for this.

Other algorithms used in unsupervised learning include neural networks, k-means clustering, probabilistic clustering methods, and more.

Semi-supervised learning

Semi-supervised learning offers a happy medium between supervised and unsupervised learning. During training, it uses a smaller labeled data set to guide classification and feature extraction from a larger, unlabeled data set. Semi-supervised learning can solve the problem of having not enough labeled data (or not being able to afford to label enough data) to train a supervised learning algorithm.



How machine learning works

UC Berkeley (link resides outside IBM) breaks out the learning system of a machine learning algorithm into three main parts

1. **A Decision Process:** In general, machine learning algorithms are used to make a prediction or classification. Based on some input data, which can be labelled or unlabeled, your algorithm will produce an estimate about a pattern in the data.
2. **An Error Function:** An error function serves to evaluate the prediction of the model. If there are known examples, an error function can make a comparison to assess the accuracy of the model.
3. **An Model Optimization Process:** If the model can fit better to the data points in the training set, then weights are adjusted to reduce the discrepancy between the known example and the model estimate. The algorithm will repeat this evaluate and optimize process, updating weights autonomously until a threshold of accuracy has been met.



Real world use cases

Customer service: Online chatbots are replacing human agents along the customer journey. They answer frequently asked questions (FAQs) around topics, like shipping, or provide personalized advice, cross-selling products or suggesting sizes for users, changing the way we think about customer engagement across websites and social media platforms. Examples include messaging bots on e-commerce sites with virtual agents, messaging apps, such as Slack and Facebook Messenger, and tasks usually done by virtual assistants and voice assistants.

Automated stock trading: Designed to optimize stock portfolios, AI-driven high-frequency trading platforms make thousands or even millions of trades per day without human intervention. Challenges of machine learning

As machine learning technology advances, it has certainly made our lives easier. However, implementing machine learning within businesses has also raised a number of ethical concerns surrounding AI technologies. Some of these include:

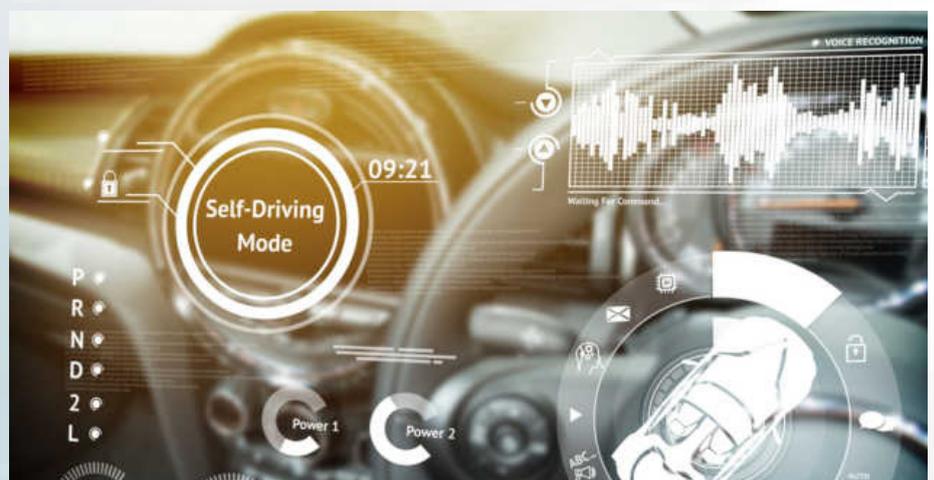
- Technological singularity

Speech recognition: It is also known as automatic speech recognition (ASR), computer speech recognition, or speech-to-text, and it is a capability which uses natural language processing (NLP) to process human speech into a written format. Many mobile devices incorporate speech recognition into their systems to conduct voice search—e.g. Siri—or provide more accessibility around texting.



Computer vision: This AI technology enables computers and systems to derive meaningful information from digital images, videos and other visual inputs, and based on those inputs, it can take action. This ability to provide recommendations distinguishes it from image recognition tasks. Powered by convolutional neural networks, computer vision has applications within photo tagging in social media, radiology imaging in healthcare, and self-driving cars within the automotive industry.

