CHIEF INVESTMENT OFFICE

Educational Series

Decoding Artificial Intelligence and its Potential Asset Management Applications

April 2024

Artificial Intelligence (AI): What it is and where we are

The idea of creating a machine that could not only learn, but continuously adapt and evolve like human beings, has existed for some time and has been a staple of science fiction since the start of the 20th century. The transition from fiction to reality can be traced to the efforts in cryptography during World War II. The idea of designing computer systems that mimic neurons in a biological brain paved the way for the development of artificial neural networks. The Turing Test, named after Alan Turing, determines whether a computer can think like a human being. Turing proposed that a computer can be said to possess AI if it can mimic human responses under specific conditions. The original Turing Test required three terminals, each of which is physically separated from the other two. One terminal is operated by a computer, while the other two are operated by humans.

Advancements in computer processing, new algorithms and big data this century have enabled the capabilities we see in today's latest version of AI.

A common characteristic of AI workloads is that they are both data- and computer-intensive. A typical AI workload involves a large, sparse matrix computation distributed across hundreds or thousands of processors Central Processing Units (CPUs), Graphics Processing Units (GPUs) or Tensor Processing Units (TPUs). These processors compute intensely and then exchange data with their peers. Data from the peers is reduced or merged with the local data and then another cycle of processing begins. In this compute-exchange-reduce cycle, any slowdown can detrimentally impact the job completion time.

While AI has been around for decades, the recent spike in enthusiasm for the technology has been caused by the rise of Generative AI (GAI) and its derivative products. In stark contrast to traditional AI, which depends on explicit programming, GAI derives insights from data, enabling it to create imaginative and unique outputs. Its capabilities extend to generating diverse forms of content, from images and music to text and even entire virtual environments.

The main difference between traditional AI and GAI lies in their capabilities and application. Prescribed rules are pretty much gone. Instead of telling the model exactly how to tell a story, design a logo, or compose a song, the model figures out the rules itself. It achieves this by discovering and learning patterns from gazillions of data points. The most well-known example is ChatGPT, which was trained on text (where most but

TRADITIONAL PREDICTIVE AI	
What It Is	• Algorithmic processes that identify patterns and trends in structured, time series data sets that might predict future behavior • Rules-based programs to perform specific tasks
AI Strenghts	• Classification, anomaly detection, prediction • Solve analytical tasks faster • Prediction, traceable, controllable
Applications	• Quantitative analysis • Risk modeling • Predictions • Portfolio optimization • Robotic Process Automation (RPA) • (Rules-based) Conversational AI

GENERATIVE AI	
What It Is	• The application of foundation models in order to create original content • Uncovers hidden patterns that may not be evident through predictive AI
AI Strenghts	• Unstructured data, which is then classified • Creating new content • No/few-shot training
Applications	• Question answering • Sentiment analysis • Information extraction and search • Content generation • Object recognition • Summarization

Source: Bain Capital Ventures, "*Sharing Our Field Notes: The State of Generative AI in Financial Services*", June 22, 2023.

6506938 4/2024

not all of the training data came from the internet), although generative AI models can be built on other data formats such as image, audio or video. Basically, AI excels at pattern recognition, while GAI excels at pattern creation.

The rise of GAI, in turn, has been driven by significant advances in two related underlying technologies: Natural Language Processing (NLP) and Large Language Models (LLMs).

NLP is a computer's ability to receive instructions and give outputs to human users in conversational language as opposed to specialized programming inputs. A key goal of NLP is the understanding of context and sentiment—having the computer differentiate between similar blocks of text based on the text surrounding it. NLP technology has also been developed over the course of decades, with models becoming ever-more sophisticated, going from lists of rules (Symbolic NLP) to machine learning-powered statistical algorithms (Statistical NLP) to deep neural networks (Neural NLP). Each advance allows the NLP models to use more data to further improve their outputs. Key use cases for NLP tech include optical character recognition and text-to-speech products, language translation, and text-to-video programs among many others.

LLMs are massive NLP models, based on neural networks, that utilize training sets of billions of inputs (parameters) to serve as the basis for the model. The LLM then (typically autonomously) ingests that data and makes inferences about the relationships of the inputs to each other, which can be used to complete tasks. Given the vast scale of inputs and the broad applicability of the rules created by the LLMs, they can be used for a wider variety of tasks than a typical NLP model. One of the first major LLM launches was Google's BERT in 2018; since then, LLM releases have come from a variety of sources (including Google, Microsoft, Meta and other AI-centric organizations) with ever-increasing levels of sophistication and capabilities.

Applications

The applications of GAI are wide-ranging, providing new avenues for creativity and innovation. In design, GAI can help create countless prototypes in minutes, reducing the time required for the ideation process. In the entertainment industry, it can help produce new music, write scripts, or even create deepfakes. In journalism, it could write articles or reports. GAI has the potential to revolutionize any field where creation and innovation are key.

