Invesco

In Silico, Part I

Understanding AI's "Hello, World!" Moment

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 first in a series in which we explore AI's explosion of interest in 2023, contextualising this moment in a brief historical overview and outlining opportunities. Subsequent articles will consider the economic, financial, and social impact of generative AI, robotics, and improvements to AI that is already embedded in our world. 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

Artificial Intelligence (AI) fundamentally alters our day-to-day interactions with the world and each other, as private citizens, as scientists and researchers, and as working professionals. AI is increasingly being embedded in our economic systems, at every point in the supply chain from R&D to sales, from procurement to marketing to after-sale analysis. While the term has been elevated to something of a buzzword or a trend to cash in on – generating hype and fear in equal measure – it's important to remember that AI already has more than a decade of real-world applications. We do not believe we are at the very beginning of an AI-driven transformation — rather, we believe we are at a point where we will see accelerating change.

Reflecting this, March 2023 was an astonishing month for generative AI. The month saw a host of new technology announcements and releases, including: OpenAI's ChatGPT4 and Microsoft's plans to integrate GPT4 functionality in its products and services, Google's early release of the Bard AI, Med-PaLM2 and the PaLM API, Baidu announced its AI, Midjourney 5 was launched, and many more companies and independent projects unveiled targets for new AI goals. Secondly, we have also seen an enormous uptick in interest surrounding AI projects, ranging from excitement to scepticism to ethical concerns, culminating most recently in an open letter to pause AI research[1]. If there is one thing that everyone can agree on, it's that AI has made a splash (Figure 1).

Figure 1: Google Trends Search Interest (31 Dec. 2018 - 1 Apr. 2023)
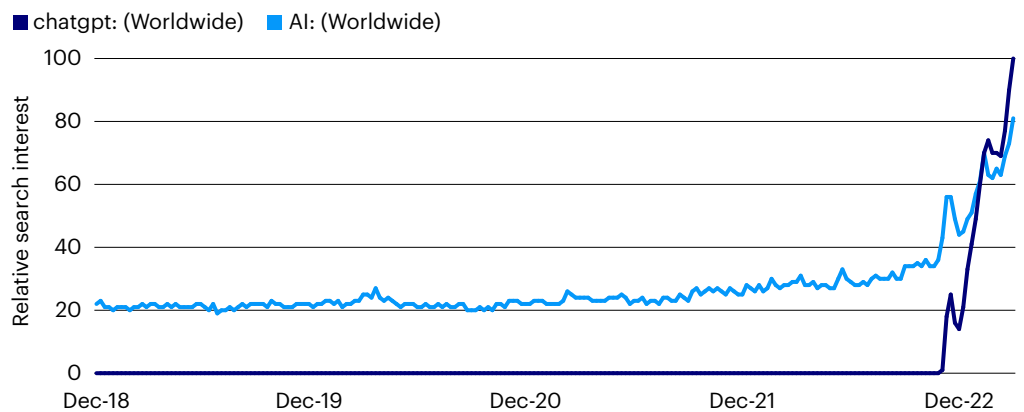
Source: Google Trends and Invesco, as at 1 April 2023. Relative search interest is normalized according to Google Trends normalization process. For more information see Google News Lab's 'What is Google Trends data - and what does it mean?' (2016)

1. Pause Giant AI Experiments: An Open Letter - Future of Life Institute

We believe that AI is among the most promising drivers of transformation, potentially bringing change to every industry and sector, with relevance for all manner of business models. AI may also herald innovations in efficiency and automation, especially for labour intensive and monotonous tasks including some kinds of physical labor, as well as white collar and creative work such as transcription, data analysis, and copywriting.

In this series, we look to explore the consequences of artificial intelligence, its potential economic and financial impacts, and its potential as a driver of change. We see three significant sub-themes for AI technology, each bringing challenges and opportunities:

1. **Generative AI:** the latest wave of creative AIs, including chatbots like ChatGPT, image generators like Midjourney, and emerging generative AI tools which aid in creating more realistic social media filters and deepfakes.