Recommendation engines: Using past consumption behavior data, AI algorithms can help to discover data trends that can be used to develop more effective cross-selling strategies. This is used to make relevant add-on recommendations to customers during the checkout process for online retailers



While a lot of public perception around artificial intelligence centers around job loss, this concern should be probably reframed. With every disruptive, new technology, we see that the market demand for specific job roles shift. For example, when we look at the automotive industry, many manufacturers, like GM, are shifting to focus on electric vehicle production to align with green initiatives. The energy industry isn't going away, but the source of energy is shifting from a fuel economy to an electric one. Artificial intelligence should be viewed in a similar manner, where artificial intelligence will shift the demand of jobs to other areas. There will need to be individuals to help manage these systems as data grows and changes every day. There will still need to be resources to address more complex problems within the industries that are most likely to be affected by job demand shifts, like customer service. The important aspect of artificial intelligence and its effect on the job market will be helping individuals transition to these new areas of market demand.

Privacy

Privacy tends to be discussed in the context of data privacy, data protection and data security, and these concerns have allowed policymakers to make more strides here in recent years. For example, in 2016, GDPR legislation was created to protect the personal data of people in the European Union and European Economic Area, giving individuals more control of their data. In the United States, individual states are developing policies, such as the California Consumer Privacy Act (CCPA), which require businesses to inform consumers about the collection of their data. This recent legislation has forced companies to rethink how they store and use personally identifiable data (PII). As a result, investments within security have become an increasing priority for businesses as they seek to eliminate any vulnerabilities and opportunities for surveillance, hacking, and cyberattack

Bias and discrimination

Instances of bias and discrimination across a number of intelligent systems have raised many ethical questions regarding the use of artificial intelligence. How can we safeguard against bias and discrimination when the training data itself can lend itself to bias? While companies typically have well-meaning intentions around their automation efforts, Reuters ([link resides outside IBM](#)) highlights some of the unforeseen consequences of incorporating AI into hiring practices. In their effort to automate and simplify a process, Amazon unintentionally biased potential job candidates by gender for open technical roles, and they ultimately had to scrap the project. As events like these surface, Harvard Business Review ([link resides outside IBM](#)) has raised other pointed questions around the use of AI within hiring practices, such as what data should you be able to use when evaluating a candidate for a role.

Bias and discrimination aren't limited to the human resources function either; it can be found in a number of applications from facial

What are the advantages and disadvantages of machine learning?

Machine learning has seen use cases ranging from predicting customer behavior to forming the operating system for self-driving cars.

