An Introduction to Artificial Intelligence

September 2023

ChatGPT, the first widely accessible Generative Artificial Intelligence model, met with an avalanche of excitement and headlines when it was released in November 2022. It is believed to be the fastest growing consumer application ever, reaching 100 million users within just two months. Since then, Generative AI has continued to make waves, from winning Art Awards to driving contentious foreign policy. Shows such as Red Dwarf and 2001: A Space Odyssey conjure an image of Artificial Intelligence as a sort of omniscient being trapped inside a computer, capable of thinking and conversing like a human. Ultimately, that is the goal of Generative AI models: to replicate and apply the processes behind thought to their, theoretically, unlimited knowledge base.

Generative AI software is based on a 'Foundation Model'. Rather than being created to tackle one specific task, a Foundation Model is trained to recognise patterns in huge quantities of data, such as textbooks or transcripts. This training can take place without any human interaction, massively improving the rate at which information can be processed. To mimic our brains' networks of billions of neurons, these AI learn the probability of one chunk of data following another, such as the probability of the word 'Intelligence' following 'Artificial'. This probability is a parameter. The quantity of parameters has grown from 117 million in the first ChatGPT model to 170 trillion in the latest. Every time it answers a question, it selects each piece of the answer based on those probabilities, giving it the ability to produce different results every time. This is not just limited to text; this can be applied to audio, numbers, and pictures, to name just a few applications. This breadth begins to explain the excitement around Generative AI. It is designed to be a general-purpose technology, to underly almost every computing process.

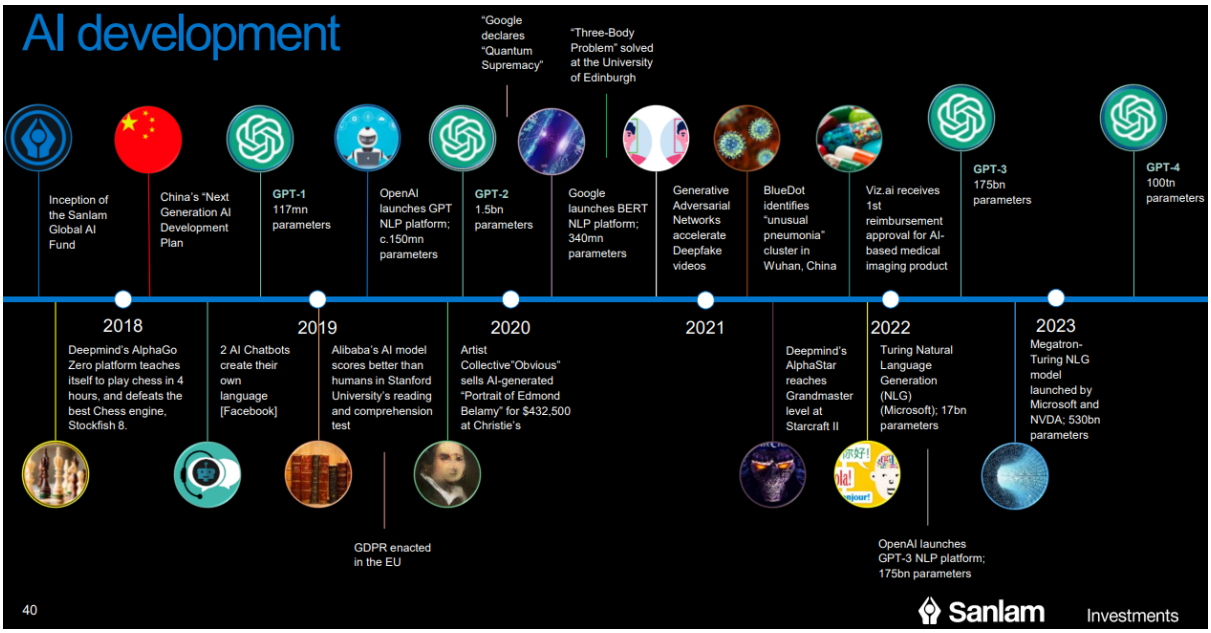

Source: Sanlam Investments

To have such broad application, AI models must have access to broad knowledge. Given the vast lakes of data involved, AI models require huge amounts of computational power. They use the most advanced computer chips available, currently designed by Nvidia and produced by Taiwan Semi-Conductor Company ('TSMC'). Some AI can run on less powerful chips but will be limited in their capabilities. The current monopoly Nvidia holds (over 80% market share) underlies the recent rise in their share price since AI hit headlines. AI software companies have been desperate to build their stock of chips to expand their systems' power and capacity. The chart below shows the scale of this power increase over time.

