


BUSINESS GUIDE

The CFO's Guide to AI and Machine Learning

These advanced technologies have possibilities for your business. Here's what you need to know.





Grab a seat and enjoy.
Read Time: 11 minutes

The CFO's Guide to AI and Machine Learning

These advanced technologies have possibilities for your business. Here's what you need to know.

Introduction

After what seems like decades of promise, artificial intelligence now presents a reality that both offers more than we expected and yet seems more dangerous than we foresaw. Within the short availability of ChatGPT, it's managed to shake up academia and intrigue and frustrate business leaders. It's been banned in Italy even as its creators and competitors release new, more powerful versions on a monthly basis.

While ChatGPT has caught the attention of the masses, other instances of AI and, more often, machine learning have made businesses more productive. With the buzz near its peak, we surveyed CFOs to learn where they think AI fits in a finance practice. Their feedback provided interesting ideas and laid bare misconceptions about what the technology offers and how it works.

That inspired us to offer guidance on how various AI and machine learning systems function and ideas to harness them for business benefit.

Early Innovations for Business

Some of the most useful AI innovations are downright dull. For instance, algorithmic approaches to optical character recognition peaked near 95% accuracy on typical business documents, with handwriting recognition far less accurate. When these systems looked at 3,000 characters on a typical page, the algorithms got 150 of them wrong. That left a lot for humans to fix — so much for relieving the drudgery of moving data from paper into the digital domain.

Today, machine learning techniques have improved algorithms. Better optical readers also help. With the aid of deep learning, [some OCR techs now claim 99.8% accuracy](#), depending on the source material.

That near-100% rate makes automating business processes like bill capture a no-brainer for many companies.

At the other end of the complexity spectrum, AI tools aimed at enterprises with gigabytes to terabytes of data may be enhancements to data warehouse and business intelligence systems, or they may be standalone products. Particularly in the latter case, adoption can be a huge commitment — one that requires a substantial ROI promise.

No one implements these products “just to see what AI is all about.” But for companies with enough data about customer behavior or other key metrics, AI-assisted systems can uncover unique insights and correlations not easily found any other way.

For firms smaller than the Fortune 2000 or not in highly data-intensive industries, it may seem like the basics of reading and classifying business documents and transactions has been the extent of AI’s benefits. And that indeed is a huge payoff, as it digitizes data and lets leaders start to automate rote, tedious tasks.

However, there are intriguing use cases that involve training AI systems on huge datasets and then applying the learnings to individual businesses. A good example: systems that predict how changing the price of a

product will affect retail sales and customer satisfaction. Determining competitive pricing used to be reasonably straightforward, even if it took some legwork. You could go a long way just by visiting competitors’ stores and seeing what stock was being added or removed and when sales were run.

Now, with online marketplaces and branded ecommerce sites, it’s much more difficult to determine what consumers will do in the face of price fluctuations across channels.

Machine learning systems can gather competitive intelligence by scouring the internet and use it, along with your own data on buyer behavior, to determine price elasticity and predict customer trends, sometimes at a personal level.

Some machine learning systems even help retailers set dynamic prices, maximizing the revenue from goods or services with finite supply, like concert tickets or limited-edition items. What will the market bear for a popup performance in an intimate venue, [a hand-painted bathtub](#), or a 50-year-old bottle of Scotch?

Behemoths of the Cloud

Many consumer AI innovations have come from the largest tech companies, and we’ve all watched them improve over time. Siri debuted with the iPhone 4 in 2011, and Amazon released Alexa in 2015. Google Assistant followed a year later. Search engines have benefitted from AI in several ways, not the least of which is the ability to correct our collective poor spelling. Now, email apps are spotting grammar issues and offering fixes or completing sentences on the fly.

These are examples of AI systems that learn and get more accurate over time.

As an area of research, AI has been greatly aided by vast and powerful cloud computing environments. Every major cloud provider now has infrastructure offerings and AI software libraries that serve as the basis for creating new AI products and facilitating AI system training at a more palatable cost.

Still, most commercial AI systems require a lot of computing horsepower — so much so that commercially viable AI products need to provide a major return on investment. In the voice recognition examples we provided above, the return is millions of people engaging on a regular basis and providing lots of personal data. In science, the return needs to be major discoveries that couldn't have practically been made otherwise or assistance with tasks that are too time-consuming and expensive for humans to do alone.