2. **A continuation of the quiet revolution:** AI appears to be seeing Moore's Law-like improvements in the background technologies that power our lives, from image stabilisation and photo processing to video upscaling, to content recommendation, and biometric identification.

3. **Rise of Robotics:** Machines have a longer history than contemporary AI and their marriage has been fruitful already. We see full-fledged autonomous robots, especially self-driving drones and vehicles as an inevitable and game-changing innovation.

AI, Then and Now

Conceptually, the idea of an artificial intelligence (AI)—a living, sentient being with an intellect like ours—stretches back into classical antiquity. Technologically, the history of AI begins in two places: first, with early automation, and then most recently somewhere between the 1930s and 50s, with a confluence of advancements in early neuroscience, mathematics, and early computation.

Since the 1950s—the same period that gave birth to the famous "Turing test"—the number of "notable" AI-related publications produced each year has grown.[2] In 1951, the first working AI programs were developed that could play checkers and chess. In the decades that followed, progress in the underlying mathematical work that underpins today's neural networks and other AI technologies, plus the invention of modern computers and exponentially increasing computing power and data storage capabilities enabled a series of fits and starts in AI development. The late 1980s brought the earliest forms of self-driving cars, which enabled the "No Hands Across America" project of 1995, where a semi-autonomous car drove from coast-to-coast in the continental US. More recently, the last decade of AI advancement has seen AI not only engage in primitive replication of human tasks but also exceed human capabilities in some domains.[3]

Since the 2010s, the locus of research has shifted from academia to something more collaborative, with industry taking the lead in recent years (Figure 2).

Figure 2: Affiliation of research teams creating notable artificial intelligence systems
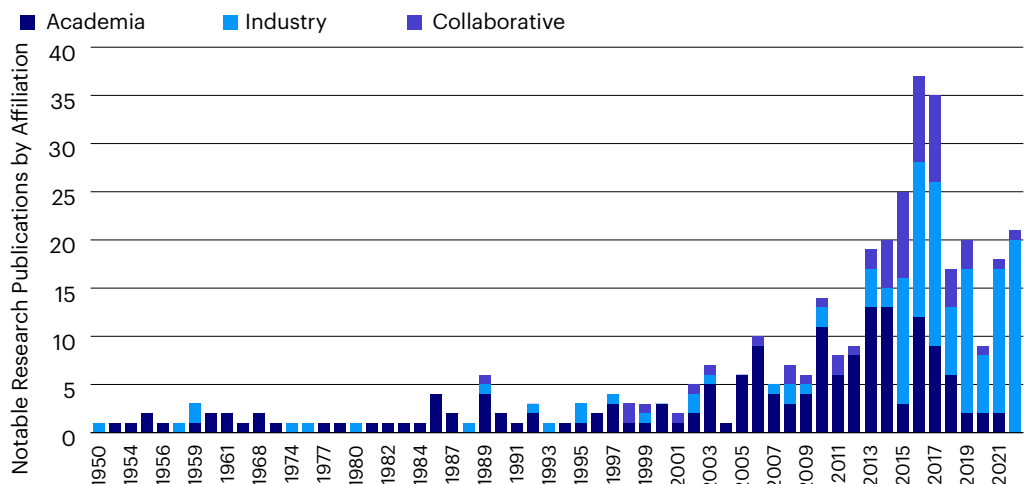
Source: Invesco and Our World in Data citing Sevilla, Villalobos, Cerón, Burtell, Heim, Nanjajjar, Ho, Besiroglu, Hobbhahn, Denain, and Dudney (2023). 'Parameter, Compute, and Data Trends in Machine Learning'

2. Invesco and Our World in Data citing Sevilla, Villalobos, Cerón, Burtell, Heim, Nanjajjar, Ho, Besiroglu, Hobbhahn, Denain, and Dudney (2023). Parameter, Compute, and Data Trends in Machine Learning. Systems are defined as "notable" by the original authors based on several criteria, such as advancing the state of the art or being of historical importance. Inclusion criteria includes both quantitative and qualitative metrics.