Computing power needed to train the best-in-class AI models has increased 300,000 times in the past decade.

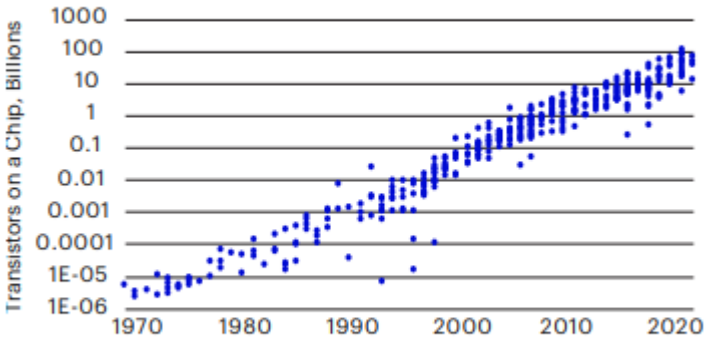

Source: AI and Compute, OpenAI – 16 May 2018, and OurWorldInData as of 19 May 2023. Note: Vertical axis uses exponential scale.

Design is not the only challenge in the creation of computer chips that are capable of processing AI. Their complexity requires engineering precision, for which the leading fabricator is TSMC. Oxford Economics forecast semi-conductors to be one of the fastest growing industries over the next 15 years. In fact, TSMC have recently announced their intention to open another fabrication plant to meet the growing demand for their most powerful chips. The chart below illustrates the economic potential for chip manufacturers, with revenue nearly doubling by 2026.

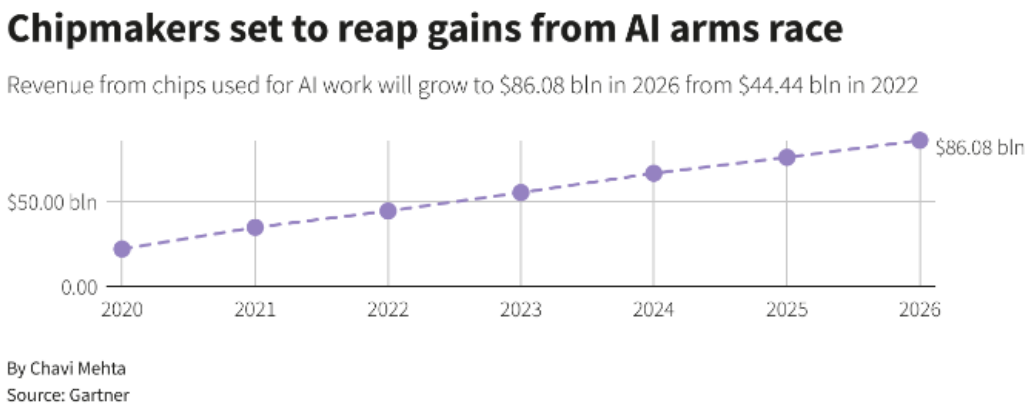

Source: Reuters

ASML is a similarly monopolistic company in the supply of lithography equipment (machines used for the fabrication of computer chips) and is positioned to grow along with the industry. Even property developers can benefit by leasing warehouses as data centres to house legions of computing equipment. Given the scale of the total addressable market for AI, the entire supply chain will benefit. But investors also need to consider who will capitalise on its application.

There are thought to be two main avenues through which Artificial Intelligence can benefit businesses: direct revenue from AI products and services, and productivity enhancements through complementing labour. The direct revenue from AI will be a fraction of the overall benefit to society and is likely to be concentrated in the most advanced companies at each stage of the supply chain. Alongside the names mentioned already, software giants such as Microsoft and Alphabet will share in this direct revenue. For example, Microsoft has announced the pricing of their Office 365 AI program, Copilot, at $30 per user, per month.