In contrast, traditional AI continues to excel in task-specific applications. It powers our chatbots, recommendation systems, predictive analytics and much more. It is the engine behind most of the current AI applications that are optimizing efficiencies across industries.

Potential Benefits and Risks

GAI is likely to have profound effects on knowledge workers (and corporate costs) across various industries, but this will not happen overnight. Unlike robotics, which has generally affected blue-collar workers, GAI is likely to mainly effect white-collar/knowledge worker jobs like marketers, software developers, writers, accountants, consultants, graphic designers, customer service reps, law clerks and even equity research analysts. While one can expect some near-term disruption to certain roles as companies assess GAI opportunities, note that paradigm shifts in technology have historically not resulted in a reduction in aggregate employment, but instead have led to the creation of new roles and reduction of labor inputs in certain existing roles. Although still early and emerging on a case-by-case basis, consensus is that efficiency gains from GAI will also drive incremental cost savings for most enterprises over the next several years, providing a tailwind to profit margins.

DID YOU KNOW

In 1956, a group of scientists gathered for the Dartmouth Summer Research Project on AI, which was the birth of this field of research. In the 1960s, the systems could solve simple problems and demonstrated natural language processing in a simulated conversation with a human. By the 1980s, expert systems and neural networks were introduced.

Computer Vision enables machines to identify and classify objects. This may occur in real time or be applied to static images. Facial recognition, used to unlock a user's phone, and mapping features that enable autonomous vacuums to learn and navigate to individual rooms, serve as examples. Driver-assisted and self-driving vehicles also utilize computer vision.

Recommendation Systems provide personalized suggestions based on a user's past choices. Common use cases include video and audio streaming services as well as retail.

Healthcare Applications can help with diagnosis through a combination of medical imaging, patient history and observed symptoms. AI can help with drug discovery through the simulation of chemical processes and analysis of large data sets. AI systems can also assist with patient monitoring.

Financial Services example uses include fraud detection, algorithmic trading and risk assessment.

Manufacturing and Smart Factories can utilize AI in multiple applications. Robots, of course, belong in this list. The use of AI shifts simple mechanical automation (i.e., the conveyor belt) to systems able to make decisions and modify or optimize a process, such as a robot painting a car part. Other uses might lie in predictive maintenance, inventory management and quality control.

While the benefit of GAI is certainly enticing, there are several systemic risks with LLMs today that must be addressed or mitigated before adoption can meaningfully scale. Among them include:

Ethical Concerns	Data privacy/ protection	Availability of training data	Hallucinations	"Cognitive" biases	Sustainability concerns
LLMs can give unethical responses (such as discriminatory or violent responses to queries), as users saw with the original launch of ChatGPT before OpenAI introduced guardrails to the model to limit its answers.	Information entered into LLMs is largely ungoverned outside of the LLM operator, which opens up possibilities for misuse of data, particularly if the data input into the LLM is of a sensitive nature (personally identifiable information, health records, sensitive governmental information, etc.)	Content owners on the internet were largely caught flat-footed by the rise of LLMs and the implications for having their data included in LLM training sets. For site publishers that have their information scraped by LLMs that could use that information for search engine-like purposes, that could mean being disintermediated and having traffic to the site cut drastically.	LLMs will frequently (confidently) give factually incorrect information, and even in cases when information is correct, can be "talked into" giving wrong information through leading user prompts.	LLMs have exhibited many traditional human cognitive biases because of incomplete or skewed training sets or the programmers' own biases seeping into the code. Inadvertent (or intentional) inclusion of cultural, gender or racial biases among other biases are possible.	LLM training, inference and operation is highly computationally intensive. At a time when major tech companies are highlighting their pushes toward sustainability and carbon footprint reduction, the transition to LLMs provides a sizable new source of energy usage.

The question remains, "Is any of this even legal?" The risks associated with GAI seem to be multiplying and include sensitive data exposure, intellectual property rights, data accuracy/bias, and phishing and social engineering attacks—though perhaps the scariest risks are the ones we are not even contemplating yet. In the near term, the big risk is the legality of LLMs that are trained on wide swaths of online data, some of which may be copyrighted. Several lawsuits have already been filed against LLM providers, which, if successful, could hamper the development of the space (think Napster in the early days of online music).

