

In-Depth Guide to Machine Learning in the Enterprise



In this e-guide

- [In-depth guide to machine learning in the enterprise](#) p. 2
- [Enterprise use cases and benefits](#) p. 3
- [Challenges](#) p. 7
- [Implementation: 6 Steps](#) p. 10
- [Management and maintenance of ML](#) p. 14
- [Future Trends](#) p. 15
- [Vendors and platforms](#) p. 18
- [Further reading](#) p. 21

In this e-guide:

Enterprises are adopting machine learning technologies at rapid rates. In this machine learning guide, we break down what you need to know about this transformative technology.

In this e-guide

- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21

In-depth guide to machine learning in the enterprise

Ed Burns, Executive Editor, SearchEnterpriseAI

Machine learning for enterprise use is exploding. From improving customer experience to developing products, there's almost no area of the modern business untouched by machine learning.

Machine learning is a pathway to creating **artificial intelligence**, which in turn is one of the primary drivers of machine learning use in the enterprise. There is some disagreement over the exact nature of the relationship between AI and machine learning. Some see machine learning as a subfield of AI, while others view AI essentially as a subfield of machine learning. In general, AI aims to replicate some aspect of human perception or decision-making, whereas machine learning can be used to enhance or automate virtually any task, not just ones related to human cognition. However you view them, the two concepts are closely linked, and they are feeding off each other's popularity.

The practice of machine learning involves taking data, **examining it for patterns** and developing some sort of prediction about future outcomes. By feeding an algorithm more data over time, data scientists can sharpen the machine learning model's

In this e-guide

- [In-depth guide to machine learning in the enterprise](#) p. 2
- [Enterprise use cases and benefits](#) p. 3
- [Challenges](#) p. 7
- [Implementation: 6 Steps](#) p. 10
- [Management and maintenance of ML](#) p. 14
- [Future Trends](#) p. 15
- [Vendors and platforms](#) p. 18
- [Further reading](#) p. 21

predictions. From this basic concept, a number of different [types of machine learning](#) have developed:

- **Supervised machine learning.** The most common form of machine learning, [supervised learning](#) involves feeding an algorithm large amounts of labeled training data and asking it to make predictions on never-before-seen data based on the correlations it learns from the labeled data.
- **Unsupervised learning.** [Unsupervised learning](#) is often used in the more advanced applications of artificial intelligence. It involves giving unlabeled training data to an algorithm and asking it to pick up whatever associations it can on its own. Unsupervised learning is popular in applications of clustering (the act of uncovering groups within data) and association (predicting rules that describe data).
- **Semisupervised learning.** In semi-supervised learning, algorithms train on small sets of labeled data and then, as in unsupervised learning, apply their learnings to unlabeled data. This approach is often used when there is a lack of quality data.
- **Reinforcement learning.** [Reinforcement learning](#) algorithms receive a set of instructions and guidelines and then make their own decisions about how to handle a task through a process of trial and error. Decisions are either rewarded or punished as a means of guiding the AI to the optimal solution to the problem.

From these four main types of machine learning, enterprises have developed an [impressive array of techniques and applications](#). Everything from relatively simple sales forecasting to today's most cutting-edge AI tools run on machine learning models. This guide to machine learning in the enterprise explores the variety of use cases for

In this e-guide

- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21

machine learning, the challenges to adoption, how to implement machine learning technologies and much more.

Enterprise use cases and benefits

Machine learning for enterprise use is accelerating, and not just at the periphery. Increasingly, businesses are putting machine learning applications at the center of their business models. The technology has enabled businesses to perform tasks at a scale previously unachievable, not only generating efficiencies for companies but also new business opportunities, as technology writer Mary Pratt explained in "[10 common uses for machine learning in business](#)." The growing use of machine learning in mission-critical business processes is reflected in the range of use cases where it plays an integral role. The following are examples:

- **Recommendation engines.** Most prominent online, consumer-facing companies today use [recommendation engines](#) to get the right product in front of their customer at the right time. Online retail giant Amazon pioneered this technology in the early part of the last decade, and it has since become standard technology for online shopping sites. These tools consider the browsing history of customers over time and match the preferences described by that history to other products the customer might not be aware of yet.
- **Fraud detection.** As more financial transactions move online, the opportunity for fraud has never been greater. That makes the need for fraud detection paramount. Credit card companies, banks and retailers are increasingly using machine learning applications to weed out likely cases of fraud. At a very basic

In this e-guide

- ▶ [In-depth guide to machine learning in the enterprise](#) p. 2
- ▶ [Enterprise use cases and benefits](#) p. 3
- ▶ [Challenges](#) p. 7
- ▶ [Implementation: 6 Steps](#) p. 10
- ▶ [Management and maintenance of ML](#) p. 14
- ▶ [Future Trends](#) p. 15
- ▶ [Vendors and platforms](#) p. 18
- ▶ [Further reading](#) p. 21

level, these applications work by learning the characteristics of legitimate transactions and then scanning incoming transactions for characteristics that deviate. The tool then flags these transactions.

