



# 5 Big Myths

## of AI and Machine Learning Debunked

Demystify the role and potential of artificial intelligence (AI) and machine learning (ML) in powering positive outcomes for the enterprise.

# It's hard to believe that the now ubiquitous terms of AI and ML date back to the 1950s.

Since then, public interest in AI and ML has waxed and waned. But the release of OpenAI's ChatGPT in late 2022, and its competitors that followed, brought generative AI into the mainstream.

Now it's more powerful and easier to use than ever before, with use cases that span across both consumer and enterprise. Generative AI can plan your next vacation, write poetry in the style of William Shatner or polish the speech you need to write for your best friend's wedding. It can also write policies for a cyber insurance application, generate code or explain the meaning of a high error in your API service.

Over the next decade, relying on AI in business will be as quotidian as flipping on a lightswitch. Gartner®'s Hype Cycle™ for Emerging Technologies 2023<sup>1</sup> claims that generative AI will reach the peak of inflated expectations in 2023. However, AI's long-term potential is underestimated, says Hao Yang, Splunk's VP of AI. Generative AI alone represents an annual \$2.6 to \$4.4 trillion in opportunity across 63 use cases, [according to a McKinsey study](#). Spending will only continue to grow: [IDC predicts](#) that worldwide spending on AI will surpass \$300 billion in 2026.

Now the ROI of those investments — a historically painful metric to quantify for AI projects — is starting to crystallize as organizations expand to more sophisticated use cases. [In a PwC survey](#), 72% of AI-mature organizations surveyed (and 59% of all other respondents) are confident in their abilities to assess the ROI of their current AI initiatives, with the ability to capture both hard and soft returns and costs.

Despite the promise of AI, an overall mistrust lingers. Fifty-two percent of organizations say that risk factors are a critical consideration when evaluating new AI use cases, [according to Gartner](#).<sup>2</sup> The old adage “you

can't protect what you don't know” rings true for AI and ML. Generative AI in particular raises deep-seated data privacy and security concerns. These AI anxieties are only natural as organizations navigate this new era of technology. The [Biden Administration's executive order on AI](#) provides a few answers, but substantive regulatory changes are likely still a ways out.

In the meantime, the AI train has left the station, with positive outcomes that are too difficult to ignore. Usage is already widespread, with 55% of respondents in the McKinsey study reporting that their organizations have adopted AI. And organizations are realizing the value of this adoption through improvements in productivity, decision-making, customer experience, innovation and beyond.

To be sure, some of AI's most far-reaching concepts (computers that can replicate the human brain entirely, fully autonomous robots and programs that design, code and upgrade themselves) are years away from reality; they're still moonshots that represent the eventual apex of AI's capabilities. But considering that AI tools can already win at Jeopardy!, are able to detect breast cancer and are logging tens of thousands of miles behind the wheel of self-driving vehicles every day, the prospect of even those moonshot concepts really doesn't seem so far-fetched.

In other words, now is the time to learn about AI and ML. Before you can develop a thoughtful strategy that considers the risks and benefits, it's important to clarify these common misconceptions.

<sup>1</sup> Gartner, [Hype Cycle for Emerging Technologies](#), Arun Chandrasekaran, Melissa Davis  
2 August 2023.

<sup>2</sup> Gartner Press Release, [Gartner Survey Finds 55% of Organizations That Have Deployed AI Take an AI-First Strategy with New Use Cases](#), July 27, 2023.

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally, and HYPE CYCLE is a registered trademark of Gartner, Inc. and/or its affiliates and are used herein with permission. All rights reserved.

## Myth 1

# AI, ML and generative AI are all the same.

As with many new technologies, artificial intelligence has created a gold rush effect across many industries.

All types of products have been described as having been built with AI, to the point where the term has become a buzzword that has seemingly lost much of its meaning. So let's try to get some of that meaning back by breaking down the term "AI" to understand what it really means.