For instance, [pharmaceutical researchers now use AI systems](#) to simulate millions of chemical compound interactions in hopes of developing new drug therapies. The AI systems involved are big, complex, and expensive, but they can model countless interactions in minutes and therefore are a logical alternative to the usual lab testing of a much smaller number of compounds selected based on currently known science. By quickly finding the 10 or 20 protein combinations that warrant human trials, AI has led to creation of drugs that may never have been found otherwise. It's a very big return on a very big investment.

Machine Learning or Artificial Intelligence?

Machine learning (ML) is a huge step toward artificial intelligence (AI), but it's not the same thing. You can show an ML system a few hundred thousand of anything from X-rays to French-to-English translations; the more “good” and “bad” examples you show the system, the more it refines its understanding of the subject.

A ML system uses models and probabilities that are refined as it sees examples to become proficient in a task.

The Origins of AI

For a solid technical look at the evolution of artificial intelligence, check out the [Wikipedia page on AI](#). It does an outstanding job documenting the history, nomenclature, and direction of AI, a computer science discipline that started in the 1950s and has progressed in fits and starts, borrowing from mathematics, statistics, economics, and even philosophy to get where it is today.

Early attempts were disappointing for two, related, reasons. First, the goals were lofty, tending toward creating a system that displayed human-like artificial general intelligence, rather than solving specific problems. Second, the computers available throughout the second half of the 20th century weren't powerful enough to support the goals of many projects. Even [DARPA has its patience and funding limits](#), so research progressed at an uneven pace until the past 25 years or so when various practical applications started to appear.

Through its history, the theory of AI was well ahead of the capabilities of computing hardware. That's been changing for the past decade or so. Hardware capable of delivering on lofty AI theories isn't cheap, but it is within the grasp of larger tech companies.

The key tenet of machine learning is that the system adapts and improves as it's given more examples to analyze. Once engineers create a ML system for a given purpose, they train it by providing examples of the sorts of items they want the system to evaluate. Engineers tweak and perfect the algorithm and eventually end up with a system that's very useful in evaluating, say, X-rays. [Providing radiologists with ML tools can lead to faster, better, more economical diagnoses.](#)

In machine learning, the system produces better results as it sees more samples, but the algorithms that do the learning don't change unless humans tinker with them. The system we've described doesn't know anything about cancer research, but it does know, statistically, what lung image aberrations look like. The algorithms use pattern matching and probabilities to guide findings, and they're highly effective.

In that learning process, ML systems need a lot of data.

Problems that have many samples and classifiable outcomes are good candidates for machine learning to solve. Those with computer-readable examples are especially ideal. Take spam filters. Coming up with training data is as easy as digging through any raw email stream hitting a busy organization's servers. If you have a few thousand email users, a machine learning system could become relatively good at spotting spam by looking through several months of messages sorted by "spam" and "not spam."

One challenge for businesses is determining whether the problem you're trying to solve creates enough data for an AI system to adequately learn, and whether that data accurately describes at all times the condition you want to test. So, in the above example, the characteristics of spam might change somewhat, but good emails will mostly continue to look like good emails, and bad ones will be relatively easy to spot. Major email vendors now claim a 99.9% success rate in identifying spam.

But let's say you want a ML system to tell you whether your company's electric bill is higher than it should be. When you start programming, perhaps the system has access to bills from the previous 12 months to learn from. So the system will have some idea of how costs vary depending on seasonality. What bad data would you give it? Perhaps bills that are 20% above or below the previous year's bill for each month would be considered bad. But then your business grows, you add more equipment, people, and computers, and the cost of electricity changes, maybe by more than 20%. The system has no basis to understand this context, so it flags all subsequent bills for human review.

It's Not AI Just Because Someone Says It Is

Because of the current fascination with AI systems, there's [a temptation to label smart algorithms as AI](#). This is a problem if ROI depends on the system scaling in a certain way or improving as it's given new data. Asking questions about what data was used to train the system and how it will learn from your data will often help identify whether you're dealing with a true ML/AI system.

Media outlets bear some responsibility for calling technology AI when it isn't, and then others repeat those claims. An example is technology from Hawk-Eye Innovations used by the professional tennis circuit to determine whether a ball is within the lines and by Minor League Baseball to call balls and strikes. Many articles have labeled the technology as AI because it uses a series of cameras and some very good computer programming to do what line judges and umpires used to do.

One place that doesn't call the system AI is the [Hawk-Eye website](#). In both use cases, human judgment is what you want to eliminate. A tennis ball is either in or it's out. Strikes are strikes and balls are balls. Calling both is a perfect use of a non-learning system. That the technology can do the job better than humans doesn't make it AI. By that standard, robot welders, painters, and even car washes would be AI, since they can all perform jobs once done by humans and do them faster and more efficiently.

Your monthly electricity bill doesn't generate nearly enough data for a machine learning algorithm to provide insights. You'd be far better off with an ERP system that could automatically look at last year's bill and see if this year's bill differs. You could set a business rule that says to notify facilities if a bill varies by more than 20% from the previous year. It's still going to flag the bill after you've added new equipment and people, but you've only invested 10 minutes to set a business rule versus spending significantly on a ML system that was never going to yield insights you didn't already have. And, you can easily adjust the business rule to reflect expected growth.

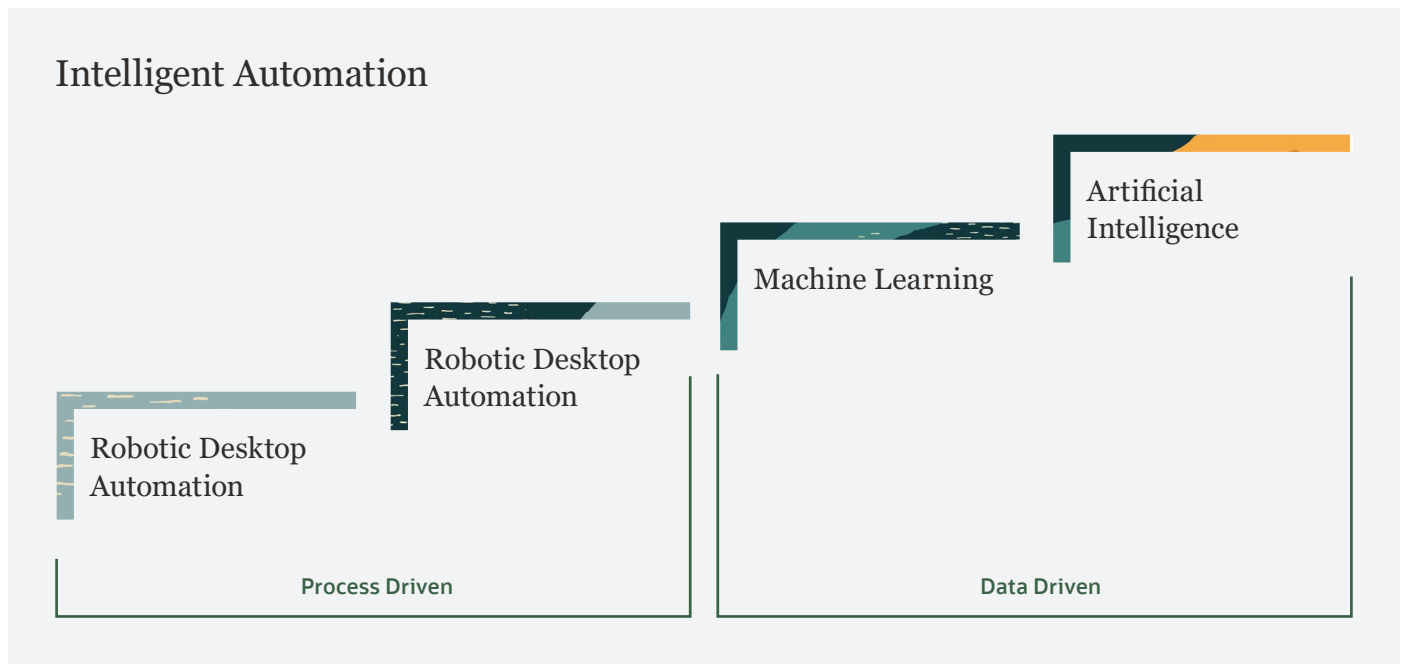
AI and ML in Finance

As AI and ML technology matures, it's more commonly being embedded in business applications. As we've discussed, for CFOs, one of the first benefits is the ability of systems to recognize, interpret, and classify business documents, storing their contents as data. [Accounts receivable and accounts payable automation](#) systems use these technologies to digitize and classify paper and digital invoices. The process is considered robotic process automation (RPA) because once documents are read, they are further processed according to business rules you set.