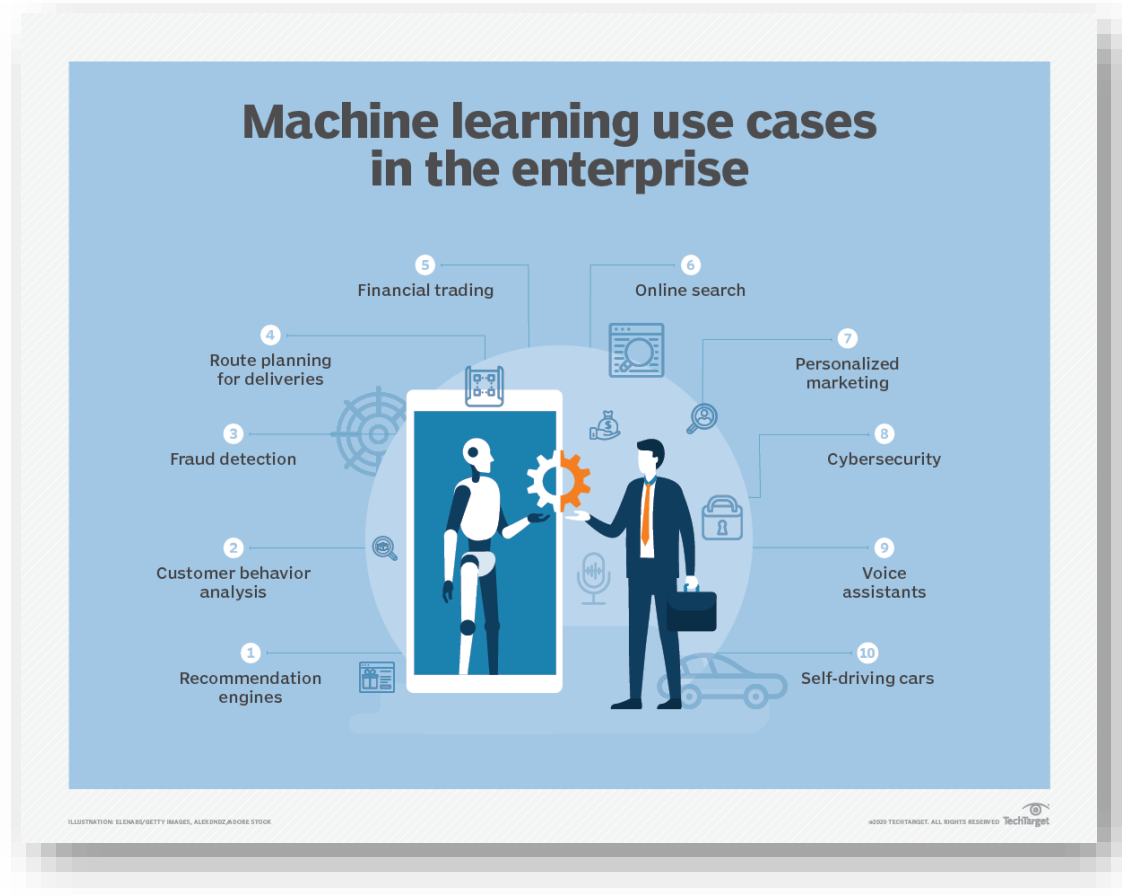
- **Customer analysis.** Most businesses today collect vast stores of data on their customers. This so-called big data includes everything from browsing history to social media activity. It's far too voluminous and diverse for humans to make sense of on their own. That's where machine learning comes in. Algorithms can troll the [data lakes](#) where enterprises store the raw data and develop insights about customers. Machine learning can even develop personalized marketing strategies that target individual customers and inform strategies for improving customer experience.
- **Financial trading.** Wall Street was one of the earliest adopters of machine learning technology, and the reason is clear: In a high-stakes world where billions of dollars are on the line, any edge is valuable. Machine learning algorithms are able to examine historical data sets, find patterns in stock performance and make predictions about how certain stocks are likely to perform in the future.
- **Virtual assistants.** By now, most people are familiar with virtual assistants from tech companies like Apple and Google. What they might not know is the extent to which machine learning powers these bots. Machine learning enters in a number of different ways, including [deep learning](#), a machine learning technique based on [neural networks](#). Deep learning plays an important role in developing natural language processing, which is how the bot is able to interact with the user, and in learning the user's preferences.
- **Self-driving cars.** This is where machine learning enters the realm of AI that aims to be on par with human intelligence. Autonomous vehicles use neural networks to learn to interpret objects detected by their cameras and other sensors, and to determine what action to take to move a vehicle down the