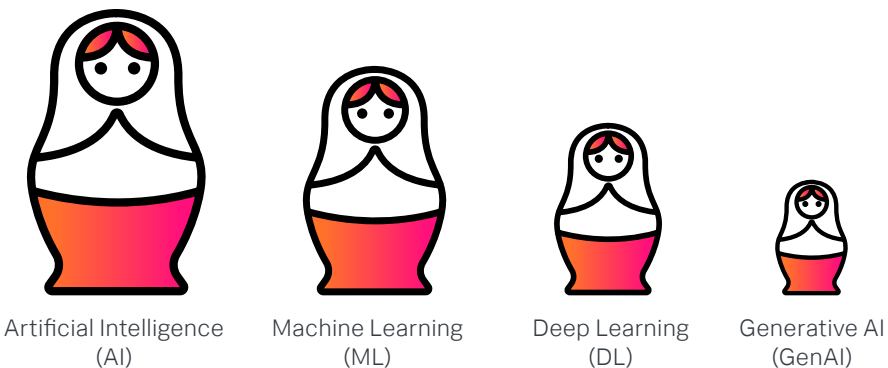
At its simplest level, AI can be split into two categories: narrow artificial intelligence (ANI) and general artificial intelligence (GAI). The names have evolved in recent years, but the terms can generally be thought of in the following ways.

### Narrow AI (ANI)

Narrow AI is a collection of technologies that rely on algorithms and programmatic responses to simulate intelligence, generally with a focus on a specific task. When you use a voice recognition system like Amazon's Alexa to turn on the lights, that's narrow AI in action. Alexa may sound smart, but it doesn't have any advanced understanding of language and can't determine the meaning behind the words you speak. The program simply listens for key sounds in your speech and, when it detects them, follows its programming to execute certain actions. To users, this can seem surprisingly intelligent — and voice recognition is far from a simple computing task — but in reality, there is no actual "thinking" going on behind the scenes. Non-player characters (NPCs) in games are another good example of ANI. While they take human-like action, they're simply following a pre-programmed series of actions designed to mimic how a human would play the game.

### General artificial intelligence (GAI)

GAI, in contrast, is intended to think on its own. The goal of GAI research is to engineer AI with learning capabilities that match or surpass human intelligence. GAI is designed to learn and adapt, to make a decision tomorrow that is better than the one it made today. None of this is easy, which is why most examples of AI you'll encounter today are the narrow form. GAI is a new, complex and varied category with numerous sub-branches, most of which are still research topics in a lab. Modern AI systems focus on solving specific tasks, such as optimization, recommendation or prediction systems and don't learn broad concepts like a human would.



**Like nesting dolls, generative AI sits within deep learning, a subset of machine learning, which sits within artificial intelligence.**

## With this distinction in mind, what is machine learning?

Put simply, machine learning is a specific type of ANI, with the goal of giving a computing device access to some store of data it can learn from — but nowhere near GAI levels. As stated above, when Alexa turns on the lights, it doesn't actually learn anything, even if it does better recognize your voice over time. When the user tells Alexa to turn off the lights, a program is then executed to carry out the command. This is an example of a rules-based approach, where the system operates on rules coded through if-then-else statements — one of the simplest forms of AI.

But ML systems can do more, too, when given a data feed (for example, temperature and tolerance information from sensors on a piece of manufacturing equipment). These systems are asked to draw conclusions about the data based on observed examples of a task. This may involve searching that data for trends, patterns and anomalies or any information that might not be obvious to a human observer. In the case of manufacturing, the machine learning algorithm would learn to send proactive alerts when temperatures exceed a certain threshold, allowing operators to take action before an issue arises.

How is all this possible? ML is a subset of AI that includes supervised, unsupervised, reinforcement and deep learning systems. Supervised machine learning algorithms and models use labeled datasets, beginning with an understanding of how the data is classified. Unsupervised models use unlabeled datasets and figure out features and patterns from the data without explicit instructions or preexisting categorizations.

Reinforcement learning, on the other hand, takes a more iterative approach. Instead of being trained with a single data set, the system learns through trial and error and feedback from data analysis.


With faster and bigger computation capabilities, ML capabilities have advanced to deep learning. Deep learning is a specific kind of ML that applies algorithms called “artificial neural networks,” composed of decision nodes to more accurately train ML systems for supervised, unsupervised and reinforcement learning tasks. Deep learning approaches are becoming more widespread, but come with high computation costs and are often harder for humans to interpret because the decision nodes are “hidden” and not exposed to the developer. Nonetheless, deep learning offers a wealth of possibilities and already has promising applications for image recognition, self-driving cars, fraud news detection and more.

## So, what is generative AI?

Generative AI, also known as GenAI, broadly falls under the category of machine learning. It simply refers to algorithms that can create content, including text, imagery, video, simulations, code, audio and more. Examples of generative AI include tools such as ChatGPT, DALL-E and Google Bard.

Chatbots that rely on generative AI have existed since the 1960s, but the introduction of the generative adversarial network (GAN) in 2014 enabled extensive innovation in GenAI. A GAN is a machine learning model composed of two neural networks: a generator and a discriminator. The former is trained to produce false data, while the latter is trained to distinguish between the false data and real examples. These two networks enter a feedback loop that enables the generator to produce more believable output. With this technology, generative AI can create incredibly realistic content — like a painting of your dog in the style of Alphonse Mucha, for example.

On the whole, the AI and ML space is constantly evolving. The important thing is understanding that these techniques can be applied to solve business problems, as long as there is data to train them.



## Myth 2

# AI is a magic wand.

As exciting as an AI-enabled hair dryer and AI-powered yoga pants sound (yes, those are real things), there is a time and a place for AI, at least as it stands today.

Data is the key to successful outcomes using AI, regardless of the level of complexity or the industry in which it's deployed.

AI would not exist without proper training, and models need to be well-built to be useful. A spam filter must be trained on how to recognize a good email message from a bad one. A voice recognition AI assistant has to listen to countless hours of spoken dialogue before it can parse what is being said with any degree of accuracy. AI-enabled factory floor initiatives are typically required to analyze several million gigabytes of data each week to make reasoned forecasts of what might happen in the future.

These are all examples of training. It's not just a game of volume, but one of quality, too. We've all heard the saying "garbage in, garbage out" and successful AI algorithms must be trained on the right data, or they simply won't be able to inform good decisions. If you were to open your email and tag all the messages from your spouse as spam, then tag all the misspelled emails from your bank as good, you'll see for yourself how quickly AI can go off the rails.

The same is true for a more advanced industrial setting. If a sensor is miscalibrated and feeds inaccurate information to an algorithm tasked with monitoring equipment, gigabytes of data will end up being useless. Or worse — the AI model will be trained with bad data, reach inaccurate conclusions and give false or misleading recommendations.

The point of this is that AI is not necessarily a cure-all. There is no "AI switch" or "AI plugin" that can take any old technology and somehow give it cognitive ability.

We still need humans in the loop. Humans have to define the problem, identify an appropriate AI technology to solve it, train the tool with the correct data and then verify that the results are valid. Even the most powerful AI tools developed to date have to be carefully managed over time so they don't run off the rails.

Many AI professionals are finding that they learn more when an AI algorithm returns the wrong answer instead of the right one. This effect is visible at both the consumer and the industrial level.

### Once an AI tool has generated results, the work isn't over.

When an AI-based spam filter miscategorizes an incoming message, the user has the chance to retrain the tool by categorizing it properly. This gives the algorithm new insight into what it might have missed the first time around, and learning from the error makes the tool incrementally more powerful. If the spam filter had not been retrained, it would be no more accurate the next time around and would likely make the same mistake again. Similarly, in a manufacturing setting, imagine that an AI tool directs that a machine be taken offline because a failure in a key part is imminent. If the part does not fail, then what?