3. Chess, Go, and other board games are 'solved' problems for AI. While AI is also posing challenges across a range of complicated games, like Diplomacy.

An Evolution of Statistics

Contrary to pop-culture conceptions of AI as a sentient being, contemporary AI research is rooted in mathematics and statistics as a means of digitally representing and analysing our world, using large bodies of data to iteratively improve performance relative to some set of human-defined goals. "Neural networks" are a prominent staple of this statistical approach to understanding our world. While there are a range of tools and approaches to AI, taking a moment to understand how a neural network could be used for image recognition helps to demystify how advanced AI technologies operate under the hood.

Imagine we want software to take an image and determine whether the image is a black and white text character, such as the number '2'. You can divide that image into its constituent pixels and feed each pixel into an input node, which will output a value like 1 if the pixel is white and 0 if the pixel is black. Then, you can feed that result into a second layer containing fewer nodes, creating from it a weighted average. For example: 10 layer-1 nodes feed into each layer-2 node, averaging the value of 10 of the original pixels to give a value to a 10-pixel area of the original image. If that 10-pixel area is all white, the value will be 1 and if it is all black it will be 0. If instead it contains a mixture of black and white—because part of the number 2 is in that area—then we will receive a number between 1 and 0 depending on the weighting of the average value of the inputs.

At the end of this process, the computer takes those weighted values and outputs a guess about what the image is. That guess is compared to labels for the data that indicate its true value. If the guess was correct, the weights are kept and if the guess is wrong the weights are changed until the guess is right. This process can then be repeated for other numbers to find weightings that allow the machine to recognise the numbers 2, 3, and so on. At scale, this trial-and-error process, and processes like it, underpin much of AI and approaches to machine learning.

Today, we have simultaneously more computational power available than ever before and far more data to process and analyze, too. These advancements enable machine learning engineers to build increasingly sophisticated models using billions of different parameters (rather than the simplistic set of pixels of our earlier example). The fundamentals of this approach are more-or-less the same: trial-and-error. However, some of the most impactful recent innovations have focused on reducing the need for human supervision during the trial-and-error process, including dispensing with the need for humans to manually label every piece of data.

Generative AI has been the subject of much of the recent AI hype, comprising a suite of products which can generate new content from old content, accepting a wide array of inputs and producing a wide array of outputs. Tools based on generative AI may be implemented into workflows, potentially improving productivity.* This is the hope underlying recent implementations of OpenAI's Chat-GPT4, where companies such as Microsoft have sought to include generative AI capabilities in existing products. Indeed, the advent and implementation of generative AI could precede meaningful uplifts in productivity that were last seen in the early days of computers.

Importantly, generative AI has its provenance in the statistical models that drive machine learning. Such models are not sentient or conscious, they do not have agency or set their own goals, and they do not think. Today's generative AI are more-or-less clever parrots and underlying them are sophisticated statistical processes with an uncanny-but-entirely-inanimate ability to produce human-like responses to human-like inputs (after having been trained on human-produced data). Fundamentally, today's generative AI systems systems are just computers of a kind.

Assessing Opportunities

AI development is a capital-intensive endeavour. In well-established domains, the barriers to entry are high and this is a natural by-product of the AI development pipeline, sketched as follows:

1. Suitable data are collected into a dataset, cleaned, and labelled. A high-quality dataset may be many terabytes in size and require many trillions of computations to train.

2. Models are trained on data, historically with human-supervised approaches although unsupervised approaches are becoming increasingly sophisticated.

3. Models are then distributed either as in-house projects, sold, or offered as AI as a Service (e.g.: BloombergGPT)[4], or they are distributed as open-source projects which may or may not be suitable for commercial applications.

