INSIDER PRO
FROM IDG

EXPERT INSIGHT INTO HOW TECHNOLOGY DRIVES BUSINESS

INSIDER EXCLUSIVE

INSIDE AI: ARTIFICIAL INTELLIGENCE in the ENTERPRISE

AI is having a profound effect on enterprises, and a variety of tools and techniques can help you **get started on your own implementation.**

POWERED BY CIO COMPUTERWORLD CSO InfoWorld NETWORKWORLD

Editor's NOTE

Artificial intelligence is fueling enormous change in business, allowing for more accurate predictive analytics than ever before, so-called autonomous data-center systems and new services in areas like finance, customer relations and healthcare.

AI can optimize business processes and IT operations, automating routine jobs and allowing staff to focus on high-value tasks.

AI, though, may be at once the most transformative and least understood areas of technology today. Adding to the confusion, AI comprises concepts and terms that overlap and create misunderstanding about the field

What is artificial intelligence?

AI is a blanket term for various algorithms, methodologies and technologies that, broadly speaking, allow software to accomplish tasks that historically were associated with human intelligence, such as responding to natural language queries, recognizing images and sounds, making decisions and predictions. Machine intelligence is a more recent synonym for AI.

Perhaps the biggest point of confusion surrounding AI involves machine learning, a term that has become well-known outside tech. Non-specialists often use it as a synonym for AI since it's responsible for some of the more spectacular recent advances in the field, such as self-driving cars. Machine learning is however a subset of AI that is itself an umbrella term, embracing methods and algorithms – such as neural networks -- that let computer systems better their performance as they ingest more data.

How you can implement AI in your own business

AI technology has been baked into many off-the-shelf packaged software products. It's used for predictive analytics in ERP, CRM and networking software from vendors like SAP, Microsoft, Salesforce and HPE. It's used for self-maintenance of Oracle's Autonomous Database. Help desk and customer service software from vendors like BMC and ServiceNow use AI to classify incidents and route tasks. An exhaustive list would fill a book-length manual.

Many enterprises, though, have rolled out their own AI-based applications. There are pre-built APIs you can incorporate into your own systems, cloud-based machine learning offerings for building data models and a wide range of programming tools you can use to build AI applications from the ground up.

To be successful with AI, you need to understand, among other things, the techniques that can be used, the tools and services that are available, the skills required and the relevant data you have. Above all you need a concrete plan for putting AI to work for specific business goals.

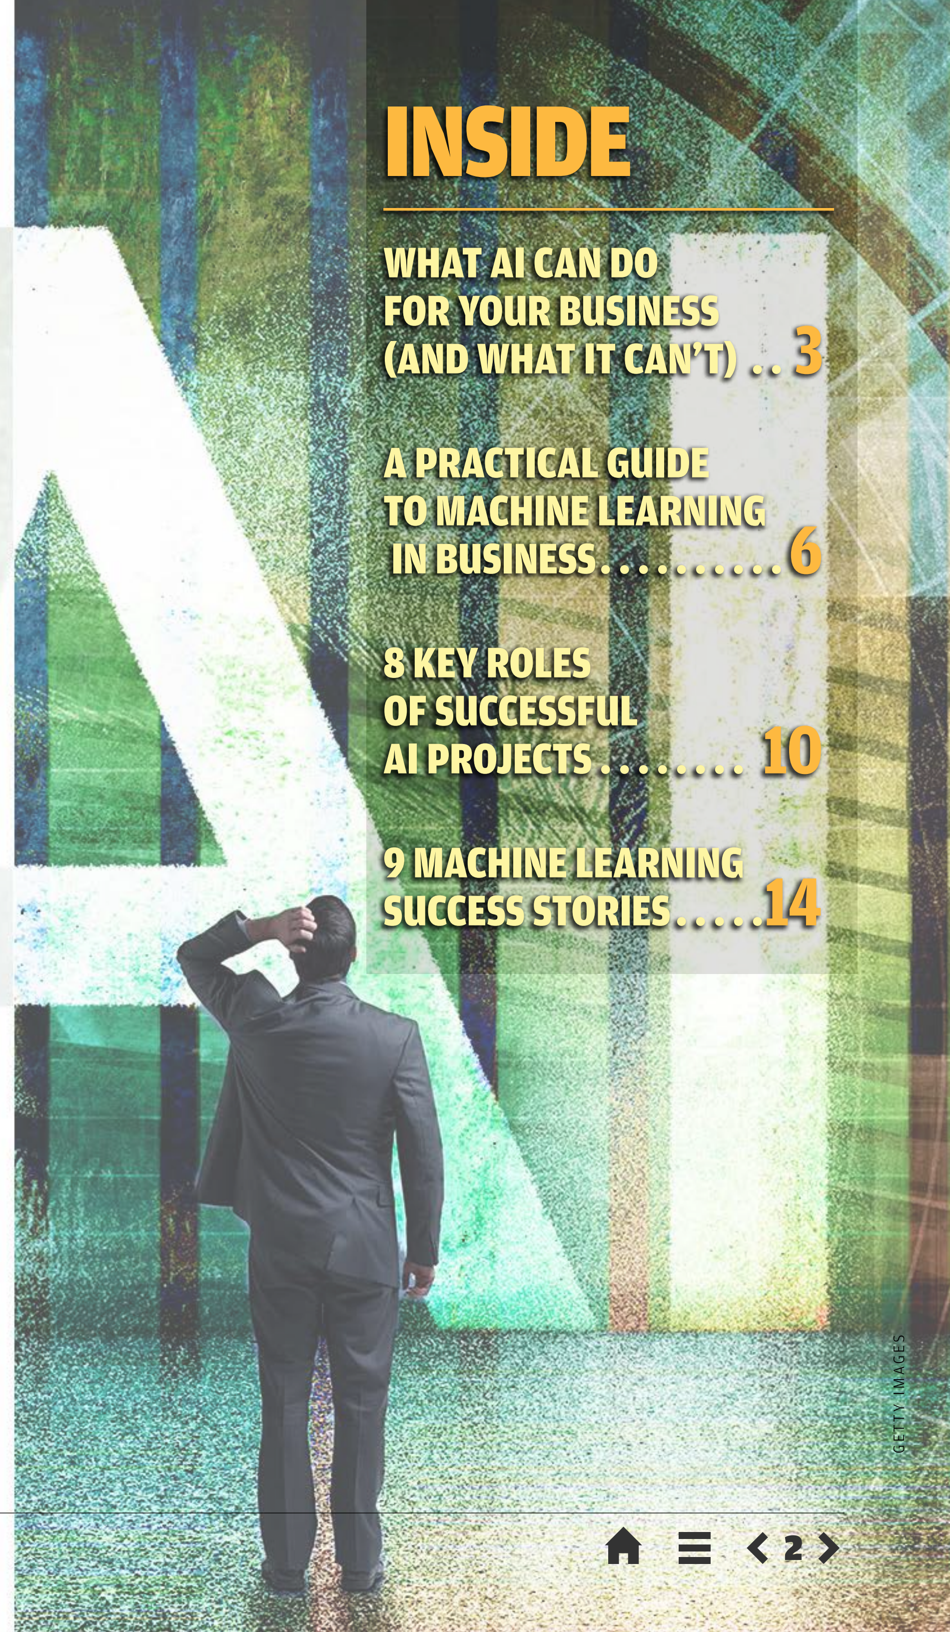

GETTY IMAGES

INSIDE

WHAT AI CAN DO FOR YOUR BUSINESS (AND WHAT IT CAN'T)

GETTY IMAGES

Artificial intelligence, machine learning, and deep learning are no silver bullets. Here's what every business should know before investing in AI.

