Invesco

In Silico, Part II

AI and the quiet revolution of machine learning

In Silico is a multi-part series discussing artificial intelligence, its economic and financial impact, and its role as a driver of change.

This article is the second in a series in which we explore artificial intelligence's (AI) explosion of interest in 2023, reviewing the existing impact of AI in today's economy and the innovation that has been happening right under our noses for decades. We also have a special feature from our own Strategy, Innovation and Planning team, highlighting how Invesco already makes use of AI.

Subsequent articles will consider the economic, financial, and social impact of Generative AI and robotics. Along the way we will highlight opportunities for investment, culminating in a comprehensive overview.

Ashley Oerth, CFA®
Senior Investment Strategy Analyst
Investment Thought Leadership

Cyril Birks
Global Thought Leadership Intern
Investment Thought Leadership

In 2023, anything that remotely looks like AI might now be called AI, just to cash in on hype. To understand where the real value of AI is, this article will look at the quiet revolution of AI that happened right under our noses. We will explore what Machine Learning and Deep Learning are and, in turn, lay the foundations for a contextualized understanding of 'generative' AI, which we will cover in Part 3.

From math to thinking machines

AI originates from mathematics, which is how we represent reality with symbols, transforming the physical and social world into a common language that can be manipulated and analyzed systematically. With mathematics, not only can we represent the size of a city, but we can also represent the relationships between each person using equations and functions.

Statistics, meanwhile, gives us the methods and processes to analyze and interpret data, allowing us to describe and predict the variability and uncertainty of our world; statistics gives us methods by which we can infer rules or ask counterfactuals. When combined with computational algorithms, statistics becomes data science, an applied, interdisciplinary field. While formal definitions are contested, data science emphasizes practical, real-world impact and is to statistics what engineering is to physics.

At its core, artificial intelligence is about mimicking some function of human intelligence, such as processing data or making a decision. AI encompasses and is largely defined by Machine Learning, which is a subfield of data science which uses computational algorithms that incrementally improve over time ('learning'). This incremental improvement is the product of an algorithm that updates itself in response to new data. A simple and commonly used example of Machine Learning is a regression, but far more complex examples iteratively improve based on an algorithm. Deep Learning, a further subset of Machine Learning, uses 'neural networks' to understand complex pattens in data by abstracting relationships and gradually learning them over time. It is most useful for processing 'unstructured' data, but it is a computationally intensive task and typically requires significant complexity before achieving useful results. Deep learning is the area of AI that has enabled many of the most exciting developments in AI, which we will explore in this edition.

Figure 1: Mapping Artificial Intelligence

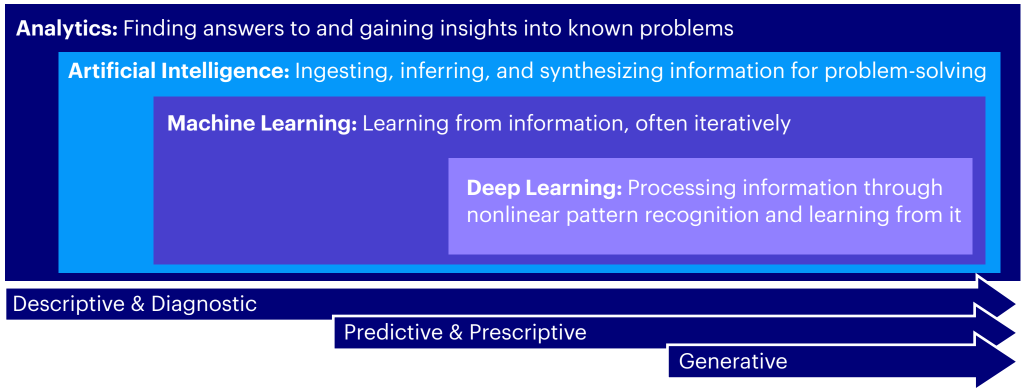

Source: Invesco. For illustrative purposes only.

Commercial applications of machine learning

Where simple algorithms might focus on a very narrow class of problems, Machine Learning describes algorithms which can work in a wide range of contexts and can adapt with 'experience.' Consider the following examples of Machine Learning and their commercial applications.

'Supervised' Machine Learning is used when the types of inputs and outputs are known. For example, you might use a regression to understand how long it will take a courier to transport some set of goods, given road conditions, traffic, and speed limits. Supervised Machine Learning helps us to understand, "How do multiple variables predict X?"