In this e-guide

- ▶ [In-depth guide to machine learning in the enterprise](#) p. 2
- ▶ [Enterprise use cases and benefits](#) p. 3
- ▶ [Challenges](#) p. 7
- ▶ [Implementation: 6 Steps](#) p. 10
- ▶ [Management and maintenance of ML](#) p. 14
- ▶ [Future Trends](#) p. 15
- ▶ [Vendors and platforms](#) p. 18
- ▶ [Further reading](#) p. 21

road. In this way, machine learning algorithms can use data to come close to replicating human-like perception and decision-making.

These are just some examples, but there are countless more. Any business process that either produces or uses large amounts of data -- particularly structured, [labeled data](#) -- is ripe for automation that uses machine learning. Enterprises across all industries have learned this and are working to implement machine learning methods throughout their processes.



In this e-guide

 In-depth guide to machine learning in the enterprise	p. 2
 Enterprise use cases and benefits	p. 3
 Challenges	p. 7
 Implementation: 6 Steps	p. 10
 Management and maintenance of ML	p. 14
 Future Trends	p. 15
 Vendors and platforms	p. 18
 Further reading	p. 21

It's not hard to see why machine learning has entered so many situations. Enterprises that have adopted machine learning are solving business problems and reaping value from this AI technique. [Here are six business benefits:](#)

- increased productivity;
- lower labor costs;
- better financial forecasting;
- clearer understanding of customers;
- fewer repetitive tasks for workers; and
- more advanced and human-like output.

Challenges

The question is no longer whether to use machine learning, it's [how to operationalize machine learning](#) in ways that return optimal results. That's where things get tricky.

Machine learning is a complicated technology that requires substantial expertise. Unlike some other technology domains, where software is mostly plug and play, machine learning forces the user to think about why they are using it, who is building the tools, what their assumptions are and how the technology is being applied. There are few other technologies that have so many potential points of failure.

The wrong use case is the downfall of many machine learning applications. Sometimes enterprises lead with the technology, looking for ways to implement machine learning,

In this e-guide

- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21

rather than allowing the problem to dictate the solution. When machine learning is shoehorned into a use case, it often fails to deliver results.

The wrong data dooms machine learning models faster than anything. Data is the lifeblood of machine learning. Models only know what they've been shown, so when the data they train on is inaccurate, unorganized or biased in some way, the model's output will be faulty.

Bias frequently hampers machine learning implementations. The many types of **bias** that can undermine machine implementations generally fall into the two categories. One type happens when data collected to train the algorithm simply doesn't reflect the real world. The data set is inaccurate, incomplete or not diverse enough. Another type of bias stems from the methods used to sample, aggregate, filter and enhance that data. In both cases, the errors can stem from the biases of the data scientists overseeing the training and result in models that are inaccurate and, worse, unfairly affect specific populations of people. In his article "[6 ways to reduce different types of bias in machine learning](#)," analyst Ron Schmelzer explained the types of biases that can derail machine learning projects and how to mitigate them.

Black box functionality is one reason why bias is so prevalent in machine learning. Many types of machine learning algorithms -- particularly unsupervised algorithms -- operate in ways that are opaque, or as a "**black box**," to the developer. A data scientist feeds the algorithm data, the algorithm makes observations of correlations and then produces some sort of output based on these observations. But most models can't

In this e-guide

- [In-depth guide to machine learning in the enterprise](#) p. 2
- [Enterprise use cases and benefits](#) p. 3
- [Challenges](#) p. 7
- [Implementation: 6 Steps](#) p. 10
- [Management and maintenance of ML](#) p. 14
- [Future Trends](#) p. 15
- [Vendors and platforms](#) p. 18
- [Further reading](#) p. 21

explain to the data scientist why they produce the outputs they do. This makes it extremely difficult to detect instances of bias or other failures of the model.