BY ISAAC SACOLICK

How can you tell whether an emerging technology such as artificial intelligence is worth investing time in when there is so much hype being published daily? We're all enamored by some of the amazing results, such as AlphaGo beating the champion Go player, voice recognition being performed by Amazon's Alexa and Microsoft's Cortana, and the image recognition done by Google Photos, Amazon Rekognition, and other photo-sharing applications.

When big, technically strong companies like Google, Amazon, Microsoft, IBM, and Apple show success with a new technology and the media glorifies it, businesses often believe these technologies are available for their own use. But is it true? And if so, where is it true?

What artificial intelligence really is, and how it got there

AI technology has been around for some time. The modern era of AI got its big start in 1968-69 when the SHRDLU natural language processing (NLP) system came out, research papers on perceptron and backpropagation algorithms were published, and the world became aware of AI through HAL in "2001: A Space Odyssey." The next major breakthroughs can be pinned to the late 1980s with the use of back propagation in learning algorithms and then their application to problems like handwriting recognition. AI took on large scale challenges in the late 1990s with the the first chatbot (ALICE), and IBM's Deep Blue beating world chess champion Gary Kasparov.

Flash-forward to today and it's easy to see why there's an exponential leap in AI results over the last several years thanks to several advances.

First, there's cloud computing, which enables running large neural networks on a cluster of machines. Instead of looping through perceptron learning algorithms one at a time and working with only one or two network layers, computation is distributed across a large array of computing nodes. This is enabling deep learning algorithms, which are essentially neural networks with a large number of nodes and layers that enable processing of large-scale problems in reasonable amounts of time.

Second, there's the emergence of commercial and open source libraries and services like TensorFlow, Caffe, Apache MXNet, providing data scientists and software developers with the tools to apply machine learning and deep learning algorithms to their data sets without having to program the underlying mathematics or enable parallel computing. Future AI applications will be driven by AI on a chip or board, driven by the innovation and competition among Nvidia, Intel, AMD, and others.

Don't confuse AI hype with AI realities

Once you have a grasp of history and an understanding of the technology, it's often useful to review where an emerging technology is in its life cycle.

Gartner has machine learning and deep learning at the peak of their hype cycles and forecasts that "general AI" (AI applied to any intelligence problem) will emerge after 2020. McKinsey has reported more than 50 percent of AI investments are coming from tech giants and startups versus businesses that happen to use technology.

Those stats should give any CIO or business executive a pause before jumping into AI investments with both feet. Although AI is certainly demonstrating a lot of promise, the commercial application of these algorithms at scale is still relatively young among enterprises. The early winners are big tech companies and startups with the talent, funding, and patience to experiment with new technologies. Most enterprises and medium-size businesses simply don't have these luxuries and are just starting their AI journeys.

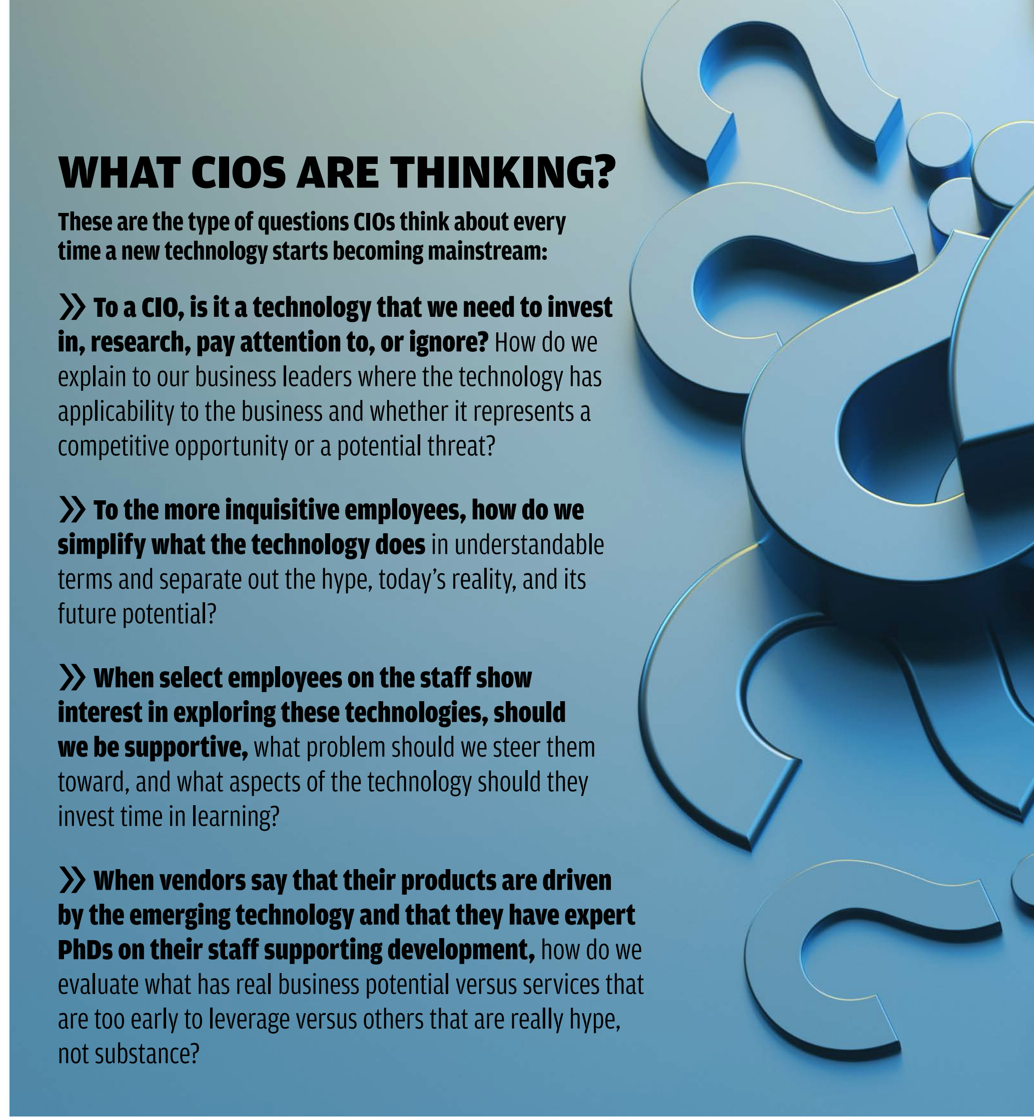

WHAT CIOS ARE THINKING?