'Unsupervised' Machine Learning is used when the inputs are not described with any particular label ('unlabeled'), and the model then learns patterns and relationships within this data. For example, you may have lots of data about how customers use your website for online shopping but may not understand the types of people that shop with you. An unsupervised Machine Learning solution will not give you demographic features about your customers, but it could cluster customers them into approximate customer segments for further human analysis. You may then choose to label those segments, based on the features of each cluster, and use that data for 'Supervised' approaches to forecast scenarios. Unsupervised approaches might also be used for marketing to current customers based on customers like them and can even be used for complex events like disease modelling or symptom monitoring.

'Reinforcement' Machine Learning is used when you desire a certain outcome, but you cannot define that outcome, and you do not have a lot of training data to use. Reinforcement can be used in cases where 'the bigger the number, the better' applies. For example: content recommendation – including video and music streaming and social media feeds – often relies on Reinforcement Machine Learning. New users of a platform are presented with a wide range of generic options. When users interact with content, their interaction data is analyzed and used as the basis for recommending new content. Over the course of potentially thousands of interactions, a Machine Learning algorithm will adapt itself to the user, recommending content that best elicits the desired behavior from the user whether it is purchasing more products, spending more time on the platform, or using a share feature. A basic Machine Learning algorithm may require explicit user-interaction to learn, like 'liking' a song or a post, and a sophisticated Machine Learning algorithm might generate its own criteria for learning based on correlations to desired behaviors that it calculates as interactions occur. Notice that the algorithm has been refined by externalizing training work to the user rather than relying on a developer.

Search engines are another common use case for Machine Learning. In the old days (actually, a mere decade ago), search results were the product of an enormous code base with hard-coded rules. If a user searched for a term, there was a pre-determined output for that term. 'Google-fu' became the slang for those who had a knack for using search terms that generated the correct result every time, rather than a heap of generic results that might contain the right answer. Since then, search engines have improved significantly without us stopping to notice. Machine Learning has allowed search results to be adaptive, predicting the user-desired search result based on user behaviour and other context cues like location, time of day, previous searches, and trending search terms or news items. Machine Learning has greatly increased the predictive ability of search engines while also being a lightweight approach to coding by removing the need for a pre-defined response to every possible input.

Importantly, for many years, machine learning required well-structured and clean data to be used. Over the last decade, so-called Deep Learning methods have emerged which are not as limited by those requirements.

Layering complexity

Deep Learning is a subset of Machine Learning which came to prominence in the 2010s and is well-suited to work with unstructured data (images, sound, and other data not in a typical tabular form). Deep Learning works with extremely large data sets, often large and complex enough to be classified as 'big data' by using 'neural networks' to understand complex pattens in data by abstracting relationships and gradually learning them over time.

Deep Learning neural networks are 'deep' because they stack multiple 'hidden layers' between the input and output layers.[1] Deep Learning has unlocked many more interesting use cases for AI, from speech recognition and language processing to image and video processing.

As with other kinds of machine learning, Deep Learning can be divided into two large groups: discriminative and generative. Discriminative AI seeks to divide data into categories and is fundamentally about finding how values of x will predict y. (A common example of this is logistic regression.) Generative AI, on the other hand, is about learning the probability distribution that generated a target value. (An example of this is Bayes' Theorem.) Deep Learning is especially relevant for generative tasks because it allows such distributions to be learned through a multitude of nonlinear abstractions. We will explore this further in In Silico, Part III.

There are two major kinds of neural networks, among others:

- **Convolutional Neural Networks (CNNs)** use 'convolution' to extract features from data. A simple way to think about convolutions is to think about how sound reverberates in a room. If you have ever been in a place so echoing that your own voice is a distraction, then you have experienced how one sound is changed by another sound – convolution helps us to describe situations like that. Convolutions are used to help understand how one feature is changed by another and in CNNs they help with tasks like object detection in images, with each layer of a CNN allowing for more complexity and detail abstraction.

  CNNs enable some vision and image technologies we use almost daily, like image upscaling, photo enhancement, noise cancellation, video stabilization, and facial recognition. If you have a phone, use a computer, watch 1080p content on a 4k screen, stream videos, or play games, you likely engage with these technologies without ever noticing. These are underlying technologies which have increasingly been powered by AI systems of one kind or another and which create enormous-but-unseen quality of life and ease-of-use improvements.

1 Just as not all Machine Learning is Deep Learning, not all Neural Networks are deep.

- **Recurrent Neural Networks (RNNs)** are neural networks best used for time series or similarly ordered data. RNNs use recurrence relations to understand how dependent relationships change outputs over time. A canonical example is a Fibonacci Sequence - where the current number depends on the previous two - though many other recurrence relationships with more complicated equations exist.

  Recurrence relationships are used in macroeconomics to model the economy, typically to estimate current unknown values using lagged variables like CPI. In RNNs the same approach can be taken to predicting text and speech. Language follows rules that inherently shape what would we expect to follow on from previous words in a sentence. As a sentence approaches its end, the set of likely words that will appear decreases. This is equally true for time series data, where our certainty about an unknown value at some time in the future can increase as we move closer to that point in time. RNNs are designed for these use cases.