Technical complexity is one of the biggest challenges to enterprise use of machine learning. The basic concept of feeding training data to an algorithm and letting it learn the characteristics of the data set may sound simple enough. But there is a lot of technical complexity under the hood. Algorithms are built around advanced mathematical concepts, and the code that algorithms run on can be difficult to learn. Not all businesses have the technical expertise in house needed to develop effective machine learning applications.

Lack of generalizability prevents machine learning from scaling to new use cases in most enterprises. Machine learning applications only know what they've been explicitly trained on. This means a model can't take something it learned about one area and apply it to another, the way a human would be able to. Algorithms need to be trained from scratch for every new use case.

➤ **Image on the next page**

In this e-guide

- ▶ [In-depth guide to machine learning in the enterprise](#) p. 2

- ▶ [Enterprise use cases and benefits](#) p. 3

- ▶ [Challenges](#) p. 7

- ▶ [Implementation: 6 Steps](#) p. 10

- ▶ [Management and maintenance of ML](#) p. 14

- ▶ [Future Trends](#) p. 15

- ▶ [Vendors and platforms](#) p. 18

- ▶ [Further reading](#) p. 21

How AI systems amplify bias

Image recognition systems that use biased machine learning data sets will inadvertently magnify that bias. Researchers are examining ways to reduce the effects.

				
COOKING	COOKING	COOKING	COOKING	COOKING
ROLE VALUE	ROLE VALUE	ROLE VALUE	ROLE VALUE	ROLE VALUE
AGENT ▶ WOMAN	AGENT ▶ WOMAN	AGENT ▶ WOMAN	AGENT ▶ WOMAN	AGENT ▶ MAN
FOOD ▶ PASTA	FOOD ▶ FRUIT	FOOD ▶ MEAT	FOOD ▶ VEGETABLES	FOOD ▶ —
HEAT ▶ STOVE	HEAT ▶ —	HEAT ▶ GRILL	HEAT ▶ STOVE	HEAT ▶ STOVE
TOOL ▶ SPATULA	TOOL ▶ KNIFE	TOOL ▶ TONGS	TOOL ▶ TONGS	TOOL ▶ SPATULA
PLACE ▶ KITCHEN	PLACE ▶ KITCHEN	PLACE ▶ OUTSIDE	PLACE ▶ KITCHEN	PLACE ▶ KITCHEN

In this example of gender bias, adapted from a report published by researchers from the University of Virginia and the University of Washington, a visual semantic role labeling system has learned to identify a person cooking as female, even when the image is male.

IMAGES (FROM LEFT): ALEXEY STORP/ADOBE STOCK, JACOB LUND/ADOBE STOCK, BERND JÜRGENS/ADOBE STOCK, GRANGER WOOTZ/GETTY IMAGES

In this e-guide

- 📖 [In-depth guide to machine learning in the enterprise](#) p. 2
- 📖 [Enterprise use cases and benefits](#) p. 3
- 📖 [Challenges](#) p. 7
- 📖 [Implementation: 6 Steps](#) p. 10
- 📖 [Management and maintenance of ML](#) p. 14
- 📖 [Future Trends](#) p. 15
- 📖 [Vendors and platforms](#) p. 18
- 📖 [Further reading](#) p. 21

Implementation: 6 steps

Implementing machine learning is a multistep process requiring input from many types of experts. Here is an outline of the process in six steps.

1. Any machine learning implementation starts with the **identification of a problem**. The most effective machine learning projects tackle specific, clearly defined business challenges or opportunities.
2. Following the problem formulation stage, data science teams should **choose their algorithm**. Different machine learning algorithms are better suited for different tasks, as explained in this article on "[9 types of machine learning algorithms](#)" by TechTarget editor Kassidy Kelley. Simple linear regression algorithms work well in any use case where the user seeks to predict one unknown variable based on another known variable. Cutting-edge deep learning algorithms are better at complicated things like image recognition or text generation. There are dozens of other types of algorithms that cover the space between these examples. Choosing the right one is essential to the success of machine learning projects.
3. Once the data science team identifies the problem and picks an algorithm, the next step is to **gather data**. The importance of collecting the right kind of and enough data is often underestimated, but it shouldn't be. Data is the lifeblood of machine learning. It supplies algorithms with everything they know, which in turn defines what they are capable of. Data collection involves complicated tasks like identifying data stores, writing scripts to connect databases to machine learning applications, verifying data, cleaning and labeling data and

In this e-guide

- ▶ [In-depth guide to machine learning in the enterprise](#) p. 2
- ▶ [Enterprise use cases and benefits](#) p. 3
- ▶ [Challenges](#) p. 7
- ▶ [Implementation: 6 Steps](#) p. 10
- ▶ [Management and maintenance of ML](#) p. 14
- ▶ [Future Trends](#) p. 15
- ▶ [Vendors and platforms](#) p. 18
- ▶ [Further reading](#) p. 21

organizing it in files for the algorithm to work on. While these are tedious and complicated jobs, their importance cannot be overstated.

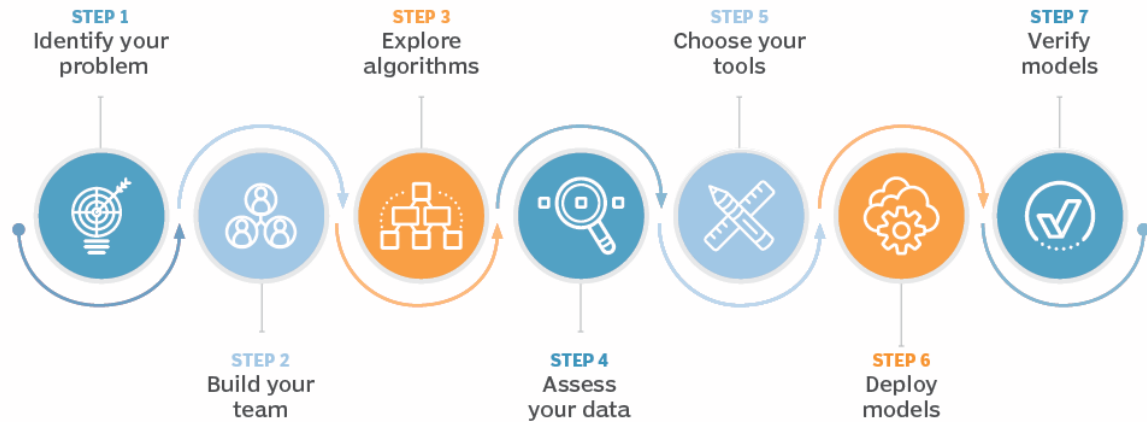
4. Now it's time for the magic to begin. Once the data science team has all the data it needs, it can start **building the model**. This step in the machine learning process will differ substantially depending on whether the team is using a supervised machine learning algorithm or an unsupervised algorithm. When the training is supervised, the team feeds the algorithm data and tells it what features to examine. In an unsupervised learning approach, the team essentially turns the algorithm loose on the data and comes back once the algorithm has produced a model of what the data looks like. Learn how to build a [neural network model](#) in this expert tip.
5. **Application development** is next. Now that the algorithm has developed a model of what the data looks like, data scientists and developers can build that learning into an application that addresses the business challenge or opportunity identified in the first step of the process. Sometimes this is very simple, like a data dashboard that updates sales projections based on changing economic conditions. It could be a recommendation engine that has learned to tailor its suggestions based on past customer behavior. Or it could be a component of cutting-edge medical software that uses image recognition technology to detect cancer cells in medical images. During the development stage, engineers will test the model against new, incoming data to make sure it delivers accurate predictions.
6. Even though the primary work is complete, now is not the time to walk away from the model. The last step in the machine learning process is **model validation**. Data scientists should verify that their application is delivering accurate predictions on an ongoing basis. If it is, there's likely little reason to make changes. However, model performance typically degrades over time. This

In this e-guide

- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21

is because the underlying facts that the model trained on -- whether economic conditions or customer tendencies -- shift as time goes by. When this happens, the performance of models gets worse. This is the time when data scientists need to retrain their models. Here, the whole process essentially starts over again.

Machine learning adoption roadmap



ICONS: ALEXDR102/ADOBE STOCK

©2020 TECHTARGET. ALL RIGHTS RESERVED TechTarget

In this e-guide

- [In-depth guide to machine learning in the enterprise](#) p. 2
- [Enterprise use cases and benefits](#) p. 3
- [Challenges](#) p. 7
- [Implementation: 6 Steps](#) p. 10
- [Management and maintenance of ML](#) p. 14
- [Future Trends](#) p. 15
- [Vendors and platforms](#) p. 18
- [Further reading](#) p. 21

Management and maintenance of ML

The management and maintenance of machine learning applications in the enterprise is one area that's sometimes given short shrift, but it can be what makes or breaks use cases.

The basic functionality of machine learning depends on models learning trends -- such as customer behavior, stock performance and inventory demand -- and projecting them to the future to inform decisions. However, underlying trends are constantly shifting, sometimes slightly, sometimes substantially. This is called *concept drift*, and if data scientists don't account for it in their models, the model's projections will eventually be off base.

The way to correct for this is to [never view models in production as finished](#). They demand a constant state of verification, retraining and reworking to ensure they continue to deliver results.

- **Verification.** Data scientists often will hold out a segment of new, incoming data and then verify the model's predictions to make sure they are close to the new, incoming data.
- **Retraining.** If a model's results start to deviate significantly from actual observed data, it's time to retrain the model. Data scientists will need to source a completely new set of data that reflects current conditions.
- **Rebuilding.** Sometimes the concept a machine learning model is supposed to predict will change so much that the underlying assumptions that went into the

In this e-guide

- 📖 [In-depth guide to machine learning in the enterprise](#) p. 2
- 📖 [Enterprise use cases and benefits](#) p. 3
- 📖 [Challenges](#) p. 7
- 📖 [Implementation: 6 Steps](#) p. 10
- 📖 [Management and maintenance of ML](#) p. 14
- 📖 [Future Trends](#) p. 15
- 📖 [Vendors and platforms](#) p. 18
- 📖 [Further reading](#) p. 21

model are no longer valid. In these cases it may be time to completely rebuild the model from scratch.

MLOps

Machine learning operations, or MLOps, is an emerging concept aimed at actively managing this lifecycle. Rather than an ad hoc approach to verifying and retraining when appropriate, [MLOps tools](#) put each model on a schedule for development, deployment, verification and retraining. It seeks to standardize these processes, a practice that's becoming more important as enterprises make machine learning a core component of their operations.


Future Trends

When we look to the future of machine learning, one overarching trend predominates. Enterprise adoption will continue to increase, bringing the technology from cutting edge to mainstream.

The trend is already well underway.

In this e-guide

- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21



A 2019 survey from analyst firm Gartner found that 37% of enterprises have adopted some form of artificial intelligence. That's up from 10% in 2015. At its current trajectory, machine learning is on a path to become a ubiquitous technology in the next few years. In its ranking of the top 10 data and analytics trends for 2020, the analyst firm named

"smarter, faster and more responsible AI" as the year's top trend. The report, noting the vital importance of machine learning and other AI techniques in providing insight into the global coronavirus pandemic, predicted that by 2024, 75% of organizations will have shifted from piloting to operationalizing AI. As a result of high rates of adoption of machine learning in the enterprise, the market for machine learning tools is growing rapidly. The analyst firm Research and Markets predicted that the machine learning market will grow to \$8.8 billion by 2022, from \$1.4 billion in 2017.

In this e-guide








- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21

The reasons for this are clear. Today's most successful companies, like Amazon, Google and Uber, put machine learning applications at the center of their business models. Rather than viewing machine learning as a nice-to-have technology, industry-leading enterprises are using machine learning and AI technologies as critical to maintaining their competitive edge, as technology writer George Lawton explored in "[Learn the business value of AI's various techniques.](#)"

[Advances in deep learning](#) -- a type of machine learning based on neural networks -- have played a huge role in bringing AI to the fore in the enterprise. Neural networks are relatively common in enterprise applications today. These advanced [deep learning techniques](#) enable models to do everything from recognize objects in images to create natural language text for product descriptions and other applications. Today, there are a number of different types of neural networks, which are designed to perform specific jobs. As technology writer David Petersson explained in "[CNNs vs. RNNs: How they differ and where they overlap](#)," understanding the uniqueness of different types of algorithms is key to getting the most out of them.