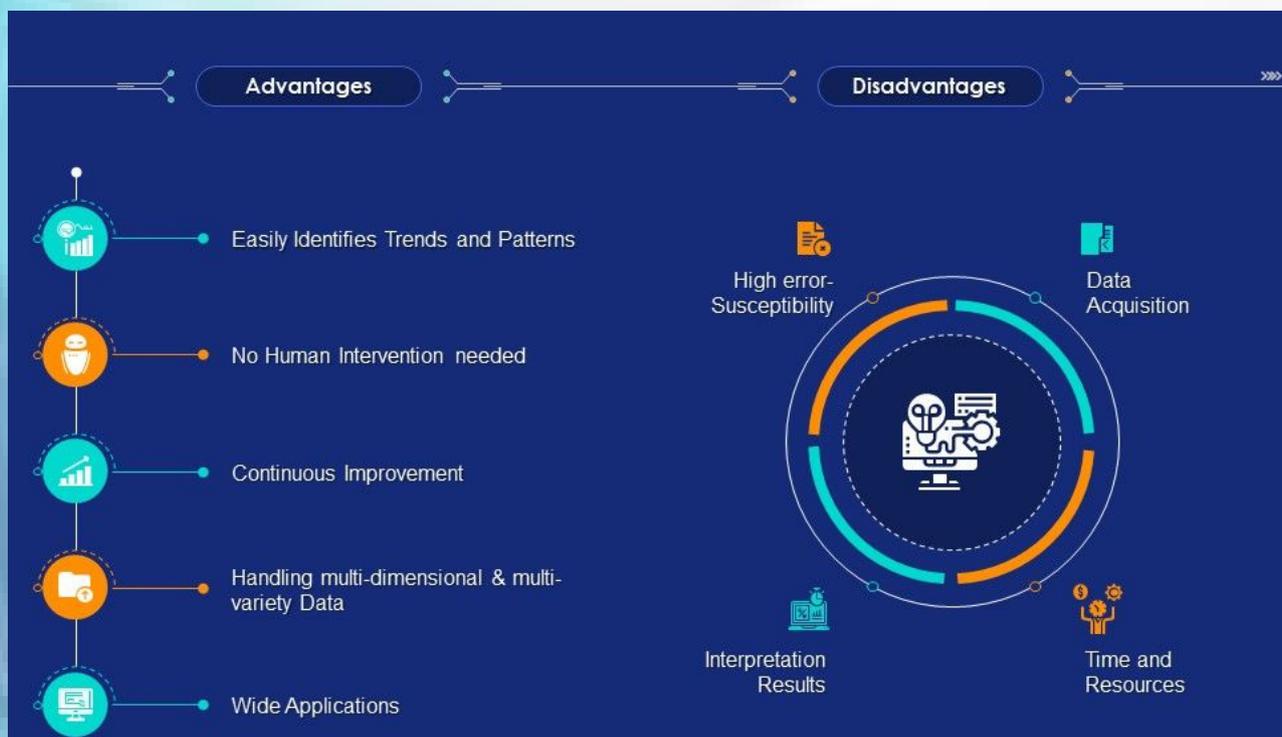
When it comes to advantages, machine learning can help enterprises understand their customers at a deeper level. By collecting customer data and correlating it with behaviors over time, machine learning algorithms can learn associations and help teams tailor product development and marketing initiatives to customer demand.

Some companies use machine learning as a primary driver in their business models. Uber, for example, uses algorithms to match drivers with riders. Google uses machine learning to surface the ride advertisements in searches.

But machine learning comes with disadvantages. First and foremost, it can be expensive. Machine learning projects are typically driven by data scientists, who command high salaries. These projects also require software infrastructure that can be expensive.

There is also the problem of machine learning bias. Algorithms trained on data sets that exclude certain populations or contain errors can lead to inaccurate models of the world that, at best, fail and, at worst, are discriminatory.

When an enterprise bases core business processes on biased models it can run into regulatory and reputational harm



What is the future of machine learning?

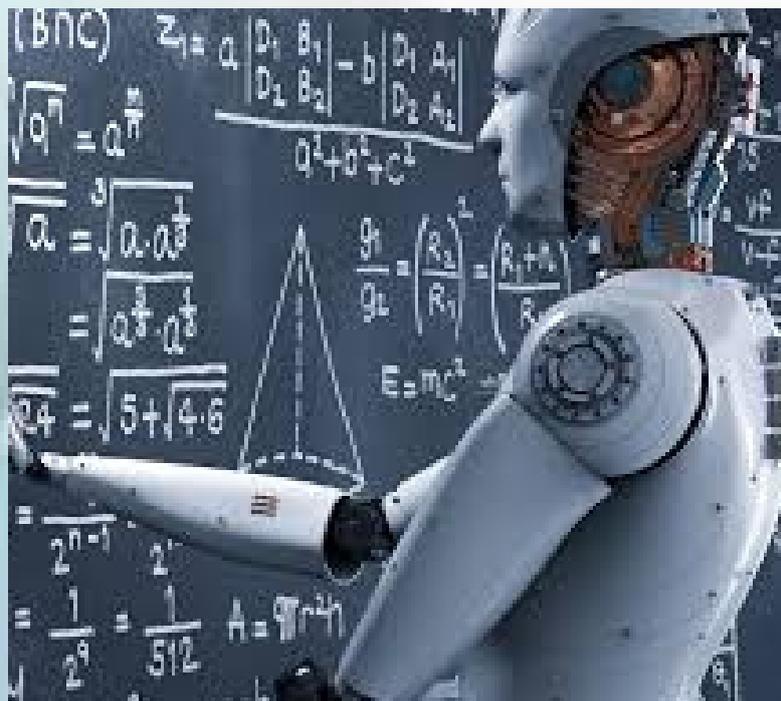
While machine learning algorithms have been around for decades, they've attained new popularity as artificial intelligence has grown in prominence. Deep learning models, in particular, power today's most

advanced AI applications.