Already all around us: AI in today's economy

AI powers the world around us in thousands of ways—through content recommendation, content monitoring, advertising, signal generation, noise filtering, spam filters, fraud prevention, cybersecurity threat detection, voice assistants, facial recognition, maps routing, ridesharing optimization, computer opponents in games, and more—and this quiet revolution of advancements has been taking place quietly before our eyes for many years.

Since 2017, the share of companies using AI in at least one business unit or function has risen from 20% to 50%, hovering between 50% and 56% in 2020 – 2022, according to a McKinsey study[2]. The average number of AI capabilities deployed has also risen over that period, from 1.9 to 3.8. Interestingly, while 50% of companies used AI in both 2020 and 2022, the average number of capabilities deployed increased from 3.1 to 3.8, suggesting that even if adoption did not become more widespread, adoption did increase.

With November 2022's launch of ChatGPT, AI feels like an 'overnight success' – but it has been decades in the making. Progress in AI has been steady for many years, enabled by three continuous advancements: new methods, the increasing availability of data, and a consistent and rapid growth of computing power. The world has seen an exponential increase in data production and capture, computing power, and model sophistication. We see no sign of this trend stopping, so we see no inherent reason for AI progress to stall,[3] whether generative or more traditional. In our next edition, we will explore Generative AI, and why there is so much buzz around it.

Figure 2A: Computing power needed to train the best-in-class AI models has increased 300,000x in the past decade.

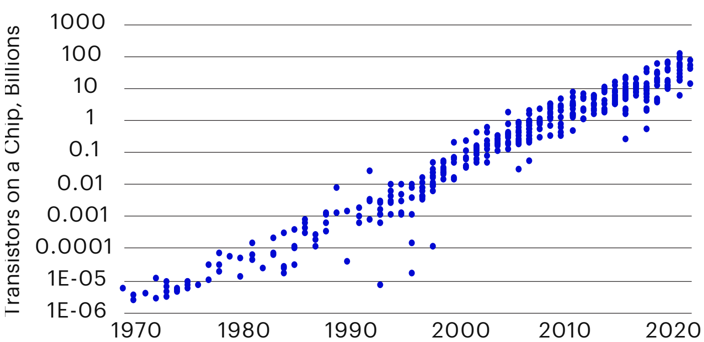

Source: AI and Compute, OpenAI – 16 May 2018, and OurWorldInData as of 19 May 2023. Note: Vertical axis uses exponential scale.

Figure 2B: Models have grown increasingly sophisticated – ChatGPT-4 may have 100 trillion parameters.

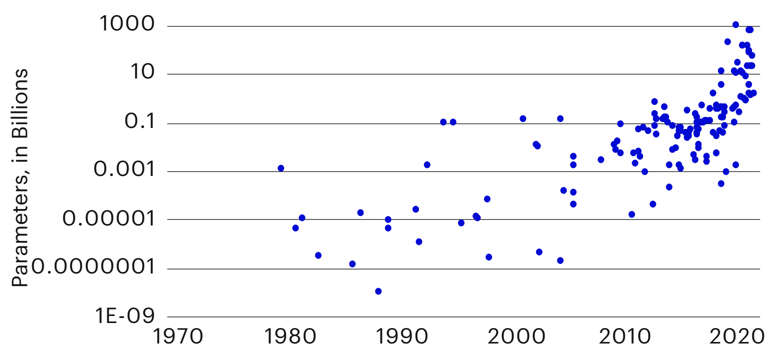

Source: Sevilla, Villalobos, Cerón, Burtell, Heim, Nanjajjar, Ho, Besiroglu, Hobbhahn, Denain, and Dudney (2023) Parameter, Compute, and Data Trends in Machine Learning via OurWorldInData. Chart shows selected models and estimated parameter counts from 1970 to July 2022. Note: Vertical axis uses exponential scale.

2 Source: McKinsey State of AI in 2022 and a half decade in review.

3 Although, it is possible that regulatory responses will apply the brakes if it is deemed necessary for responsible innovation.

AI at Invesco

Invesco's technology team, Strategy Innovation and Planning, helps enable AI-driven solutions across functions by reviewing problem statements and requests for expertise and working across departments to build capabilities.

AI sentiment analysis

One of our most recent innovations has been the development of a sentiment scoring system that uses natural language processing and machine learning to extract sentiment signals from quarterly earnings calls. Because stock prices are often volatile around earnings calls, we have identified them as a potential source of alpha.