It is now viewed as inevitable that a large amount of knowledge work will be automated. Even some creative fields are being infiltrated by machine learning-driven AI applications. This is raising questions about the future of work. In a world where machines are able to manage customer relations, detect cancer in medical images, conduct legal reviews, drive shipping containers across the country and produce creative assets, what is the role of human workers? Proponents of AI say automation will free people up to pursue more creative activities by eliminating rote tasks. But

In this e-guide

 In-depth guide to machine learning in the enterprise	p. 2
 Enterprise use cases and benefits	p. 3
 Challenges	p. 7
 Implementation: 6 Steps	p. 10
 Management and maintenance of ML	p. 14
 Future Trends	p. 15
 Vendors and platforms	p. 18
 Further reading	p. 21

others worry that an incessant drive for automation will leave little room for human workers.

Vendors and platforms

Enterprises looking to deploy machine learning have no shortage of options. The machine learning space features strong competition between open source tools and software built and supported by traditional vendors. Regardless of whether an [enterprise chooses machine learning software](#) from a vendor or open source tool, it is common for applications to be hosted in the cloud computing environments and delivered as a service. There are more vendors and platforms than one article could name, but the following list gives a high-level overview of offerings from some of the bigger players in the field.

Vendor tools

- Amazon Sagemaker is a cloud-based tool that allows users to work at a range of levels of abstraction. Users can run pretrained algorithms for simple workloads or code their own for more expansive applications.
- Google Cloud is a collection of services that range from plug-and-play AI components to data science development tools.
- IBM Watson Machine Learning is delivered through the IBM cloud and allows data scientists to build, train and deploy machine learning applications.

In this e-guide

- [In-depth guide to machine learning in the enterprise](#) p. 2
- [Enterprise use cases and benefits](#) p. 3
- [Challenges](#) p. 7
- [Implementation: 6 Steps](#) p. 10
- [Management and maintenance of ML](#) p. 14
- [Future Trends](#) p. 15
- [Vendors and platforms](#) p. 18
- [Further reading](#) p. 21

- Microsoft Azure Machine Learning Studio is a graphical user interface tool that supports building and deploying machine learning models on the Microsoft cloud.
- SAS Enterprise Miner is a machine learning offering from a more traditional analytics company. It focuses on building enterprise machine learning applications and productionalizing them quickly.

Open source

- Caffe is a framework is specifically engineered to support the development of deep learning models -- in particular, neural networks.
- Scikit-learn is an open source library of Python code modules that allow users to do traditional machine learning workloads like regression analysis and clustering.
- TensorFlow is a machine learning platform built and open sourced by Google. It is commonly used for developing neural networks.
- Theano was originally released in 2007 and is one of the oldest and most trusted machine learning libraries. It is optimized to run jobs on GPUs, which can result in fast machine learning algorithm training.
- Torch is a machine learning library that is optimized to train algorithms on GPUs. It is built primarily to train deep learning neural networks.

A more exhaustive list of vendor offerings can be found in this expert [overview of machine learning platforms](#).

In general, most enterprise machine learning users consider open source tools to be more innovative and powerful. However, there is still a strong case for proprietary tools, as vendors offer training and support that is generally absent from open source

In this e-guide

- In-depth guide to machine learning in the enterprise p. 2
- Enterprise use cases and benefits p. 3
- Challenges p. 7
- Implementation: 6 Steps p. 10
- Management and maintenance of ML p. 14
- Future Trends p. 15
- Vendors and platforms p. 18
- Further reading p. 21

offerings. Many of today's vendor tools support use of open source libraries, allowing users to have the best of both worlds.

Further reading

In this e-guide

- 📖 In-depth guide to machine learning in the enterprise p. 2
- 📖 Enterprise use cases and benefits p. 3
- 📖 Challenges p. 7
- 📖 Implementation: 6 Steps p. 10
- 📖 Management and maintenance of ML p. 14
- 📖 Future Trends p. 15
- 📖 Vendors and platforms p. 18
- 📖 Further reading p. 21

📖 About SearchEnterpriseAI

Smart organizations are using artificial intelligence (AI) to introduce new business opportunities, identify new revenue streams, enhance the customer and employee experience, and transform processes.

Where are you with tapping into the power of AI? No matter the answer, you've come to the right place.

SearchEnterpriseAI.com is here to help ensure that your AI investments aren't perceived as a waste of money, but instead an important source of the business value that you crave..

For further reading, visit us at
<http://SearchEnterpriseAI.com/>

Images: Fotalia

©2020 TechTarget. No part of this publication may be transmitted or reproduced in any form or by any means without written permission from the publisher.