What happens if security-focused AI blocks your traveling sales force from accessing the network because it wrongly assumed they were hackers? Because of the logical nature of AI, a developer can determine why the AI made these specific decisions and can work backwards to determine what data it relied on in the process. This may reveal flaws in the data, an error of logic processing or some other bug that would otherwise go unnoticed. In situations like these, developers can take necessary actions to correct the code, retrain the models or do whatever it takes once these flaws are identified.

What is critical is that humans are in the loop, because AI is still just a tool — yes, a really advanced tool — not a magic wand.



Any sufficiently advanced technology is indistinguishable from magic.

Arthur C. Clarke

## Myth 3

# You need a Ph.D to benefit from AI and ML.

These are dazzlingly complex technologies that, under the hood, are far from accessible to the layperson. While building your own models may require a team and doctorate, embedded AI can help anyone.

You don't need an advanced degree to benefit from the applications of AI and ML and there is no use case too small for smart technology.

Nonetheless, it's important to understand the difference between building an artificial intelligence solution from the ground up and implementing existing AI tools within your organization. The first of these is extremely difficult. The second is getting easier every day. Consider all of the technological tools you use in the course of a day: an email client, or productivity tools like your digital assistant or spreadsheets. They're not simple technologies, but you're able to master them without knowing what's happening under the hood.

The same thing is happening to AI, as tools are becoming increasingly accessible. There's been an uptick in the breadth and quality of self-service analytics platforms, enabling non-technical employees to output analysis without relying purely on data scientists. Beginners can build their own machine learning models that have pre-built algorithms and intuitive interfaces, and are aimed at developers who don't necessarily have robust backgrounds in data science or statistics.

Today, you often don't even have to build models to gain value from embedded AI. There's a chance that AI and ML may be under the hood powering the features or functionality you're already using.

Google or ChatGPT require absolutely zero data science knowledge to gain results. They may be great at gathering search results or generating poetry, but what if you want to make your manufacturing process or incident response more efficient? Chances are they're not going to be able to help you. AI can still be extremely effective at those tasks and money is pouring into the so-called democratization of AI, with the launch of open-source and commercial tools from legacy and emerging tech companies. Some technologies have been tailored for specific tasks like helping catch fraudsters with behavioral biometrics. While we might not be at the level of pure point-and-click when it comes to implementing an AI system, these tools are now accessible to anyone willing to learn.

You may also be under the impression that your business or specific use case is too small or insignificant to merit an investment in AI, or that your environment is just not complex enough to benefit from the technology. In many cases, this is not true. Even small businesses and other environments limited in scope and scale can benefit from the results provided by AI and ML. In a small business or business unit setting, you might task AI with identifying anomalous security events, testing and developing better social media ads, automating and improving customer service requests or searching for patterns around when and why competitors are changing prices or product offerings. All of these AI-driven activities are readily accessible to even single-owner operations.

## Myth 4

# AI and ML will replace me.

Worried that AI will eliminate your job in the near future? You're not alone. [McKinsey recently suggested](#) that by 2030, an additional 12 million occupational transitions may be needed, as machines become increasingly capable of doing work previously reserved for humans.

But sometimes buried in these reports are a few key details that should calm the masses.

There is a fundamental truth to AI: it can't operate in a vacuum. It requires humans in the loop to develop, deploy, manage and maintain it. That means jobs. McKinsey predicts that generative AI will enhance the way STEM, creative, and business and legal professionals work rather than eliminating a significant number of jobs outright. And according to our [CISO report](#), 86% of CISOs believe that genAI will alleviate skills gaps and talent shortages on their security teams.

### Why won't AI decimate employment as we know it?

To date, all AI is narrow AI — and most say that GAI won't be reality anytime soon. In many cases, an AI and a human work hand in hand: The AI scours the data, looking for details that would otherwise take months of time to uncover, and the human checks that the AI's results are on target. As problems get increasingly difficult, the need for a human in the loop becomes even greater, not smaller.

Additionally, AI doesn't have innate knowledge of business strategy, process or implementation. As Greek philosopher Heraclitus said, "change is the only constant," and it's especially true when it comes to companies. Whether it's a strategy shift or migrating business processes to more flexible systems, these transitions are often lengthy and require significant human intervention and soft skills to successfully implement. Additionally, a substantial amount of institutional knowledge about the business, industry and competitive environment are needed. Only through a collaborative effort can an AI tool effectively recommend process automation and reconfiguration activities — all of which ensure that people will remain a key part of the puzzle for the long haul.

For all of its intelligence, AI isn't always right. AI can not only go wrong — it can go catastrophically wrong. When that happens, a human (with intuition, experience and the ability to react quickly) is invariably needed to overcome the problem or pick up the pieces. What's worse is that AI tools often don't even know they have made a mistake, which requires even more work from a human who can figure out how to prevent them.

## Myth 5

# You need to have all the data and create your own models to take advantage of AI and ML.


Spoiler: You don't need terabytes worth of data or teams of data scientists to take advantage of AI and ML.

While it's true that AI as a rule thrives on a large and accurate pool of data, your business doesn't need to have all of that data in house for AI to be useful.

A tool that monitors and analyzes social media collects its data from external sources as it goes. An AI system that relies on data feeds such as ambient temperature, housing prices or neighborhood demographics typically pulls all of this information from publicly available sources. There's really no need that's "too small." Remember that a tiny improvement in a key business vector can have a huge impact on the bottom line. A system that reduces production mistakes by a fraction of a percent or correctly recommends a price increase of just a few pennies could equate to millions of dollars in avoided costs or additional profits. The challenge is largely in identifying these opportunities.

For in-house data, it's also possible for organizations to train models with unstructured data. Unstructured data tends to be qualitative data (think audio, video, surveillance imagery and so on) whereas structured data is highly organized and is more or less the same as quantitative data. But businesses can indeed analyze structured and unstructured data, thanks to machine learning algorithms such as pattern classification, text-mining and natural language processing, and thanks to lower costs for computing power and hard drive space as well.

If you're thinking of using generative AI, chances are it probably doesn't make sense from a cost (millions of dollars) or time perspective to try and create a model from scratch that will compete with GPT4. But that doesn't mean you can't take advantage of the vast strides being made in AI. With today's options, you can take advantage of the burgeoning marketplace of AI offerings and look for a domain-specific solution that suits your needs, or look for something pre-built that you or your team can tailor.



**have already achieved either moderate or substantial benefits from their work with these technologies.**



# Truth: AI is here to stay.

While AI can be a game changer that takes your business to the next level, taking your first steps with AI and machine learning does not have to be a monumental undertaking. With AI now embedded within enterprise software, you can gain the benefits and efficiencies that AI can drive without having to build it yourself.

If you want a specific tailored solution, numerous tools on the market let you experiment with AI in a sandbox, targeting small “problem areas” that might have long stymied your attempts at improvement. The important thing is that you need to get started soon, so that you can gain the value driven by AI before your competition jumps ahead of you in the marketplace.

AI is already having a profound impact on the bottom line of businesses that were early movers into the field. Companies see improvements in customer satisfaction, decreases in manufacturing downtime and better overall employee productivity. There’s no blanket AI tool and no single metric that will improve once these tools are implemented, but it’s up to you to determine where to target AI based on the specific challenges you see in your organization.

[Learn more](#)

Of course, getting there means overcoming some hurdles. You may have to educate nervous staff members about the realities of AI and job displacement. Turn those fears around by showcasing how AI can improve their work lives, free up their time to work on more interesting initiatives, and brighten their future career prospects. Smaller businesses may also need to overcome the sentiment that AI is a game that only the largest of enterprises can play, which is where some targeted pilot projects can really help.

While AI is already showcasing real-world results, the future of these tools is even more exciting. It’s a journey, however, that you need to begin today.