These are the type of questions CIOs think about every time a new technology starts becoming mainstream:

» **To a CIO, is it a technology that we need to invest in, research, pay attention to, or ignore?** How do we explain to our business leaders where the technology has applicability to the business and whether it represents a competitive opportunity or a potential threat?

» **To the more inquisitive employees, how do we simplify what the technology does** in understandable terms and separate out the hype, today's reality, and its future potential?

» **When select employees on the staff show interest in exploring these technologies, should we be supportive,** what problem should we steer them toward, and what aspects of the technology should they invest time in learning?

» **When vendors say that their products are driven by the emerging technology and that they have expert PhDs on their staff supporting development,** how do we evaluate what has real business potential versus services that are too early to leverage versus others that are really hype, not substance?

GETTY IMAGES

AI *is* a highly disruptive technology, so you should not ignore it. Just proceed judiciously and avoid getting hypnotized by the AI hype.

For example, *when* voice becomes a better human machine interface than screens and keyboards for some applications, many businesses will have to upgrade their user experiences with the technology.

And *when* we begin to pull intelligence from spoken language, audio, and video as effect ively as we can with more structured data, using these capabilities will provide significant competitive advantages to a large array of businesses.

When is the operative word.

Most businesses, particularly small and medium-size enterprises, should aim to be fast followers, not early adopt-

EXAMPLES OF AI IN ACTION

As you learn about AI capabilities, look for tools and practical example to help evaluate applications of AI and their maturity.

» Forrester defines nine essential AI technologies and proposes a building-block model that starts with hypothesis and research and ends with three levels of pragmatic application.

» Workday published its maturity model that outlines four stages of AI evolution: automating, informing, discovering and transforming.

» There are many examples by industry including financial services, healthcare, casual dining, and manufacturing.

On the other hand, if the solution requires significant cognitive evaluation and thinking, you're heading into an immature AI space.

ers. That means paying attention and even experimenting with AI in these relatively early days, but waiting to rely on AI until the technology is sufficiently mature, proven, and able to deliver at scale for a given application.

Clean data needed

In any case, don't be fooled when a vendor says something like, "Just throw your data at our AI" and expect expert intelligence to be returned. It won't happen. A prerequisite for running AI successfully is that you need large amounts of relatively clean data to train AI solutions and evaluate outputs.

One reason autonomous vehicles are possible is the 4,000 GB data generated from one hour of driving, from lidar and other sensors found in these cars. That's a lot of data being used to make what are really just a handful of fundamental decisions on whether the car should turn, speed up, slow down, or fully stop.

Many successful AI solutions fall into this same category of churning large amounts of data into a finite number of decisions. In image recognition, for example, am I looking at a picture that contains you in it, or not? In collaborative filtering, is a newly published article more relevant to you based on your past reading experiences versus other reading options? When evaluating a transaction, does it have similar patterns to fraudulent transactions?

The AI "inside the box" is trying to approximate a curve to make these decisions. In deep learning, for example, the number of layers and neurons in the network can approximate highly complex curves to differentiate outcomes. To develop this network, you need a large, tagged data set so the network can be trained by comparing its computed results against your tagged result with the desired outcome. The errors are then used to tune the network using backpropagation or other learning algorithms, and the exercise is repeated multiple times across all tagged data until the network stabilizes to an optimized curve. These are supervised learning solutions, developed using a training set.

If the data isn't tagged, networks can use unsupervised learning approaches. For example, when Google's DeepMind was used to learn to play the Atari game Breakout, it used the score to evaluate outcomes.

Going beyond data sets, your organization needs a data integration and automation capability so you can move data into and out of any AI processing engine. If your organization is used to having people run scripts manually to push data around, I strongly suggest first investing in automation before diving into AI solutions.

Your options for experimenting with AI

Once you have business opportunities identified and large, cleansed data sets available, you are ready to consider an AI journey. Those first two steps are prerequisites for preparing your organization for artificial intelligence. The main next steps are to consider the type of AI solution and implementation. If you have the talent, you can experiment with TensorFlow or one of the other AI engines. If you don't have the expertise, think twice about trying to recruit for it; the tech giants are paying huge salaries for scarce AI talent, and so the costs are enormous to just get into the game.

A second option is to use vendors that have embedded AI in their solutions. One example is Salesforce Einstein, an AI platform that can perform forecasting and other functions on top of CRM data stored in Salesforce. Likewise, you can look at industry-specific solutions, such as Synechron's Neo for financial tech (fintech).

Once you settle on one or more approaches, it's important to set realistic expectations with stakeholders. Investing in AI requires a commitment to agile experimentation because you're likely to encounter many dead ends and experiments that require many runs before they are optimized. Set those expectations for budget, timeframe, and talent upfront. ◆

ISAAC SACOLICK *is the author of Driving Digital: The Leader's Guide to Business Transformation through Technology. Sacolick is a recognized top social CIO, a long-time blogger at Social, Agile and Transformation and CIO.com, and president of StarCIO.*

GETTY IMAGES

A PRACTICAL GUIDE TO MACHINE LEARNING IN BUSINESS

GETTY IMAGES

Machine learning is poised to have a profound impact on your business but the hype is sowing confusion. Here's a look at what machine learning is and how it can be used today.

BY MARY BRANSCOMBE

Machine learning is transforming business. But even as the technology advances, companies still struggle to take advantage of it, largely because they don't understand how to strategically implement machine learning in service of business goals. Hype hasn't helped, sowing confusion over what exactly machine learning is, how well it works and what it can do for your company. Here, we provide a clear-eyed look at what machine learning is and how it can be used today.

What is machine learning?

Machine learning is a subset of artificial intelligence that enables systems to learn and predict outcomes without explicit programming. It is often used interchangeably with the term AI because it is the AI technique that has made the greatest impact in the real world to date, and it's what you're most likely to use in your business. Chatbots, product recommendations, spam filters, self-driving cars and a huge range of other systems leverage machine learning, as do "intelligent agents" like Siri and Cortana.

Instead of writing algorithms and rules that make decisions directly, or trying to program a computer to "be intelligent" using sets of rules, exceptions and filters, machine learning teaches computer systems to make decisions by learning from large data sets. Rule-based systems quickly become fragile when they have to account for the complexity of the real world; machine learning can create models that represent and generalize patterns in the data you use to train it, and it can use those models to interpret and analyze new information.

