Artificial intelligence and machine learning: It's still about the data

Federated Hermes

Federated Hermes MDT Advisers

Damien Zhang, CFA
Head of Research

Key takeaways

- **Teaching machines how to "think" has been a goal since the early days of computing, and recent advances in generative AI (artificial intelligence) have shown how close we may be to achieving it.**
- **ChatGPT and the interest it spawned is the latest step in the broader field of artificial intelligence, which has had a long and colorful history.**
- **Generative AI has substantial promise, but the current generation of models poses significant challenges for use in investing.**
- **As this field evolves, MDT Advisers will seek to identify ideas worth importing from AI research to improve our investment process.**

As a leader in the use of machine learning in the investment field, MDT Advisers (MDT) has been a quantitative investing practitioner since the 1980s and a proponent of technology in stock selection and portfolio construction. The recent excitement around generative AI applications like ChatGPT has led to many questions about our thoughts on AI in general and its application in investment management. We address some of those questions below, but always welcome the opportunity to discuss technological advances, especially as they relate to investment management.

Teaching machines how to "think" has been a tantalizing goal since the early days of computing, and recent advances in generative AI have brought us closer than ever to achieving it. By leveraging vast datasets and computing power, researchers have developed models that enable computers to have conversations, draw pictures and perform other human-like tasks that were thought impossible even a few years ago.

Generative AI did not appear out of nowhere. It is the latest step in the broader field of AI, which has had a long and colorful history of progress and setbacks. Some early successes in the lab failed to materialize as practical applications. Some promising ideas led to dead ends. Some seemingly dead ends, such as the humble Perceptron, were later resurrected as core components of the neural networks that undergird modern generative AI models.

Many early attempts at AI centered around rules programmed by humans. These approaches ran up against the difficulty of managing the complex tangle of rules that inevitably result from trying to navigate the nuances of the real world. What has instead become the dominant approach is to have a machine learn by distilling the rules from large datasets containing inputs and desired outcomes. The result of that learning can be encoded as an equation, a forest of decision trees, a neural network or any other flexible model. The better the model is at capturing the true nature of the rules at play, the more useful that model is.

Generative AI is the result of this latter approach taken to the extreme. With more data, more computational power and larger models than ever, we have not yet reached the limits of this approach. But even as generative AI provides the most compelling proof of the effectiveness of machine learning, it also illustrates some of its pitfalls.

As generative AI upends our notions of what a machine can do, it is worthwhile to explore how it could be used to enhance our own field of investment management. At MDT, we have a long history of incorporating machine learning innovations while seeking to avoid the pitfalls.

Artificial intelligence

- A field of computer science that covers a wide range of algorithms and approaches, with many subfields
- Focused on the development of machines that mimic functions associated with the human mind and that can perform human tasks such as understanding speech, playing games and driving cars
- Colloquially, "AI" is often used as a nebulous term that encompasses big data, machine learning or artificial neural networks

The generative AI revolution

Much of the current buzz around AI stems from the release of ChatGPT in 2022. Its remarkable ability to converse and respond to a wide array of topics with seemingly human-like intelligence has provoked questions from practicality to ethics. Can it do my homework? Can it take over my job? Does it exhibit consciousness?

Under the hood, the original ChatGPT used a variant of a large language model (LLM) called GPT-3.5, which was trained on a dataset containing hundreds of billions of words collected from the internet. Perhaps surprisingly, given its apparent capabilities, the model is trained only to predict the next word in a sentence as accurately as it can. Some have even called it a glorified autocomplete. Though complex in the details, the training itself amounts to having the model guess the next word, checking how big of a mistake it made, and adjusting some of its billions of model parameters ("neurons") so it makes a smaller mistake in the future.

Yet this contrast is emblematic of the success of machine learning's data-driven approach. The simple high-level technique of applying massive amounts of data to train a very large model made of simple individual components can lead to impressive results — generally, the bigger the model, the better the result — if there is enough data.

Just as a model can be trained to predict the next word, it can also be trained to predict the next sound, the next pixel or some other combination, leading to audio, image and multi-modal generative models. These models are also improving, and a new generation of open-source models now rivals even the best proprietary implementations.

One particular strength of generative LLMs is their ability to summarize long blocks of text, which could help investors parse lengthy regulatory filings for critical bits of information. A less obvious but more intriguing possibility is for the model to distill that summary into a rating or other quantitative metric.

For example, suppose we would like to assign a sentiment score to a news article about a company. Historically, such tasks have used rudimentary machine learning models or counted the number of positive- and negative-sounding words. Now these tasks can be accomplished by telling the LLM to read the relevant article and then asking it to respond with a number for sentiment. Once you know what you want to see, you have the potential to use generative AI to unlock that information from regulatory filings, earnings calls, satellite photographs and other kinds of unstructured data. While this has a lot of promise, there are problems as well.

The "but"…

One only needs to read a few passages of AI-written content to notice that it is a simulacrum — when asked to write an essay, the program may line the words up, but the text may lack humor, nuance or deeper insight. It can struggle to solve simple math problems. Sometimes the model will even come up with an excuse to avoid doing what you asked. Missing the mark can be harmless and amusing if there is nothing at stake, but investing based on unverified AI input presents a real problem.

These AI systems can sometimes generate misleading or factually incorrect responses with complete conviction. Termed "hallucinations," these responses continue to afflict large AI models in part because they are trained on vast swathes of human-generated data on the internet that is itself replete with inaccuracies presented as fact. In this regard, the models are remarkably human-like. While some adjustments can be made to the models after they are trained, it is impossible to manually check the entire training data to remove inaccurate information.