This process uses transcripts from FactSet®, which identifies each paragraph by speaker and segment (e.g., Presenter, Q&A). The dataset comprises > 200,000 transcripts from over 10,000 companies since 2008. Our system calculates sentiment scores based on pre-defined lists of positive and negative words tailored specifically to a financial context. For example, the term 'liabilities' might be perceived as negative in a general discussion but is considered a neutral standard term in finance. To generate a sentiment score, we extract the total number of positive words in the text and the total number of negative words, and calculate a weighted score.

Earnings calls typically consist of three distinct sections: CEO's prepared remarks, CFO's prepared remarks and analyst questions and answers (Q&A). We score these sections separately. In addition, we calculate sector-normalized scores, comparing a company to its universe and sector. We also calculate a score comparing the current call to the previous one. For example, if the raw sentiment score was 34 in the previous quarter and is now 16, the 'change score' would be -18. We calculate this as 'surprises', and changes are likely to move stock prices considerably. If the company has long been doing well, much of that may already be reflected in the stock price. However, new information that materially differs from that of the previous quarter may shift sentiment – and stock prices – dramatically.

AI Aspect-based sentiment analysis

Recently, the technology team introduced their latest product, the Hotel Review Model, designed specifically for our European Hotel Strategy team. Leveraging advanced machine learning and natural language processing techniques, the Invesco Hotel Review Model undertakes a comprehensive analysis of online reviews contributed by hotel customers, encompassing comments, feedback, and ratings. To pull comprehensive insights into customer sentiment from the reviews, the model surpasses conventional sentiment analysis methods by employing a technique known as Aspect-Based Sentiment Analysis. In contrast to conventional sentiment analysis, which treats the text as a whole and assigns a singular sentiment label or score, Aspect-Based Sentiment Analysis deconstructs the reviews into distinct features and effectively captures not only the sentiment associated with each feature but also the underlying reasons behind such sentiment. This approach facilitates the extraction of actionable insights from the reviews, empowering the team with a more nuanced understanding of customer sentiment.

The model provides Invesco with detailed insights into common complaints and positive feedback, such as noise or air conditioning complaints, hotel cleanliness, and customer service. The tool also allows the team to monitor customer sentiment and track changes in feedback over time, providing valuable information for investment decisions.

Investment risks

The value of investments and any income will fluctuate (this may partly be the result of exchange rate fluctuations) and investors may not get back the full amount invested.

Important information

This document is intended only for Professional Investors in Hong Kong, for Institutional Investors and/or Accredited Investors in Singapore, for certain specific sovereign wealth funds and/or Qualified Domestic Institutional Investors approved by local regulators only in the People’s Republic of China, for certain specific Qualified Institutions and/or Sophisticated Investors only in Taiwan, for Qualified Professional Investors in Korea, for certain specific institutional investors in Brunei, for Qualified Institutional Investors and/or certain specific institutional investors in Thailand, for certain specific institutional investors in Malaysia upon request, for certain specific institutional investors in Indonesia and for qualified buyers in Philippines for informational purposes only. This document is not an offering of a financial product and should not be distributed to retail clients who are resident in jurisdiction where its distribution is not authorized or is unlawful. Circulation, disclosure, or dissemination of all or any part of this document to any unauthorized person is prohibited.

This document may contain statements that are not purely historical in nature but are "forward-looking statements," which are based on certain assumptions of future events. Forward-looking statements are based on information available on the date hereof, and Invesco does not assume any duty to update any forward-looking statement. Actual events may differ from those assumed. There can be no assurance that forward-looking statements, including any projected returns, will materialize or that actual market conditions and/or performance results will not be materially different or worse than those presented.

All material presented is compiled from sources believed to be reliable and current, but accuracy cannot be guaranteed. Investment involves risk. Please review all financial material carefully before investing. The opinions expressed are based on current market conditions and are subject to change without notice. These opinions may differ from those of other Invesco investment professionals.

The distribution and offering of this document in certain jurisdictions may be restricted by law. Persons into whose possession this marketing material may come are required to inform themselves about and to comply with any relevant restrictions. This does not constitute an offer or solicitation by anyone in any jurisdiction in which such an offer is not authorised or to any person to whom it is unlawful to make such an offer or solicitation.

This document is issued in the following countries:

- in Hong Kong by Invesco Hong Kong Limited景順投資管理有限公司, 45/F, Jardine House, 1 Connaught Place, Central, Hong Kong. This document has not been reviewed by the Securities and Futures Commission.
- in Singapore by Invesco Asset Management Singapore Ltd, 9 Raffles Place, #18-01 Republic Plaza, Singapore 048619.
- in Taiwan by Invesco Taiwan Limited, 22F, No.1, Songzhi Road, Taipei 11047, Taiwan (0800-045-066). **Invesco Taiwan Limited is operated and managed independently.**