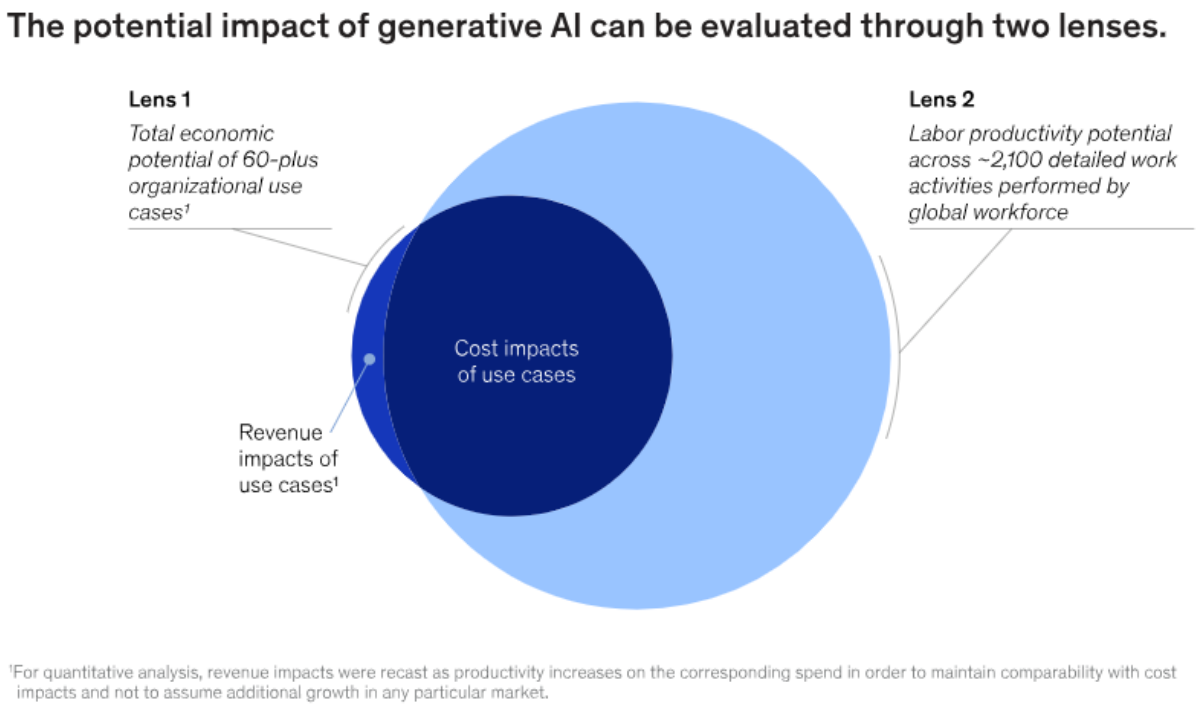

Source: McKinsey & Company

So, AI has much to offer, but what should we all be most excited about? The answer is simply the substantial productivity gains. While it is impossible to forecast the total impact, most economists conservatively expect somewhere between a 0.5% and 1% increase in economic growth. Given the weak productivity growth over the last decade, shown below, this is a substantial boost. Factors such as the rate of adoption and the regulatory environment will heavily affect AI's impact, and those impacts will not be spread evenly across the economy but concentrated in certain industries. While some jobs may be aided by AI, others may become obsolete. Computer programming is thought to be among the most influenced. In a trial set up by McKinsey, they found developers completed simple tasks in half the time when given AI assistance, although more complex projects saw less improvement.

Source: Oxford Economics

Perhaps the most profound application for AI is in healthcare. Given AI's strength in pattern recognition, diagnoses can be made earlier, faster and with greater accuracy than when made by humans. By pooling data from hospitals, AI can use knowledge from the most experienced doctors across the world. It can even help to better model how DNA affects our biology or predict what future strains of flu will look like.