Todays' largest and most well-known players have advantages either because they control access to the highest quality datasets or because they have sufficient capital and assets to train, deploy, and support computationally intensive AI. Some industry leaders go so far as to make custom hardware and firmware solutions for their AI products, including custom chips. A single, top of the line commercially accessible GPU would take hundreds or thousands of years to train a model like GPT4, while many hundreds of highly specialized, industry grade processing units can make new models accessible in just a couple of years or less. Although AI can be cheap to use when built, training can be very capital intensive. Having access to lots of data or substantial capital represents a defensible moat long-before considering the innate advantages that come with a big name's brand recognition or pre-established routes to market.

The very cutting edge of AI research is defined by those large players who can deploy sufficient capital to do more with more, scrappy innovators who have worked out how to do more with less (often through more focused product offerings), and those adventurous few who seek out new niches or use-cases altogether typically as part of collaborations between academia and industry.

An emerging domain, like synthetic biology (SynBio), is one example where AI can make a difference and where established players have little to no existing footprint. Indeed, AI can be applied to already-existing '-omics' data sets and is already deployed in research including de novo antibody generation[5], biomarker identification, protein folding, describing novel bioassays, and more. Many of these advancements are blended, with academic institutions using open-access tools from the likes of DeepMind, Meta, and OpenAI. Collaborative projects, like BLOOM, also exist. Meanwhile, some institutions like Stanford, are releasing their own AI strictly for academic research (prohibiting commercial use)[6].

New entrants to AI may seek opportunities elsewhere, beyond the current purview of major technology firms, established players, and Tier-1 research institutions. The last decade has seen a significant increase in the number of AI start-ups, including an increase in new entrants each year and an increase in funding[7] (Figure 3).

Figure 3: Number of Private Investment Events and Number of Newly Funded AI Companies
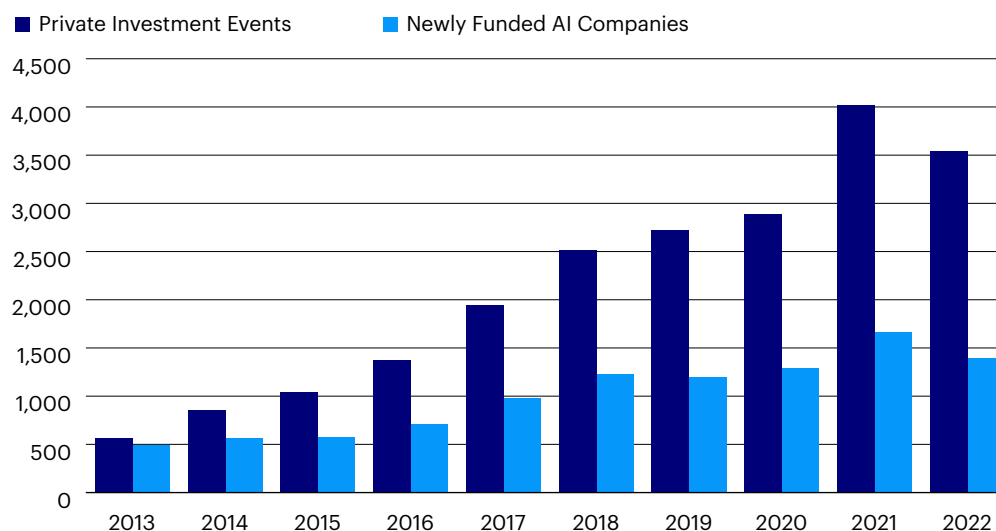
Source: Invesco and Stanford AI Index Report 2023 Data Repository citing Netbase Quid 2022

5. Unlocking de novo antibody design with generative artificial intelligence | bioRxiv
6. Stanford CRFM
7. Stanford 2023 AI Index Report citing NetBase Quid 2022 - AI Index Report 2023 – Artificial Intelligence Index (stanford.edu)