Machine learning is suitable for **classification**, which includes the ability to recognize text and objects in images and video, as well as finding associations in data or segmenting data into clusters (e.g., finding groups of customers). Machine learning is also adept at **prediction**, such as calculating the likelihood of events or forecasting outcomes. Machine learning can also be used to generate missing data; for example, the latest version of CorelDRAW uses machine learning to interpolate the smooth stroke you're trying to draw from multiple rough strokes you make with the pen tool.

To comment on this story, visit Insider Pro's Twitter page.

At the heart of machine learning are algorithms. Some, such as regressions, k-means clustering and support vector machines, have been in use for decades. Support vector machines, for example, use mathematical methods for representing how a dividing line can be drawn between things that belong in separate categories. The key to effective use of machine learning is matching the right algorithm to your problem.

Neural networks

A neural network is a machine learning algorithm built on a network of interconnected nodes that work well for tasks like recognizing patterns.

Neural networks aren't a new algorithm, but the availability of large data sets and more powerful processing (especially GPUs, which can handle large streams of data in parallel) have only recently made them useful in practice. Despite the name, neural networks are based only loosely on biological neurons. Each node in a neural network has connections to other nodes that are triggered by inputs. When triggered, each node adds a weight to its input to mark the probability that it does or doesn't match that node's function. The nodes are organized in fixed layers that the data flows through, unlike the brain, which creates, removes and reorganizes synapse connections regularly.

Deep learning

Deep learning is a subset of machine learning based on deep neural networks. Deep neural networks are neural network that have many layers for performing learning in multiple steps. **Convolutional deep neural networks** often perform image recognition by processing a hierarchy of features where each layer looks for more complicated objects. For example, the first layer of a deep network that recognizes dog breeds might be trained to find the shape of the dog in an image, the second layer might look at textures like fur and teeth, with other layers recognizing ears, eyes, tails and other characteristics, and the final level distinguishing different breeds. **Recursive deep neural networks** are used for speech recognition and natural language processing, where the sequence and context are important.

There are many open source deep learning toolkits available that you can use to build your own systems. Theano, Torch and Caffe are popular choices, and Google's TensorFlow and Microsoft Cognitive Toolkit let you use multiple servers to build more powerful systems with more layers in your network.

Microsoft's Distributed Machine Learning Toolkit packages up several of these deep learning toolkits with other machine learning libraries, and both AWS and Azure offer VMs with deep learning toolkits pre-installed.

Machine learning in practice

Machine learning results are a percentage certainty that the data you're looking at matches what your machine learning model is trained to find. So, a deep network trained to identify emotions from photographs and videos of people's faces might score an image as "97.6% happiness 0.1% sadness 5.2% surprise 0.5% neutral 0.2% anger 0.3% contempt 0.01% disgust 12% fear." Using that information means working with probabilities and uncertainty, not exact results.

Probabilistic machine learning uses the concept of probability to enable you to perform machine learning without writing algorithms at all. Instead of the set values of variables in standard programming, some variables in probabilistic programming have values that fall in a known range and others have unknown values. Treat the data you want to understand as if it was the output of this code and you can work backwards to fill in what those unknown values would have to be to produce that result. With less coding, you can do more prototyping and experimenting; probabilistic machine learning is also easier to debug.

This is the technique the Clutter feature in Outlook uses to filter messages that are less likely to be interesting to you based on what messages you've read, replied to and deleted in the past. It was built with Infer.NET, a .NET framework you can use to build your own probabilistic systems.

Cognitive computing is the term IBM uses for its Watson offerings, because back in 2011 when an earlier version won Jeopardy, the term AI wasn't fashionable; over the decades it's been worked on, AI has gone through alternating periods of hype and dismissal.

Watson isn't a single tool. It's a mix of models and APIs that you can also get from other vendors such as Salesforce, Twilio, Google and Microsoft. These give you so-called "cognitive" services, such as image recognition, including facial recognition, speech (and speaker) recognition, natural language understanding, sentiment analysis and other recognition APIs that look like human cognitive abilities. Whether it's Watson or Microsoft's Cognitive Services, the cognitive term is really just a marketing brand wrapped around a collection of (very useful) technologies. You could use these APIs to create a chatbot from an existing FAQ page that can answer text queries and also recognise photos of products to give the right support information, or use photos of shelf labels to check stock levels.

Many "cognitive" APIs use deep learning, but you don't need to know how they're built because many work as REST APIs that you call from your own app. Some let you create custom models from your own data. Salesforce Einstein has a custom image recognition service and Microsoft's Cognitive APIs let you create custom models for text, speech, images and video.

That's made easier by **transfer learning**, which is less a technique and more a useful side effect of deep networks. A deep neural network that has been trained to do one thing, like translating between English and Mandarin, turns out to learn a second task, like translating between English and French, more efficiently. That may be because the very long numbers that represent, say, the mathematical relationships between words like *big* and *large* are to some degree common between languages, but we don't really know.

Transfer learning isn't well understood but it may enable you to get good results from a smaller training set. The Microsoft Custom Vision Service uses transfer learning to train an image recognizer in just a few minutes using 30 to 50 images per category, rather than the thousands usually needed for accurate results.

Build your own machine learning system

If you don't want pre-built APIs, and you have the data to work with, there's an enormous range of tools for building machine learning systems, from R and Python scripts, to predictive analytics using Spark and Hadoop, to specific AI tools and frameworks.

Supervised learning

Many machine learning techniques use **supervised learning**, in which a function is derived from labelled training data. Developers choose and label a set of training data, set aside a proportion of that data for testing, and score the results from the machine learning system to help it improve. The training process can be complex, and results are often probabilities, with a system being, for example, 30 percent confident that it has recognized a dog in an image, 80 percent confident it's found a cat, and maybe even 2 percent certain it's found a bicycle. The feedback developers give the system is likely a score between one and zero indicating how close the answer is to correct.

It's important not to train the system too precisely to the training data; that's called overfitting and it means the system won't be able to generalize to cope with new inputs. If the data changes significantly over time, developers will need to retrain the system due to what some researchers refer to as "ML rot."

Machine learning algorithms – and when to use them

If you already know what the labels for all the items in your data set are, assigning labels to new examples is a **classification** problem. If you're trying to predict a result like the selling price of a house based on its size, that's a **regression** problem because house price is a continuous rather than discrete category. (Predicting whether a house will sell for more or less than the asking price is a classification problem because those are two distinct categories.)