To harness this potential, companies need internal data. The unverifiable nature of public data, such as that on social media or the internet, necessitates a database of trustworthy data. AI learn patterns, not logic or rules. If they are told that 2+2=3, that is what they will answer. To be able to trust the output of AI, we need to trust the input. Clearly, this hands an advantage to larger companies or simply companies whose business is data. It will be much easier for AI to spot patterns in consumer behaviour in Tesco than in your local corner shop.

It is our task to consider how best to monetise this trend within the context of a multi-asset portfolio. While there are specific AI equity funds available, these are very specialist strategies. Focusing on a single theme can lead to excessive volatility as every stock is correlated. This means that if the market has a loss of faith in AI because of governance concerns or semi-conductor providers and engineers stumble in their journey to deliver the tools for AI to flourish, the value of such a fund with an AI focus can fall dramatically. Our allocations to funds investing in the US and internationally naturally provide exposure to the leading AI beneficiaries. In fact, given the concentration of the top seven stocks in US and Global Indices, many index-linked funds will also give access to the theme. The managers of these underlying funds are well placed to recognise and exploit any prevalent investment themes, of which AI generally tops the list. Additionally, we have exposure to those

companies further down the supply chain, such as TSMC, within our Asian and Emerging Market equities. Ensuring that our clients capture the benefit of the future productivity gains will be more nuanced. We have seen companies within the supply chain rapidly rise in value, but the long-term economic growth will be gradual and far-reaching.

To give everyone a taste of what Generative AI is capable of, I asked ChatGPT to write our conclusion. It certainly doesn't lack self-confidence, but the output is impressive none-the-less.

In summary, the journey of generative AI from its groundbreaking debut to its potential across myriad industries signifies a transformative chapter in our technological narrative. The impact on economics, technology, and healthcare is substantial, promising unparalleled advancements while demanding a nuanced consideration of ethical and societal implications. As we navigate this path, embracing AI's power to augment human ingenuity and innovation, we embark on a journey that promises to reshape the fabric of our existence in ways both profound and unprecedented.

In a less eloquent nutshell, I'd say that we broadly welcome the AI revolution and believe that it will benefit us all.

Ollie Rubinstein, Investment Analyst, City Asset Management

Ollie is an Investment Analyst, responsible for meeting and researching fund managers, assisting all teams with data queries and researching the macro-economic side of investing. Having recently graduated from Bath University, his interest in numbers and background in economics gives him an extra level of insight. Keen to continue his professional development, Ollie is currently studying for his IMC and CFA exams

Jargon Buster

For further clarification of any of the terms used, please see the glossary below:

Artificial Intelligence (AI): The simulation of human intelligence processes by machines, especially computer systems, involving tasks such as learning, reasoning, problem-solving, and decision-making.

Artificial Intelligence Products and Services: Offerings created using AI technology, providing solutions, insights, or automation.

ChatGPT: The first widely accessible Generative Artificial Intelligence model, designed for generating human-like text responses and interactions.

Data Pooling: The aggregation of data from multiple sources to create a larger dataset for analysis.

Data Processing: The manipulation and organization of data through various methods to produce desired outcomes or insights.

Equity Funds: Investment funds that primarily allocate capital to stocks and shares of companies.

Foundation Model: A base model trained to recognize patterns in vast amounts of data, serving as the foundation for more specialized AI models.

Generative AI: Artificial Intelligence that generates new content, such as text, images, or audio, based on patterns learned from large datasets.

Index Tracker: An investment fund designed to replicate the performance of a specific market index.

Lithography Equipment: Machines used in the fabrication of computer chips to imprint intricate patterns onto semiconductor wafers.

Multi-Asset Portfolio: A collection of different types of investments, such as stocks, bonds, and real estate, designed to manage risk and achieve returns.

Neural Network: A computer system modelled after the human brain's interconnected network of neurons, used in AI for tasks like pattern recognition and machine learning.

Parameter: A value used to control the behaviour or output of a system, often learned from data in AI models.

Pattern Recognition: The ability of AI to identify and interpret recurring patterns or trends in data.

Productivity Gains: Increases in efficiency, output, and resource utilization, leading to improved economic performance.

Semiconductor: A material with electrical conductivity between a conductor and an insulator, used in computer chips and electronic devices.

Volatility: The degree of variation in the price of a financial instrument, indicating its risk level.