The diagram below makes a distinction between process-driven automation, which can include everything from spelling and grammar checkers on the desktop to those image recognition/optical character recognition systems we mentioned earlier, and data-driven automation, which depends on ML and AI to offer insights and guide decisions.

Robotic process automation happens when finance team members use a system that knows and follows business rules to accomplish tasks. For expense reports, if you're using a system that captures receipt images, classifies expenses, and enforces reimbursement rules, you've got a good example of RPA.

In our findings from CFOs, many want to leapfrog from limited automation to artificial intelligence. But you can't skip the process automation phase because this is where you teach systems the business rules to follow with the data they take in. Without automation, there's no way for the AI system to know what to do — it needs both digitized data and knowledge of your processes to accomplish a desired job on its own.




Many teams, for example, spend a lot of time executing repetitive tasks. Whether it's managing accounts receivables or payables, three-way matching expenses, running payroll, closing the books, or any of many other monthly functions, automation is an affordable way to save tremendous resources while enforcing your business processes. Not only is that helpful for finance efficiency, it's essential for AI, since AI systems can work only with digital data.

Right now, automation is the place to start if you're looking to channel finance team resources to more strategic tasks, like scenario or demand planning, FP&A, and other data analysis.

The quickest way to reduce days sales outstanding (DSO) is to automate most of the work that goes into accounts receivable. Billing and collection actions happen more quickly and predictably, the system generates the data you need to get constant updates on how your DSO is tracking, and you can get early warnings on accounts that are pushing DSO in the wrong direction so you can deal with problems early on.

Once a finance team has defined processes and digitized business data, automated tasks can follow the business rules required to complete them. Now it makes sense to start considering opportunities to use data-driven machine learning and AI to further improve or even automate operational decisions.

But before diving into data-driven intelligent automation, pause to take stock and avoid mistakes you may have made in the past.



The market for data-driven AI is shaping up much like the business application market has over the past 30-plus years. Vendor A might excel at AI-assisted supply chain management. Vendor B is great at customer sentiment analysis. Vendor C can apply AI to finance and accounting. Vendor D has a great solution for AI-assisted inventory management. You get the idea. Engaging with all those vendors leads to lots of standalone systems that each optimize exactly one piece of the business management puzzle.

Going challenge by challenge to find ML systems that can help is daunting and probably no wiser than it was to go problem by problem to find applications to help solve business challenges.

But don't discount obvious wins, particularly if your business is slowed by one intractable problem. Just bear in mind that data-driven intelligence is often best served by collecting many data streams into a single data warehouse that's capable of complicated "what-if" data analysis and that offers techniques like data visualization along with ML tools. Combining finance data with operational data, web analytics, lead-generation data, store or warehouse traffic information, customer satisfaction metrics, and other business insights lets analysts uncover trends unique to your company that wouldn't otherwise be discoverable.

The key is to keep the advantages of process automation as you move into data-driven analysis.

Ideally, specialty systems that help with a particular aspect of business can easily — that is, without a lot of custom code or ETL work — use the data store in which you've chosen to keep your business data. Look for those that do.

Here's where ChatGPT stirs the imagination. It's expert in a lot of things, maybe enough so that it or technology like it could subvert this whole evolutionary journey from paper to digital to process automation to data-driven AI nirvana.

Enter the world of deep learning.

Go Deep

In ML, the algorithm remains constant unless a programmer intervenes while the system improves as it is shown more examples through guided learning. The next step is to let the algorithm develop as it learns. This is the domain of artificial neural networks that seek to use huge arrays of computers to mimic the function of the human brain in a process known as "deep learning."

In a neural network, the idea is to create layers of artificial neurons that solve a problem by successively getting into more detail — pretty much how the human brain works.

Think of it as each layer of the network analyzing one feature of the data it's learning. If it's looking at pictures of humans, the first layer might determine edges and lines within the picture, while the next layer understands more about contours. Eventually you get far enough into the layers of the network that it can identify eyes, noses, and ears. If the goal is facial recognition, the algorithm refines itself so that it can identify individuals after seeing many millions of varied images of millions of people.

Just like our own brains, the facial recognition process needs to work on massive amounts of data in parallel to get results quickly enough to be useful. The CPU of a typical computer, like your laptop or phone, might have eight cores, meaning it can process eight streams or run eight algorithms simultaneously. It does so very quickly, but in a task like image processing, data needs to be processed in parallel — not eight streams, but 100,000. In a human brain, millions of neurons may be working at the same time to allow us to instantly recognize the things we see.

A revelation in 2012 was that we could use graphical processing units (GPUs) rather than general-purpose CPUs to analyze pictures. That technology, combined with very large data sets and some algorithmic advances, particularly in the learning phase, led to vastly improved accuracy in identifying images after training.

ChatGPT-3 Training Data

Dataset	Number of Tokens	Proportion Within Training
Common Crawl	410 billion	60%
WebText2	19 billion	22%
Books1	12 billion	8%
Books2	55 billion	8%
Wikipedia	3 billion	3%

Data: [Wikipedia](#)

GPUs have hundreds to thousands of cores and can access memory very fast, meaning they can efficiently process lots of data in parallel. In many ways, they're ideally suited to neural net designs. In particular, they're far better suited to analyzing the very large training sets than are CPUs. Datasets that can be analyzed in days by GPUs would take years with the same number of CPUs.

Over the past decade, Nvidia and other vendors have refined their architectures to create chips specifically designed for deep-learning applications. Now, deep learning is leading to the sorts of mind-blowing results we see from ChatGPT. The generative AI system can pass medical and bar exams because everything it needs to know to do so can be found on the internet and in other learning datasets. It's also a very good coder and writes in clear English with few errors.

Reduce to Practice for Business

ChatGPT works on a huge scale. Microsoft reportedly spent hundreds of millions just to create the computing infrastructure where the system was trained and now

lives. So far, details about GPT Version 4 have not been released, but the Version 3 training dataset was roughly 3 terabytes of text, as shown above.

To put this in context, the King James version of the Bible runs 4.13 megabytes, meaning the algorithm studied text amounting to more than 800,000 King James Bibles. In that training set listed above are medical and legal texts along with the many, many programming language examples and explanations on the web plus a lot of general-knowledge content.

And while V3 was trained only with text, V4's training will include images.

You can have a conversation with ChatGPT-3. It will answer questions and, beyond that, build on previous responses. For purposes of answering basic questions and writing content for business or other purposes, it's fairly good. If you need a refresher on business practices — say, the pitfalls of intercompany transactions — it can provide that.

We asked ChatGPT a question. Here's its response.

Q: Can you tell me 5 things to watch out for when accounting for intercompany transactions?

A: Sure, here are 5 things to watch out for when accounting for intercompany transactions:

1. Proper documentation: All intercompany transactions should be properly documented and recorded. This documentation should include invoices, purchase orders, and payment receipts.
2. Elimination of intercompany transactions: Intercompany transactions should be eliminated from consolidated financial statements to avoid double-counting or misrepresentation of the company's financial position.
3. Intercompany pricing: The appropriate pricing for intercompany transactions should be determined based on market rates or other industry standards. This helps to ensure accurate financial reporting and avoid tax implications.
4. Intercompany loans: Intercompany loans should be properly structured and documented, with clear repayment terms and interest rates.
5. Legal and regulatory compliance: All intercompany transactions should comply with applicable legal and regulatory requirements, including tax laws, accounting standards, and transfer pricing regulations. Failure to comply with these regulations can result in penalties and legal issues for the company.

That's a pretty good list of concerns, and of course you can dig further. But this version of ChatGPT can't audit your books and flag where intercompany transactions haven't been handled properly because it doesn't have access to your books and wouldn't inherently know how to read them.

Finance and other business leaders would love an AI system that checks things like intercompany

transactions or, better yet, one that checks by following the FASB and GAAP rules that ChatGPT seems to understand.

Asking ChatGPT to describe how to comply with accounting rules is one thing. Having a system like it that applies those rules to your books is, unfortunately, another.

Here's what ChatGPT has to say about ASC 842:

Q: What steps should I take to ensure my company is ASC 842 compliant?

A: ASC 842 is an accounting standard that affects the way companies report leases in their financial statements. To ensure your company is ASC 842 compliant, you should take the following steps:

1. Proper documentation: All intercompany transactions should be properly documented and recorded. This documentation should include invoices, purchase orders, and payment receipts.
2. Elimination of intercompany transactions: Intercompany transactions should be eliminated from consolidated financial statements to avoid double-counting or misrepresentation of the company's financial position.
3. Intercompany pricing: The appropriate pricing for intercompany transactions should be determined based on market rates or other industry standards. This helps to ensure accurate financial reporting and avoid tax implications.
4. Intercompany loans: Intercompany loans should be properly structured and documented, with clear repayment terms and interest rates.
5. Legal and regulatory compliance: All intercompany transactions should comply with applicable legal and regulatory requirements, including tax laws, accounting standards, and transfer pricing regulations. Failure to comply with these regulations can result in penalties and legal issues for the company.



It understands the regulation, so creating a system that does the work seems well within reach. ChatGPT is already advising you not to try it without software. Tech companies are scrambling to release products that leverage ChatGPT by building systems that access the technology through APIs or that use similar technology. Those systems could read and interpret your books and apply an AI-driven accounting tool to comply with FASB rules and your business practices, though the technology used probably won't be generative AI.

Your AI accountant is somewhere on the horizon.

The question will then be: How much do you trust your AI accountant, and how will you review its work to ensure the rules you've set for managing your books are followed?

CFOs and controllers will start out with caution, so understanding how to monitor work AI does for you will be a big part of adoption.

Ethical and Practical Concerns

Ethical concerns abound, particularly as AI systems start making decisions that affect lives in substantial ways or create output that encroaches on the copyrights and intellectual property ownership of humans.

Discussions of these issues are important, complicated, and nuanced. Ethicists, lawmakers, and technologists are raising warnings about potential issues that could arise from a headlong rush toward widespread use of AI in all sorts of applications. An [open letter encouraging a pause](#) in development of "Giant AI Experiments" starts out with these two sentences: AI systems with human-competitive intelligence can pose profound risks to society and humanity, as shown by extensive research and acknowledged by top AI labs. As stated in the widely-endorsed [Asilomar AI Principles](#), Advanced AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources.

The letter has been signed by more than 28,000 people so far, including many leading thinkers in technology, business, and academic research. Many of the signers have pointed to the unintended negative consequences on society of unregulated social media and postulate that unregulated AI could have a much greater effect.

How will this affect your industry? Here's a short list of papers to consider:

- [AI in healthcare](#): As in many areas, some of the lead concerns in healthcare revolve around whether existing legal structures will be sufficient to assess responsibility for the actions of AI systems.
- [AI in law enforcement](#): Two questions tend to dominate law enforcement concerns. The first is around privacy. AI has the potential to use huge amounts of data to let law enforcement track what citizens are up to. There's an obvious potential for abuse of such power. The second concern revolves around implicit bias. Training sets may not include enough data about minority groups. As a result, false positives, including in facial recognition, are more likely in those groups, possibly leading to improper actions by law enforcement.
- [AI ethics in the legal profession](#): Lawyers have a professional obligation to fully understand and agree with and stand behind the output of their practice, whether it's creating contracts, offering opinions on disputes, or citing case law that would guide actions on behalf of a client. AI offers a way to much more quickly reach legal opinions. However, particularly in deep-learning systems, the tools operate as a "black box," which is to say the algorithm used to create output is purely mathematical and therefore not intelligible to humans in a way such that we can know if the output was derived "in the right way." As a result, lawyers can't just trust what AI tells them. Vetting is required.

- [AI ethics in banking](#): Here again, the black box nature of systems leads to concerns. AI will advise bankers on issues like whether clients are suitable for loans, or whether the bank's risk portfolio is well managed. But without understanding how conclusions were reached, it's impossible to know whether implicit bias played a part. In more complex matters of risk management, AI systems could propose portfolio structures that are so complex that bankers and auditors cannot assess whether they'll meet the risk goals of the institution.

This very short list provides a sampling of the questions that must be addressed industry by industry and use case by use case. [Stanford University offers a paper on AI ethics](#) that explores the general topic, and [Harvard Business Review offers another on how AI might be regulated](#).

The Bottom Line

We're only beginning to understand the ramifications of AI, both good and bad. There are many, many questions to be answered about its use, but the promise is alluring. Vendors won't necessarily have all the answers, so smart decision-makers will carefully consider their choices and watch efforts to restrict what AI is allowed to do.

