

ORACLE
NETSUITE

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.

Learn more at [NetSuite.com](https://www.netsuite.com)



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.

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.


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.



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.



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.