If you don't know all the labels, you can't use them for training; instead, you score the results and leave your system to devise rules that make sense of the answers it gets right or wrong, in what's known as **unsupervised learning**. The

AI in the clouds

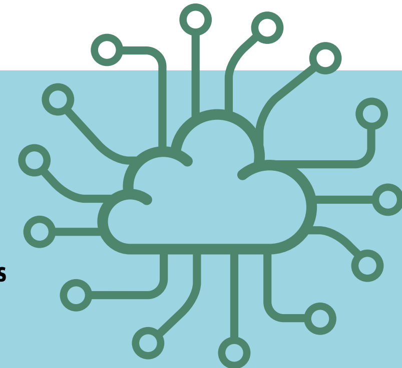

GETTY IMAGES

Rather than set up your own infrastructure, you can use machine learning services in the cloud to build data models. With cloud services you do not need to install a range of tools. Moreover, these services build in more of the expertise needed to get successful results.

- **Amazon Machine Learning** offers several machine learning models you can use with data stored in S3, Redshift or R3, but you can't export the models, and the training set size is rather limited.
- **Microsoft's Azure ML Studio** has a wider range of algorithms, including deep learning, plus R and Python packages, and a graphical user interface for working with them. It also offers the option to use Azure Batch to periodically load extremely large training sets, and you can use your trained models as APIs to call from your own programs and services. There are also machine learning features such as image recognition built into cloud databases like SQL Azure Data Lake, so that you can do your machine learning where your data is.

GETTY IMAGES

As machine learning becomes more widely used, you'll need to explain why your machine learning-powered systems do what they do.

most common unsupervised learning algorithm is **clustering**, which derives the structure of your data by looking at relationships between variables in the data. Amazon's product recommendation system that tells you what people who bought an item also bought uses unsupervised learning.

With **reinforcement learning**, the system learns as it goes by seeing what happens. You set up a clear set of rewards so the system can judge how successful its actions are. Reinforcement learning is well suited to game play because there are obvious rewards. Google's DeepMind AlphaGo used reinforcement learning to learn Go, Microsoft' Project Malmo system allows researchers to use Minecraft as a reinforcement learning environment, and a bot built with OpenAI's reinforcement learning algorithm recently beat several top-ranked players at Valve's Dota 2 game.

The complexity of creating accurate, useful rewards has limited the use of reinforcement learning, but Microsoft has been using a specific form of reinforcement learning called **contextual bandits** (based on the concept of a multi-armed slot machine) to significantly improve click-through rates on MSN. That system is now available as the Microsoft Custom Decision Service API. Microsoft is also using a reinforcement learning system in a pilot where customer service chatbots monitor how useful their automated responses are and offer to hand you off to a real person if the information isn't what you need; the human agent also scores the bot to help it improve.

Combining machine learning algorithms for best results

Often, it takes more than one machine learning method to get the best result; **ensemble learning** systems use multiple machine learning techniques in combination. For example, the DeepMind system that beat expert human players at Go uses not only reinforcement learning but also supervised deep learning to learn from thousands of recorded Go matches between human players. That combination is sometimes known as **semi-supervised learning**.

Similarly, the machine learning system that Microsoft Kinect uses to recognize human movements was built with a combination of a **discriminative** system — to build that Microsoft rented a Hollywood motion-capture suite, extracted the position of the skeleton and labelled the individual body parts to classify which of the various known postures it was in — and a **generative** system, which used a model of the characteristics of each posture to synthesize thousands more images to give the system a large enough data set to learn from.

Predictive analytics often combines different machine learning and statistical techniques; one model might score how likely a group of customers is to churn, with another model predicting which channel you should use to contact each person with an offer that might keep them as a customer.

Navigating the downsides of machine learning

Because machine learning systems aren't explicitly programmed to solve problems, it's difficult to know how a system arrived at its results. This is known as a "black box" problem, and it can have consequences, especially in regulated industries.

As machine learning becomes more widely used, you'll need to explain why your machine learning-powered systems do what they do. Some markets — housing, financial decisions and healthcare — already have regulations requiring you to give explanations for decisions. You may also want **algorithmic transparency** so that you can audit machine learning performance. Details of the training data and the algorithms in use isn't enough. There are many layers of non-linear processing going on inside a deep network, making it very difficult to understand why a deep network is making a particular decision. A common technique is to use another machine learning system to describe the behavior of the first.

You also need to be aware of the dangers of **algorithmic bias,** such as when a machine learning system reinforces the bias in a data set that associates men with sports and women with domestic tasks because all its examples of sporting activities have pictures of men and all the people pictured in kitchens are women. Or when a system that correlates non-medical information makes decisions that disadvantage people with a particular medical condition.

Machine learning can only be as good as the data it trains on to build its model and the data it processes, so it's important to scrutinize the data you're using. Machine learning also doesn't understand the data or the concepts behind it the way a person might. For example, researchers can create pictures that look like random static but get recognized as specific objects.

There are plenty of recognition and classification problems that machine learning can solve more quickly and efficiently than humans, but for the foreseeable future machine learning is best thought of as a set of tools to support people at work rather than replace them. ◆

MARY BRANSCOMBE *is a freelance journalist who has been covering technology for over two decades.*

GETTY IMAGES

8 KEY ROLES OF SUCCESSFUL AI PROJECTS

As enterprises further develop artificial intelligence projects, they are finding that some roles are essential to success. But the right talent can be hard to find.

BY MARY BRANSCOMBE

Artificial intelligence offers ample opportunities to reap business value. When done right, AI can help improve sales, optimize operations, and free up staff for higher-value work. It can help reduce costs and empower organizations to create new products and pursue new markets.

And enterprises are diving in. According to a recent Deloitte survey, 55 percent of IT executives say their companies launched six or more AI-related pilot projects in 2018, up from 35 percent in 2017. More than a third have invested over $5 million in cognitive technologies, and 56 percent expect AI to transform their companies within the next three years.

But getting there isn't easy, and certain key skills are required — but hard to find. Here we take a look at eight key roles for AI success, according to those making early forays into artificial intelligence for business.

1 AI researchers

It might seem counterproductive for the average enterprise to get involved in research. After all, AI researchers are often the PhDs who perform fundamental research that could, someday, lead to a breakthrough in machines' abilities to think. Plus, going after AI researchers means competing against universities and tech giants like Google and Microsoft for near unicorns who might not immediately conjure business benefits.

But there's always hope that a breakthrough will catapult them into the lead. This promise alone may speak to the

Predictive analytics **combines different ML and statistical techniques;** one model might score how likely a group of customers is to churn, with another predicting which channel you should use to contact a customer.

high demand for AI researchers. According to Deloitte's survey, 30 percent of IT execs view finding AI researchers among their top priorities, more than any other role.

"People want that shiny object," says Vivek Katyal, global leader for analytics and data risk at Deloitte Risk and Financial Advisory. "But will that shiny object make a difference to what they're truly after?" Unless a company wants to be the next Facebook, he says, maybe not.

But many business executives making funding decisions don't understand the difference between AI research and AI applications, he says. "It's not the data scientists funding these projects."

For those companies where AI is critical to their core business, however, research is not a luxury, but a necessity. AppTek, for example, was founded around 30 years ago as a speech recognition company. The entire field of speech recognition has been transformed by AI, and AppTek has had to invest in research to keep up. For example, its latest published research focuses on identifying different speakers in a conversation.

"That's a real commercial need," says Mike Veronis, the company's chief revenue officer. "We did that to solve the problem, and to push capabilities."

2 AI software developers

AI software developers take fundamental research, such as the latest developments in deep learning or generative adversarial networks, and turn them into usable products. Some companies leave this work up to the big vendors, relying on commercial platforms rather than developing their own approaches to AI. But even if companies are using known AI techniques, they might still want to build their own platforms, says Deloitte's Katyal. This in part may explain the high demand for AI software developers, a top priority for 28 percent of respondents to Deloitte's survey.

One reason to build your own is the "black box" problem of current AI frameworks. Without the ability to see the source code of off-the-shelf products, some companies, especially in regulated areas like finance or healthcare, might rather pursue their own course.

"Maybe I should develop something on my own, where I know what I've built, I own the code, I control everything about it," Katyal says. "That discussion is very prevalent." When they build their own AI software, they can also get a better understanding of the built-in biases of the tools, he adds.

That is the case for AppTek as well. Instead of having a black box commercial system that can't easily be tuned, it gets a product that can be customized as needed, in addition to having unique features based on the company's own research. "We can adapt and train and continuously improve the speech recognition engine," says AppTek's Veronis.

3 Data scientists

When companies think about overcoming AI challenges, they typically think of creating new AI algorithms, Katyal says. But they would probably get more value from improving their data. "That is the usual barrier to functional AI," he says.

This makes data scientists the most important AI role of all, according to Katyal. Sought by 24 percent of respondents, data scientists prep a company's data for use in AI systems. They also identify the data a company needs to meet its goals — data that is either generated internally or gathered from third parties. Data scientists can also spot when data is missing, know when there isn't enough data of a particular type, and recognize when a data set is biased or out of date.

They are also the ones who identify the right algorithms to use on their data sets, train and tune those algorithms, and work with subject matter experts to validate the results.

"In the old days, they would have been advanced statisticians," says Katyal. "They are the users of the AI research and the AI software."

Data scientists are at the heart of Sumitomo Mitsui Banking Corp.'s recent AI projects. SMBC, a global financial company and the second largest bank in Japan by assets, is using AI to improve customer service in its data centers, to make it easier for employees to find information, and to better identify potential corporate customers.

The bank already had a data management department and data scientists on staff, says Akinobu Funayama, the bank's executive director. At first, the data scientists would manually set up use cases, identify the data points most relevant to those use cases, and create the algorithms to analyze the data. For example, when scoring potential new customers for profitability, data scientists would look at thousands of factors to see if any turn out useful.

The entire process would take two to three months per use case, translating into 10 to 15 use cases per year. Using technology from dotData to help identify data points most useful for creating new algorithms, SMBC has reduced the time it takes to create a new model to just a few hours. This has increased the number of use cases the bank can tackle to about a 100 per year, enabling it to apply AI to more areas in the bank, including finance, treasury, and compliance.

"We are working on improving the performance of the whole group," says Funayama.

The data scientists are still critical to the process, he says, but instead of doing repetitive feature engineering work, they are now tackling a much wider array of business use cases of AI technology.

4 User experience designers

As AI is incorporated into more products and services, user experience design is becoming increasingly important. Instead of opening menus or clicking buttons, people now expect to be able to ask plain-English questions, or have applications deduce what they need from context.

"We've always thought of user experience as being web-driven or mobile-driven," says Brandon Ebken, CTO at Insight, Tempe, Ariz.-based technology consulting firm. "In the AI world, we're interfacing with chatbots or Siri or Cortana, with voice." It's created a whole new type of user experience design, he says, and is a critical piece when creating new AI-powered tools.

"The connection between AI-powered things and human experience is evolving," Deloitte's Katyal agrees. "I think that's the next revolution, one that we're already starting to see."

As new tools are created, people have to be able to use them, and that can require new kinds of interfaces, as well as accompanying changes in the way that an application or business processes are structured.

To find people with these skills, companies should look for experts in customer service, he says.

5 Change management experts

Change management is the single most overlooked aspect of AI deployments, says Deloitte's Katyal. And it's not just enterprise employees that benefit from change management, but also users and customers. "It's the hardest thing," Katyal adds. "This is an area most ignored and undervalued in the enterprise."

Still, change management experts remain in high demand, a top needed skill for 22 percent of respondents to Deloitte's survey. AI projects can have a large impact on knowledge workers, who may refuse to accept AI recommendations if they have not been involved in the development of the solution, according to Deloitte.

"The fundamentals of fostering organizational change can get lost amid excitement around pilots, grassroots experiments, and vendor-driven hype," says the Deloitte report.

Moreover, 63 percent of IT managers surveyed said that, to cut costs, their company wants to use AI to automate as many jobs as possible — further underscoring the need for change management expertise.

GETTY IMAGES

Even when a company uses outside vendors for much of their AI functionality, **having in-house business expertise is critical.**

6 Project managers

Many AI projects are plagued with issues because they are often not managed with the same rigor that companies use with more mature technologies. Project managers capable of leading AI implementations can help integrate AI into a company's roles and processes and help measure and prove business value, a top-three challenge for 39 percent and 30 percent of Deloitte survey respondents, respectively. They can also deal with skill shortages in other areas related to AI.

It's difficult enough to find data scientists, much less data scientists who are also software engineers, user interface designers, security professionals, and subject matter experts. Because of this, AI projects include complex teams of people, says Marty Young, managing director at Slalom Build, the technology consulting division of Seattle-based Slalom.

Project managers are needed to wrangle all these roles. Moreover,

GETTY IMAGES

project managers will help multidisciplinary teams move AI from experimental pilot projects to becoming just another aspect of software engineering and the software lifecycle, says Steve Herrod, managing director at General Catalyst Partners, a venture capital firm focusing on high-tech startups. Herrod was previously CTO at VMware.

"We shouldn't lose sight of the project and program managers that need to understand the unique aspects of the models and fit them into the broader software releases that they must be part of," he adds.

As the field advances, there will be an even broader range of roles that will be relevant, such as people to handle audit and certification-related questions, Herrod says.

That's going to create more work, and more need, for project managers.

7 Business leaders to interpret AI results

Even when a company uses outside vendors for much of their AI functionality, having in-house business expertise is critical.

That was the case for Spoton Logistics, an India-based shipping company looking to use AI to help with customer service, sentiment analysis, and automation in the finance department. For example, one specific use case is to solve the company's "first mile" and "last mile" address problem.

"India addresses are not standard," says Satya Pal, the company's head of business engineering. It only gets worse when the company is working with addresses that have not been fully filled in. "This takes away the possibility of central planning and vehicle utilization."