Figure 4: Total Global Investment in AI (in Billions of US Dollars)
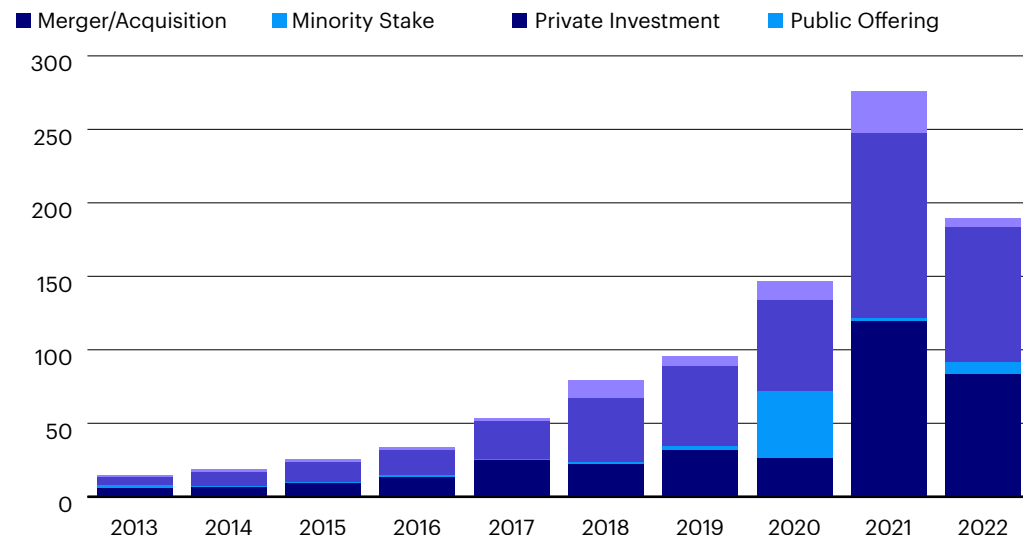
■ Merger/Acquisition   ■ Minority Stake   ■ Private Investment   ■ Public Offering

Source: Invesco and Stanford AI Index Report 2023 Data Repository citing Netbase Quid 2022

Cautious investors may favour a 'picks and shovels' approach, leaning towards the enablers of AI rather than the architects. If so, it is worth considering the underlying infrastructure that makes research possible: hardware and training data. Suppliers of raw materials or machinery for semiconductor manufacturing, semiconductor manufacturers themselves, and those who build the surrounding hardware and software architecture for computer processors, graphics processing units, and other specialized processors (such as tensor processing units, or TPUs) are likely to benefit if they position themselves for the future of AI. Meanwhile, many of the traditional players in Big Tech have become something like Data Giants, hoarding petabytes of valuable customer information and consumer and user behaviour which might be repackageable as training data (legislative and regulatory hurdles notwithstanding).

AI may also be packaged in a 'AI as a Service' business model, with potentially wide-reaching implications for improving existing ways of doing business and, in the process, likely bolstering productivity wherever it is sensibly deployed inside organisations.

While we are excited for the future of AI-driven innovations, we acknowledge that, at least currently, the pace of innovation currently outstrips legislative and regulatory capabilities. While pockets of academia and independent think tanks have called for more ethical oversight, including a precautionary approach to innovation, AI has largely been able to innovate free from onerous constraints. Although current legislation, like data protection and privacy laws, often do cover use cases for AI, it seems likely that AI will become an increasingly regulated industry, which may present some bumps in the road for the future of AI development.

Looking Ahead

In 2023, AI has been framed as the innovation to change our lives. A future full of AI is viewed with a mixture of optimism and pessimism, ranging from hopes of a path to technological utopia to fears of human obsolescence. In this series, we hope to step beyond these binary viewpoints by exploring the impacts of AI with an emphasis on economic opportunities and risks. We are confident that AI has already left an indelible mark on recent technological progress, with more to come. In future articles in this series, we will dive deeper into the explosion of interest in AI and look back on the quiet revolution that has long been underway, powering so much of the world around us already.

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 issued in: