



# AI Facts vs Myths

Myth	Fact
------	------



Artificial intelligence (AI), machine learning (ML), and deep learning (DL) are all the same thing

AI solves a task usually requiring human intelligence. ML solves a specific AI-task by learning from data, making it a strict subset of AI. DL solves an ML-problem by using neural networks (NNs) as its algorithm. Once again, we have a strict subset relationship, as can be seen in the illustration.



AI is a new technology

The term AI was coined in the 1950s by computer scientist John McCarthy, and many modern AI algorithms and methods have been around for decades. AI is getting considerable attention today since increases in data availability and computing power have made it significantly more effective at solving various problems, such as image recognition. For instance, in the mid-2010s, the AI-enabled system AlphaGo received international recognition for beating the world's best players in Go, a game that was thought to be too complex for a computer to succeed at relative to the top human players.



AI is on the verge of automating all jobs.

Similar to many past technologies, AI will automate some jobs while creating new ones. Currently, there are many jobs that AI cannot automate, such as those involving social and creative intelligence. For instance, an AI system can help diagnose a patient's illness, but it cannot respond to their questions or respond to negative reactions with human empathy. Jobs that complement AI programming will also be created. For instance, AI systems rely on considerable amounts of data, so there will be demand for people that find, generate, and clean this data.



AI systems are all blackboxes and cannot be regulated.

Some AI systems are simple and relatively easy to comprehend, while others employ complex computer programming. Techniques are under development to help better explain and interpret different AI systems and algorithms. There are various approaches to regulating these AI systems. One approach is to require a certain level of interpretability and explainability for how an AI reached a particular conclusion. Another method is to judge an AI system based on outcomes rather than how it comes to a decision. Other approaches include third-party auditing and impact assessments before deploying an AI system. These regulatory approaches are not mutually exclusive, and policymakers should decide which to use based on the specific use case and risk factors.



Modern AI is superintelligence.

Researcher Nick Bostrom defines superintelligence as "any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest." AI systems today are not close to meeting this definition. Modern AI is generally good at accomplishing a single task, such as recognizing images but does not have vast domain expertise in multiple fields.



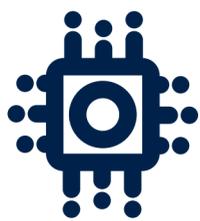
AI beating a world champion at (insert boardgame) is proof that we are close to superintelligence.

AI systems have defeated the best human players in complex strategy boardgames like chess and Go. However, this is not proof that AI is nearing superintelligence. These boardgames provide the AI with perfect information and static rules, while the real world is more dynamic and uncertain. An AI's success in one domain does not necessarily make it good at others.



AI has massive computational power.

AI depends on the host system's computational processing power. The more processing power a system has, the more an AI can execute its programming. But processing power is finite and expensive. Advanced deep neural networks become cost-prohibitive or financially unsustainable. This leads to a reality check among AI practitioners; at some point, computational power is not available to do the more advanced AI experiments. The computational advancements emphasized by Moore's Law—the observation and forecast by Intel co-founder Gordon Moore that computing power doubles roughly every two years—show signs of slowing down, making it harder for AI programmers to create more sophisticated AI tools using existing AI algorithms.



AI is inherently bias free, since computers are not biased.

AI can become biased based on the data and assumptions that go into it. For instance, an AI system trained with data reflecting historical biases in hiring decisions will adopt those biases. If not properly addressed, AI bias can perpetuate and accentuate existing societal inequities and harm vulnerable populations.



Nothing can be done about AI bias, because it is inherent to the technology.

AI systems can become biased based on data inputs and assumptions programmers make. Awareness about bias can help people address it through methods such as testing, improving data quality, and having a diversity of viewpoints represented when designing, evaluating, and deploying an AI system. More academic research into the causes of bias and bias mitigating strategies can help further identify ways to reduce bias.



AI is good at proving causality.

Correlation does not mean causation. AI is very good at recognizing patterns and finding correlations in data, but it is not good at proving causality. For instance, an AI system may identify that smoking and cancer are correlated, but it will have difficulty inferring that smoking causes cancers. Designing AI that can make causal inferences is an area of active interest among AI researchers.



AI can learn on its own without any human help.

AI systems are not self-aware. They require humans to design them, give them tasks, and generally provide them with data to "learn." Data scientists spend significant time structuring and cleaning data to ensure it is usable in an AI system.



AI works just like a human brain.

Neuroscientists still don't know all the intricacies of the human brain, and neither do computer scientists. While artificial neural networks (a subset of AI) were partly inspired by how the human brain works, the functioning of the human brain is currently more complex in many ways and more of a mystery. Another significant difference is that AI runs on computer hardware, while the human brain is a biological system.



AI can predict the future with certainty.

AI algorithms can be used for forecast analysis. However, according to Princeton professors Arvind Narayanan and Matt Salganikto, "even a cursory review of the literature reveals that state-of-the-art predictive accuracies fall well short of expectations." AI systems often fall short due to the limits of modeling and data. They are especially bad at predicting unusual or novel events, such as how a new pandemic would affect a modern economy.