The company decided to use outside vendors for much of the work, instead of building the technology in-house. However, the business leaders who were needed to interpret the AI results were on the company's internal team. They had business knowledge of the specific problem the company was trying to solve, and an understanding of various AI models and frameworks, he says. For example, they were able to understand the applications of classification models versus reinforcement learning, and supervised versus unsupervised learning.

"Generally, they were from a computer science background with Python knowledge," he says. Some additional training was required, but this was usually independent research as well as AI-related online courses.

This allowed them to determine which AI approach was best suited to solve particular products and validate progress.

8 Subject matter experts

Because off-the-shelf AI tools don't always work for all use cases, subject matter experts are key. Take, for example, product recommendation engines, which are typically designed around the needs of online retailers, says Michael Rigney, SVP of client solutions at EnergySavvy, a software company focusing on the utility industry.

Online retailers collect data about the shopping habits of their customers and can compare that to the shopping habits of other customers. But past purchases aren't useful metrics for those who get electricity from the local utility company. Here, expertise from companies such as EnergySavvy can help.

"We know how to identify which customers are benefitting from energy-efficiency projects, how much they're benefitting, and who else would be similar to those customers and would benefit as well," says Rigney. That has helped EnergySavvy to serve clients such as NationalGrid in Massachusetts, he says.

The new AI capabilities are responsible for the vast majority of the company's recent revenue growth, says EnergySavvy's VP of marketing Ryan Warren. "Most of our new customer growth, the lines of business that are the future of our company, are fundamentally tied to technologies that are underpinned by AI." ◆

MARY BRANSCOMBE *is a freelance journalist who has been covering technology for over two decades.*

9 Machine learning success stories

IT leaders share how they are using artificial intelligence and machine learning to generate business insights and new services.

BY CLINT BOULTON

GETTY IMAGES

Artificial intelligence (AI) and machine learning, once considered pie-in-the-sky projects for companies, have become mainstream.

More enterprises are harnessing the technology, which mimics the behavior of the human mind, to woo customers and bolster business operations.

Fifty-nine percent of 106 IT and business professionals told Gartner they had deployed AI to date and had, on average, four AI/ML projects in place. Respondents also said they expect to add six more projects in the next 12 months, and another 15 within the next three years. By 2022, those organizations expect to have an average of 35 AI or ML projects in place, says Gartner analyst Whit Andrews, who adds that the volume of planned projects "surprised" him.

Improved customer experience and task automation through virtual assistants for customer service and internal decision-making are among the most popular projects, Andrews says.

But projects alone do not make for a broader AI strategy, as a larger survey of 2,473 organizations conducted by IDC found that only 25 percent have developed an enterprise-wide AI strategy.

IDC found that as many as 25 percent of those it surveyed copped to a 50 percent failure rate, owing to lack of skilled staff and unrealistic expectations. But that isn't stopping companies from trying. Here CIOs who are experimenting with, building and implementing AI and ML technologies shared their ML use cases and some practical advice.

1

Insurer leverages chatbots for CX

At mutual insurance giant Guardian, CIO Dean Del Vecchio is testing AI and ML to improve both customer experience (CX) and improve employee productivity.

One CX pilot project includes using Amelia, a virtual assistant made by IPsoft, to help automate the onboarding process and to answer benefits questions, freeing up CX workers to focus on more complex cases. Educated with business logic from Guardian, Amelia answers such questions as, "What does a critical illness mean?" and other inquiries that Guardian's millions of customers may have.

Internally, Guardian is testing how IBM Watson can help the company better understand how customers are interacting with the business. As part of the pilot, IBM's cognitive computing software reads and prioritizes the millions of emails Guardian gets each day, a scan that happens in minutes, freeing up staff from hours spent manually reading and prioritizing the messages.

KEY ADVICE: ML projects should be considered within the greater context of the business. At Guardian, AI/ML-enabled process automation comprises just a fraction of a broader digital transformation Del Vecchio is leading to streamline operations while improving CX. This includes a move to Amazon Web Services, adoption of agile development, and the construction of an API marketplace.

To comment on this story, visit Insider Pro's Twitter page.

2

This blouse was recommended by AI

You could argue that Zulily wouldn't have a viable business without ML, which the retailer of women's apparel uses to personalize merchandize offers to its customers.

The software relies on dozens of signals — such as purchase history, time spent browsing a selection, taps and swipes on its mobile app, and social media behavior — to determine whether to send shoppers an offer via push notification or email, says Zulily CIO Luke Friang. "As we get to know you, we become better at listening to you."

While Zulily builds custom algorithms, it relies on a host of open source technologies, such as Hadoop, TensorFlow, and H2O, to round out its ML stacks.

KEY ADVICE: AI projects benefit from sound cloud architectures. Zulily's personalization efforts, for example, rely on a critical handshake between Amazon Web Services and Google Cloud Platform, in which GCP serves up recommendations that, when selected, whisk a consumer to Zulily's ecommerce platform, powered by AWS. Such a connection underscores how both GCP and AWS are fueling competition that will benefit CIOs in the long run, Friang says.

3

Meet JiLL: AI assistant for the corporate office

You wouldn't think a commercial real estate provider would have much use for AI on the face of it. But Jones Lang LaSalle (JLL) in June teamed up with Google on JiLL, a voice assistant that allows office employees to set up meetings, find colleagues, look up train schedules or fill service requests via voice or text.

For example, JiLL can handle such requests as, "Hey JiLL, book a weekly meeting with my team," or "Hey JiLL, find me an open desk on the 3rd floor this afternoon," according to JLL Chief Digital Product Officer Vinay Goel.

Goel says in a statement that JiLL takes into account JLL's datasets about buildings, user interactions and transactions with physical spaces, which are processed within GCP, as well as Kubernetes containers. "Over time, we expect JiLL to become an essential platform for hundreds of skills that help employees improve their daily productivity," Goel says.

KEY ADVICE: For organizations seeking to make a shift to services, virtual assistants can be a worthwhile investment. JLL's JiLL is part of a strategy to leverage technology to provide value-added services, ideally to attract more corporate customers. JLL plans to add additional skills and open up the platform to third-party capabilities, part of a marketplace strategy designed to boost adoption. More broadly, the initiative suggests virtual assistants, for years the toys of general-purpose consumption, have veered into vertical markets.

GETTY IMAGES

4

ML to marry dynamic pricing with convenience

Convenience store chain Wawa plans to use ML to dynamically change prices based on competitive factors. This could be a big boost in helping Wawa personalize offers for loyalty program customers, says CIO John Collier.

"We want to embrace data and algorithms so we're not setting a price but setting the rules," Collier adds. The trick lies in balancing improvements to the user experience with the cost to provide that experience, he adds.