Machine learning platforms are among enterprise technology's most competitive realms, with most major vendors, including Amazon, Google, Microsoft, IBM and others, racing to sign customers up for platform services that cover the spectrum of machine learning activities, including data collection, data preparation, data classification, model building, training and application deployment.

As machine learning continues to increase in importance to business operations and AI becomes more practical in enterprise settings, the machine learning platform wars will only intensify.

Continued research into deep learning and AI is increasingly focused on developing more general applications. Today's AI models require extensive training in order to produce an algorithm that is highly optimized to perform one task. But some researchers are exploring ways to make models more flexible and are seeking techniques that allow a machine to apply context learned from one task to future, different tasks



what is Artificial Intelligence

By

Ed Burns, Executive Editor
Nicole Laskowski, Senior News Director
Linda Tucci, Editor at Large-Strategic Initiatives
Artificial intelligence is the simulation of human intelligence processes by machines, especially computer systems. Specific applications of AI include expert systems, natural language processing, speech recognition and machine vision.

How does AI work?

As the hype around AI has accelerated, vendors have been scrambling to promote how their products and services use AI. Often what they refer to as AI is simply one component of AI, such as machine learning. AI requires a foundation of specialized hardware and software for writing and training machine learning algorithms. No one programming language is synonymous with AI, but a few, including Python, R and Java, are popular.

In general, AI systems work by ingesting large amounts of labeled training data, analyzing the data for correlations and patterns, and using these patterns to make predictions about future states. In this way, a chatbot that is fed examples of text chats can learn to produce lifelike exchanges with people, or an image recognition tool can learn to identify and describe objects in images by reviewing millions of examples

AI programming focuses on three cognitive skills: learning, reasoning and self-correction.

1. Learning processes. This aspect of AI programming focuses on acquiring data and creating rules for how to turn the data into actionable information. The rules, which are called algorithms, provide computing devices with step-by-step instructions for how to complete a specific task.

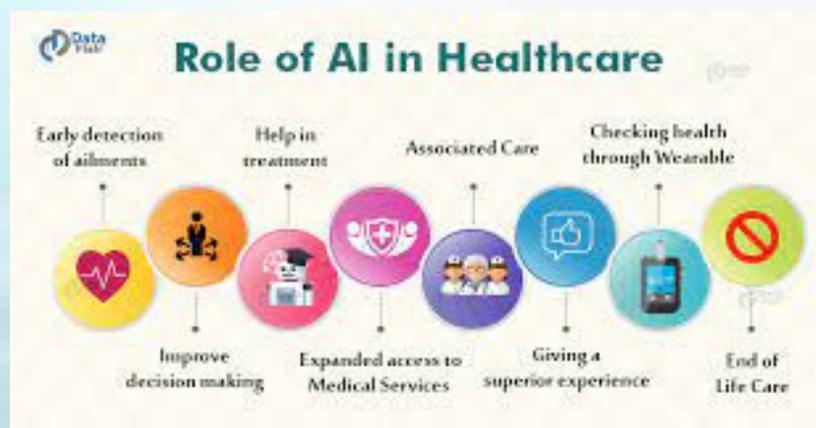
2. Reasoning processes. This aspect of AI programming focuses on choosing the right algorithm to reach a desired outcome.

3. Self-correction processes. This aspect of AI programming is designed to continually fine-tune algorithms and ensure they provide the most accurate results possible.



Why is artificial intelligence important

AI is important because it can give enterprises insights into their operations that they may not have been aware of previously and because, in some cases, AI can perform tasks better than humans. Particularly when it comes to repetitive, detail-oriented tasks like analyzing large numbers of legal documents to ensure relevant fields are filled in properly, AI tools often complete jobs quickly and with relatively few errors.