AI in Investment Management

The investment universe is very wide ranging, with many asset managers focused on managing exchange-traded funds, mutual funds, hedge funds, private equity funds, etc. All these firms would be expected to be following the developments in the world of AI and looking to incorporate them into parts of their businesses where it makes sense. AI has the potential to apply to various aspects of an asset manager's operations, from its research division to its middle and back office and even its investor relations and marketing functions. A firm may have added automation to its filling out of requests for proposal in the marketing area, or it may be doing a first pass on resumes it receives in the human resource department with the help of AI tools; but from an investment management perspective, one would be most interested in how AI is used in the investment process, in how assets may be selected for portfolios. And it is quantitative asset managers, those utilizing models and algorithms in the investment process to make portfolio decisions, who may be depending the most on AI and machine learning if that is a core element of the framework they use. (Note that quantitative asset managers typically do not look to create content, at least as part of their investment process, so the latest techniques in GAI are less likely to apply to them. They are much more likely to focus on predictive AI methodologies.)

The investment process incorporates many stages, from making asset return predictions to portfolio construction and risk management. Each stage is important to generating satisfactory performance, but the first stage of determining expectations in terms of what different assets in a manager's universe should return is critical. After all, if one is making poor predictions of whether assets or markets should be going up or down and by how much, good portfolio construction and risk management cannot improve results by much. It is in this stage of the investment process that the use of predictive AI and machine learning techniques can be a meaningful differentiator.

For quantitative asset managers, there are two main routes to use in how expected asset or market returns are estimated. One is traditional modeling, where an analyst does deep quantitative research to understand all the factors affecting the asset's performance. For an individual stock, that analysis would involve collecting data about the company from diverse sources, including information from its financial statements, the industry and sector in which it operates, macro variables affecting it, etc. For a market like crude oil, the input data could include fundamental information about demand, supply, future production, inventory, indicators from the options market, economic indicators related to global growth, inflation, interest rates, etc. The analyst would then look to determine the relationships between these variables and the return of the asset, coding those relationships into a model to make predictions on a go-forward basis. Alternatively, a predictive AI or machine learning framework could be used to derive those relationships by looking at all the data and determining how they influence the asset's return. Both methods are valid, with the latter machine learning approach being the more recent one to be put to use by a number of managers in the last decade or so.

There are advantages and disadvantages to using the machine learning route.

In terms of advantages, with machine learning, one does not need to specify the relationships among the input variables and the return one is forecasting. The algorithm makes those determinations using the data. If the relationships change over time, it has the ability to learn about them as new data becomes available. If the underlying process changes, it can adapt, not immediately, but over time, without the need for anyone to intervene and reprogram the model. It can also uncover more complex relationships that exist within the data, some non-linear, which might be more difficult or even impossible for a human programmer to pick up using traditional modeling methods. In these respects, the AI approach could lead to superior predictions and better performance.

Among the disadvantages of using AI in the asset return prediction phase of the investment process is the fact that it needs a very large amount of data to understand the underlying process in something as complex as an asset or market. The data available to the algorithm can be limited. The data necessary for a machine learning model to understand even a simple, stationary process is meaningful; it takes significantly more data to understand assets and markets where the process and underlying relationships may be changing through time. Additionally, while the steps programmed into a traditional model are clear and one can go back to them to understand a given forecast, a machine learning model cannot explain itself in the same way, making it more difficult to make sense of its output.

Despite these limitations, a small number of asset managers are using predictive AI in that most crucial "asset return forecasting" phase of the investment process. With the availability of more data, those results would be expected to improve. In identifying such managers, one should focus on the people building the models (their background, pedigree, and experience in building machine learning models), their research process (from the acquisition and cleaning of data to designing the algorithms in accordance with the scientific method) and the live performance of their models over time. Due diligence teams looking to uncover skilled asset managers remain optimistic about the prospects of qualified investment professionals with machine learning expertise delivering strong returns in diverse asset classes utilizing a range of quantitative strategies. It is a developing field.

Important Disclosures

Investing involves risk, including the possible loss of principal.

This material does not take into account a client's particular investment objectives, financial situations, or needs and is not intended as a recommendation, offer, or solicitation for the purchase or sale of any security or investment strategy. Merrill offers a broad range of brokerage, investment advisory (including financial planning) and other services. There are important differences between brokerage and investment advisory services, including the type of advice and assistance provided, the fees charged, and the rights and obligations of the parties. It is important to understand the differences, particularly when determining which service or services to select. For more information about these services and their differences, speak with your Merrill financial advisor.

This information should not be construed as investment advice and is subject to change. It is provided for informational purposes only and is not intended to be either a specific offer by Bank of America, Merrill or any affiliate to sell or provide, or a specific invitation for a consumer to apply for, any particular retail financial product or service that may be available.

The Chief Investment Office ("CIO") provides thought leadership on wealth management, investment strategy and global markets; portfolio management solutions; due diligence; and solutions oversight and data analytics. CIO viewpoints are developed for Bank of America Private Bank, a division of Bank of America, N.A., ("Bank of America") and Merrill Lynch, Pierce, Fenner & Smith Incorporated ("MLPF&S" or "Merrill"), a registered broker-dealer, registered investment adviser and a wholly owned subsidiary of Bank of America Corporation ("BofA Corp.").