KEY ADVICE: Collier says that he would prefer to have dynamic pricing in place today, but there's one problem: The emerg-

ing technologies won't interface well with his legacy systems. Because of this, modernization is essential to any AI strategy, Collier says, adding that Wawa is overhauling its legacy systems as part of a digital transformation. "We're investing heavily in our data strategy."

5

Credit reporting firm builds ML analytics engine

At credit card reporting giant Experian, a digital transformation paved the way for a new strategic product that leverages ML capabilities at its core: Ascend Analytics On Demand, a self-service analytics platform that enables companies to build predictive models to determine critical factors such as whether any of 220 million consumers qualify for the amount of credit they are asking for.

Customers can run sophisticated analysis against that data in a matter of minutes versus what currently takes several weeks, says Alex Lintner, Experian's president of consumer information services. Ideally, the tool will enable consumers to receive qualifications for credit in their moment of need.

Ascend comes as AI technologies are becoming pervasive in almost every new software product and service, according to Gartner.

"Customers want the ability to see enormous sets of information in real-time," says Experian Global CIO Barry Libenson, who oversaw the construction of the platform, which uses Hadoop and other analytics tools. "Gone are the days when we could prescribe things. They want them in real-time, when they want, in the way they want."

KEY ADVICE: You can't build new analytics platforms on legacy software and expect them to perform well. To support Ascend, Experian took a hybrid cloud approach and invested in open source tools, including containers, API engines and microservices. Experian also standardized the way it builds and consumes software, with applications and code that can be reused by its employees and customers around the globe.

6

Credit card company fights fraud with ML

Like credit monitoring companies, credit card companies are forever fighting fraudsters.

But at a time when many experts decry digital as the bane of our online privacy and security, ML and AI tools can make services more secure than plastic credit cards, says Ed McLaughlin, Mastercard's president of technology and operations.

Mastercard uses multiple layers of ML and AI to weed out consumers with malicious intent. At the heart of its safeguards lies an in-memory database that has saved Mastercard from an estimated $1 billion in fraud losses since 2016, McLaughlin says. The software uses more than 200 attributes to anticipate and head off fraud.

That core system, combined with tokenization, biometrics, deep learning and other newfangled approaches, has helped Mastercard maintain its reputation for facilitating billions of dollars' worth of secure transactions.

KEY ADVICE: Humans are the weakest link when it comes to cybersecurity. "What's most important is to take the human out of the loop" as much as possible, McLaughlin says, adding that ML, AI, and natural language processing software are all critical components in Mastercard's toolkit.

GETTY IMAGES

At a time when many experts decry digital as the bane of our online privacy and security, **ML and AI tools can make services more secure than plastic credit cards.**

-ED MCLAUGHLIN, PRESIDENT OF TECHNOLOGY AND OPERATIONS, MASTERCARD

7

Racing company taps ML analytics for car insights

Mercedes-AMG Petronas Motorsport is using ML capabilities to help visualize race car performance.

The company collects multiple channels of data on its Formula One race cars, sometimes as many as 10,000 data points per second, to make critical decisions, says Matt Harris, the company's head of IT.

Mercedes-AMG Petronas uses Tibco software to visualize the impact variables such as weather, tire temperature and fuel quantity have on its cars. The software also enables engineers to analyze details such as the performance of, and wear on, a car's gears. Drivers typically change gears 100 times per lap, and every time the driver makes a gear change, Tibco collects around 1,000 data points.

"When you visualize that data, you can actually make the gear box last longer, or more importantly, make harsher gear changes," Harris says. "You can then find that if you put the gear box into one particular mode, it's roughly 50 milliseconds faster per lap. Cars can be separated by thousandths of a second in qualifying, so 50 milliseconds matters."

To comment on this story, visit Insider Pro's Twitter page.

Harris says that Mercedes-AMG Petronas is building ML algorithms to help "do things that humans can't, or that is a prohibitively expensive way to do a piece of work." Harris believes these capabilities will eventually become a key enabler of the team's competitive advantage.

KEY ADVICE: Why build something that isn't your core competency? Before landing on Tibco, Mercedes-AMG Petronas used homegrown visualization software that proved too inefficient to maintain over time. By leaning on Tibco, Mercedes-AMG Petronas can focus on its strength: building high-performing cars. "The big thing is allowing people to be creative and think about solving problems, rather than coding software to visualize the problem," Harris says.

8

ML that predicts when employees will leave

Like most automotive repair companies, Caliber Collision has long had a turnover issue, as mechanics, painters and customer support staff tend to come and go at a rapid clip, sometimes as much as 40 percent a year across its 600-plus locations.

Part of the problem, Caliber found, is that its shops sometimes didn't have enough cars for staff to fix, resulting in inconsistent pay. This got CIO Ashley Denison wondering: What if Caliber could predict when an employee might leave and stage an intervention?

Caliber began working with tech consultant Sparkhound, which created software that pulls employee data from Caliber's Workday human resources software and mashes it up with Microsoft PowerBI to create a custom-built regression model that anticipates whether an employee might mull leaving. Then Caliber reaches out to staff, via digital surveys or personal contact, to check in.

For example, if an employee's pay dips over a couple of weeks, regional managers at Caliber can ensure they get more cars to work on. Conversely, if an employee appears to be overloaded with work, the company can reallocate some jobs to his or her coworkers. The result? Caliber is saving as much as $1 million a year reducing turnover.

KEY ADVICE: Heading off problems to save money is a practical way to use ML algorithms at a time when they tend to be overhyped and oversold. "It makes it a whole lot easier to keep employees once they come in the door," Denison says.

> * **Why build something that isn't your core competency?** Before landing on Tibco, Mercedes-AMG Petronas used homegrown visualization software that proved too inefficient to maintain.

9

AI as product and business enabler

Adobe Systems is harnessing ML to analyze help-desk tickets for system-failure trends, and then to proactively fix issues before they result in significant downtime, CIO Cynthia Stoddard told CIO.com at the 2018 MIT CIO Sloan Symposium.

If the system sees events that suggest an outage could occur, the system can be proactive to eliminate or mitigate those events before they trigger failures.

Called HAAS, for "healing-as-a-service," the tool catches and remediates anything from failed integrations with Adobe's ERP to faulty data feeds intended for the company's various analytics systems. Stoddard says HAAS has reduced fix times from 30 minutes performed manually by humans to 1 minute. She estimates it has saved Adobe 330 hours of remediation in the past several months. Using reports detailing the issues, Adobe engineers are then able to create permanent fixes.

"If you know you have to fix something and you know how to fix it, then you can automate it," Stoddard said. "It's been a tremendous benefit." The work builds on the ML-based diagnostic testing framework Stoddard's team created in 2017.

KEY ADVICE: Using ML to identify patterns is the key to creating self-healing capabilities. "If you know how you fixed it you can put a self-healing component in there and take the human element out of the equation," Stoddard says.

CLINT BOULTON *is a senior writer for CIO.com, covering IT leadership, the CIO role, and digital transformation.*