This has helped fuel an explosion in efficiency and opened the door to entirely new business opportunities for some larger enterprises. Prior to the current wave of AI, it would have been hard to imagine using computer software to connect riders to taxis, but today Uber has become one of the largest companies in the world by doing just that. It utilizes sophisticated machine learning algorithms to predict when people are likely to need rides in certain areas, which helps proactively get drivers on the road before they're needed. As another example, Google has become one of the largest players for a range of online services by using machine learning to understand how people use their services and then improving them. In 2017, the company's CEO, Sundar Pichai, pronounced that Google would operate as an "AI first" company.

Today's largest and most successful enterprises have used AI to improve their operations and gain advantage on their competitors

what are the advantages and disadvantages of artificial intelligence

Artificial neural networks and deep learning artificial intelligence technologies are quickly evolving, primarily because AI processes large amounts of data much faster and makes predictions more accurately than humanly possible.

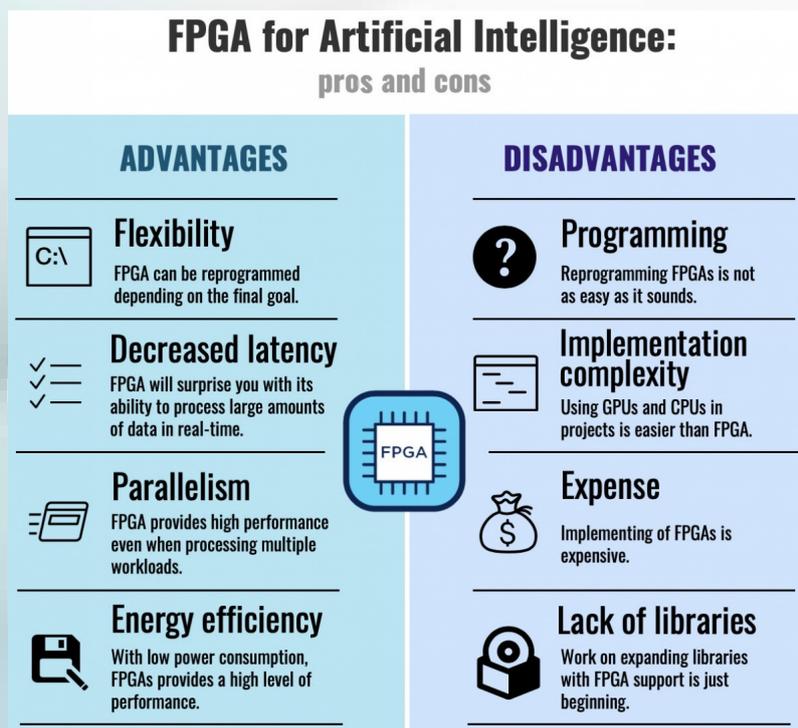
While the huge volume of data being created on a daily basis would bury a human researcher, AI applications that use machine learning can take that data and quickly turn it into actionable information. As of this writing, the primary disadvantage of using AI is that it is expensive to process the large amounts of data that AI programming requires.

Advantages

Good at detail-oriented jobs;
Reduced time for data-heavy tasks;
Delivers consistent results; and
AI-powered virtual agents are always available.

Disadvantages

Expensive;
Requires deep technical expertise;
Limited supply of qualified workers to build AI tools;
Only knows what it's been shown; and
Lack of ability to generalize from one task to another.

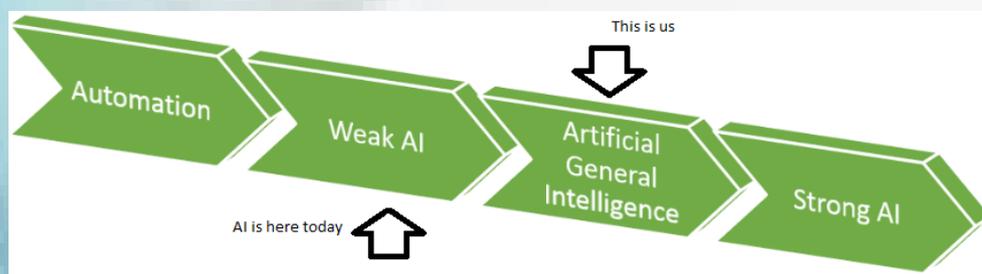
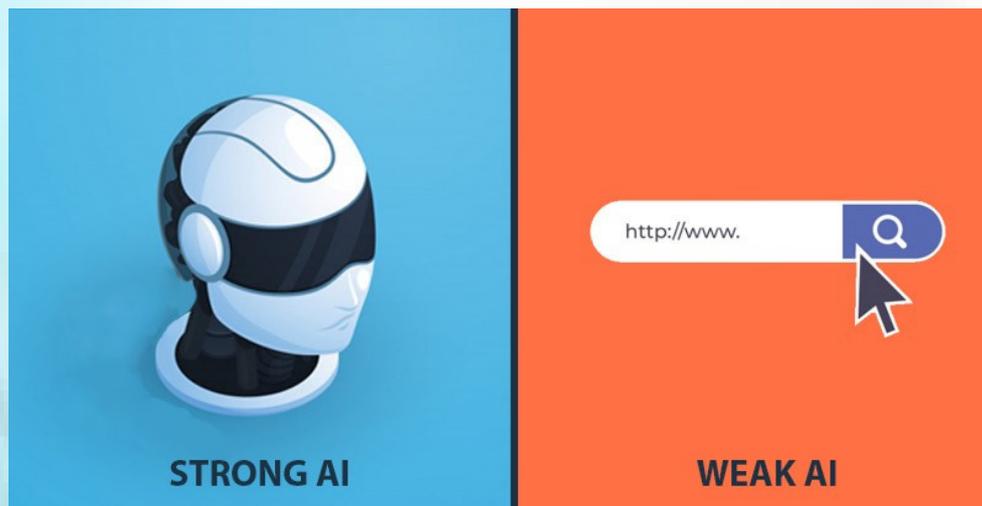


Strong AI vs. weak AI

AI can be categorized as either weak or strong.

Weak AI, also known as narrow AI, is an AI system that is designed and trained to complete a specific task. Industrial robots and virtual personal assistants, such as Apple's Siri, use weak AI.

Strong AI, also known as artificial general intelligence (AGI), describes programming that can replicate the cognitive abilities of the human brain. When presented with an unfamiliar task, a strong AI system can use fuzzy logic to apply knowledge from one domain to another and find a solution autonomously. In theory, a strong AI program should be able to pass both a Turing Test and the Chinese room test.



What are the 4 types of artificial intelligence?

Arend Hintze, an assistant professor of integrative biology and computer science and engineering at Michigan State University, explained in a 2016 article that AI can be categorized into four types, beginning with the task-specific intelligent systems in wide use today and progressing to sentient systems, which do not yet exist. The categories are as follows:

Type 1: Reactive machines. These AI systems have no memory and are task specific. An example is Deep Blue,

the IBM chess program that beat Garry Kasparov in the 1990s.

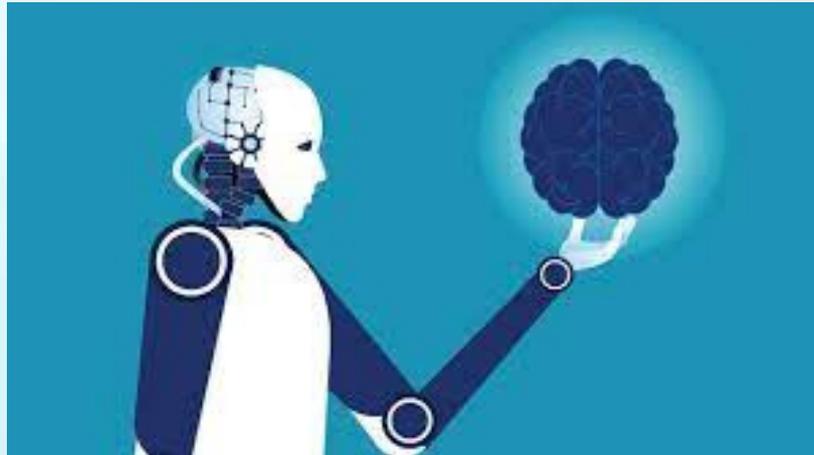
Deep Blue can identify pieces on the chessboard and make predictions, but because it has no memory, it cannot use past experiences to inform future ones.



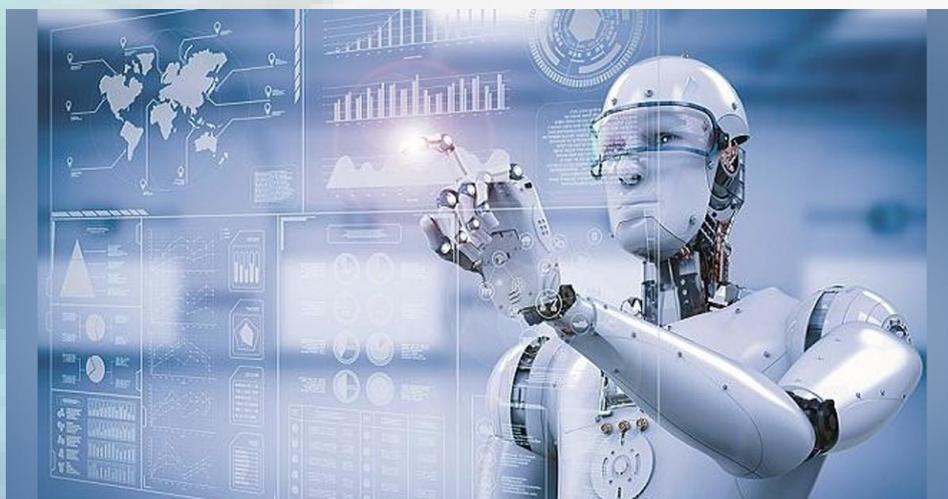
Type 2: Limited memory. These AI systems have memory, so they can use past experiences to inform future decisions. Some of the decision-making functions in self-driving cars are designed this way.



Type 3: Theory of mind. Theory of mind is a psychology term. When applied to AI, it means that the system would have the social intelligence to understand emotions. This type of AI will be able to infer human intentions and predict behavior, a necessary skill for AI systems to become integral members of human teams.



Type 4: Self-awareness. In this category, AI systems have a sense of self, which gives them consciousness. Machines with self-awareness understand their own current state. This type of AI does not yet exist.



What are examples of AI technology and how is it used today?

AI is incorporated into a variety of different types of technology.

Here are six examples:

Automation. When paired with AI technologies, automation tools can expand the volume and types of tasks performed. An example is robotic process automation (RPA), a type of software that automates repetitive, rules-based data processing tasks traditionally done by humans. When combined with machine learning and emerging AI tools, RPA can automate bigger portions of enterprise jobs, enabling RPA's tactical bots to pass along intelligence from AI and respond to process changes.

Machine learning. This is the science of getting a computer to act without programming. Deep learning is a subset of machine learning that, in very simple terms, can be thought of as the automation of predictive analytics. There are three types of machine learning algorithms:

Supervised learning. Data sets are labeled so that patterns can be detected and used to label new data sets.

Unsupervised learning. Data sets aren't labeled and are sorted according to similarities or differences.

Reinforcement learning. Data sets aren't labeled but, after performing an action or several actions, the AI system is given feedback.

Machine vision. This technology gives a machine the ability to see. Machine vision captures and analyzes visual information using a camera, analog-to-digital conversion and digital signal processing. It is often compared to human eyesight, but machine vision isn't bound by biology and can be programmed to see through walls, for example. It is used in a range of applications from signature identification to medical image analysis. Computer vision, which is focused on machine-based image processing, is often conflated with machine vision.

Natural language processing (NLP). This is the processing of human language by a computer program. One of the older and best-known examples of NLP is spam detection, which looks at the subject line and text of an email and decides if it's junk. Current approaches to NLP are based on machine learning. NLP tasks include text translation, sentiment analysis and speech recognition.

Robotics. This field of engineering focuses on the design and manufacturing of robots. Robots are often used to perform tasks that are difficult for humans to perform or perform consistently. For example, robots are used in assembly lines for car production or by NASA to move large objects in space.

Researchers are also using machine learning to build robots that can interact in social settings.

Self-driving cars. Autonomous vehicles use a combination of computer vision, image recognition and deep learning to build automated skill at piloting a vehicle while staying in a given lane and avoiding unexpected obstructions, such as pedestrians.