Whether AI-generated signals can be used for investing is partly an empirical question of whether the errors they make in hallucination or incorrect interpretation are outweighed by the value of the signals they extract.

Additionally, the reasoning process that these AI models use to arrive at answers is a black box. You could ask a model how it arrived at an answer, but as the model is incapable of genuine introspection, the answer will still be based on a best guess of the likeliest next word. Thus, if an AI-generated signal indicates a counterintuitive trade, it will be difficult to understand the true chain of reasoning that led to that trade.

Finally, for systematic investors, backtesting is an important tool that gives essential information about the effectiveness of a strategy. Generative AI models are trained on data up to the present day, and their enormous sizes mean that they can memorize a great deal of historical information, including market information. Using a signal generated from such a model risks look-ahead bias: contaminating simulated trading in a past time period with information from the future.

Machine learning

- Emerged as a separate field at a time when it was difficult to secure funding for artificial intelligence research
- The goal of machine learning is to produce the most accurate prediction given the available data
- Typical techniques include decision trees and deep neural networks

Our approach at MDT

Generative AI holds substantial promise, but even the current generation of models faces significant challenges for use in investing. However, it is possible to avoid the pitfalls while reaping substantive benefits from advances in AI.

MDT has long employed machine learning in its investment modeling process, but we have insisted on a transparent and accountable process. Decision trees and the technology we have built around them allow us to understand and explain the rationale behind every trade. A commitment to maintaining high-quality historical data enables us to build models that seek to accurately replicate history and make unbiased forecasts.

At the same time, some of the techniques developed in support of training generative AI models can be useful beyond their original scope. Some of these methods can be imported to train decision-tree models and can improve the accuracy of their forecasts. Generative models, for instance, are trained using an optimization method called gradient descent,* which we have also found useful in helping us refine the boundaries of the market segments where the best alpha (excess performance) opportunities can be found.

Generative AI models also feature tens to hundreds of billions of parameters, so despite the substantial dataset on which they are trained, they still require a range of techniques to help them avoid overfitting. Core among these techniques are regularization (mathematical penalties to reduce model complexity) and randomization (introducing noise to improve the model's robustness to unseen data). We have found both techniques useful in our context as well.

By selectively adopting the best and most relevant ideas from the rapidly progressing AI research, we are positioned to preserve our core model's accountability and transparency while benefiting from the groundbreaking work.

Looking forward

What's remarkable is the speed at which innovation occurs: What's true about AI and its applications today is unlikely to be true tomorrow. We routinely evaluate new advances in machine learning and incorporate effective ideas. We will continue to explore AI to enhance the alpha-generating signals in our portfolio construction process. As this field evolves, we will seek to identify ideas worth importing from AI research to improve our investment process.

*At the heart of gradient descent is a mathematical function that measures how well the model's predictions align with real world observations. Larger gradients indicate areas with greater potential for learning, and the descent algorithm prescribes how to improve the model step by step by making small adjustments in the direction of greatest potential improvement.

Artificial intelligence and machine learning: It's still about the data

Federated Hermes

A brief history of AI

- **1956:** "Artificial intelligence" is coined by John McCarthy
- **1960:** The Perceptron, an early precursor to neural networks based on biological principles, is invented
- **1960s:** DARPA provides open-ended funding for AI research
- **1970s:** Funding is cut to AI research after a "web of increasing exaggeration," causing the first "AI winter"
- **1980s:** Development of "expert systems" using handcrafted rules sparks new interest in AI, promising to make knowledge "easily replicated [...] and essentially immortal"
- **1980s:** Invention of the "backpropagation" algorithm enables training of "artificial neural networks"
- **1990s:** Expert systems prove too difficult to maintain for most business uses. Most commercial vendors fold or exit the business, leading to a second AI winter.
- **1992:** IBM develops TD-Gammon using reinforcement learning and an artificial neural network
- **1997:** IBM's Deep Blue, using a brute-force classical AI approach, defeats chess world champion Garry Kasparov
- **2011:** IBM's Watson question-answering system wins *Jeopardy!*
- **2017:** Google DeepMind's AlphaGo, using reinforcement learning, a deep artificial neural network and classical AI techniques, defeats the Go world champion
- **2017:** Google publishes the transformer architecture, laying the groundwork for the next generation of AI models
- **2020:** OpenAI (partnering with Microsoft) releases GPT-3, a state-of-the-art large language model (LLM) built on the transformer architecture scaled to billions of parameters
- **2021:** Microsoft releases GitHub Copilot, an AI code assistant, one of the first commercial products to use GPT-3
- **2022:** IBM sells Watson Health unit after struggling with profitability
- **2022:** Stability AI and OpenAI release groundbreaking text-to-image models Stable Diffusion and DALL-E 2
- **2022:** OpenAI releases ChatGPT, bringing a chat interface to a version of GPT-3 enhanced with human feedback through reinforcement learning
- **2023:** Competition heats up in LLMs as OpenAI releases GPT-4, Google releases Bard, Microsoft releases Bing Chat and Meta releases LLaMA

The quantitative models and analysis used by MDT may perform differently than expected and negatively affect performance.

Investing in equities is speculative and involves substantial risks. The value of equity securities will rise and fall. These fluctuations could be a sustained trend or a drastic movement.

Past performance is no guarantee of future results.

G85114-06 (5/24)
MDT Advisers, a Federated Advisory Company