James Sanders – Portfolio Manager
January 2022

# On the Horizon #3 – Artificial Intelligence

Artificial intelligence has fascinated humankind since the ancient Greeks – in around 700 B.C. Hesiod explored the idea of robots in the story of Talos, and of an artificial woman in that of Pandora, both made by Hephaestus who, according to Homer, also made automated servants out of gold. More recently, and more impactfully, George Boole set out to “investigate the fundamental laws of those operations of the mind by which reasoning is performed” in the 1850s, creating Boolean algebra, the foundation of all modern programming languages.

While many of today’s approaches to AI were conceived decades ago (neural networks were simulated as far back as the 1950s), it took a few decades of Moore’s Law before the field of AI could advance beyond the theoretical, and even then there were bumps along the way. Setbacks in the 70s and 80s led to ‘AI Winters’, as AI’s failure to live up to its hype led to capital fleeing the sector.

The data creation associated with the rise of social media and Web 2.0 in the 2000s, alongside increasingly powerful computational resources, acted as major accelerants to the commercialisation of AI. Perhaps the most ubiquitous application is the recommender algorithm - a tool used by all social media and digital content giants to personalise our online experience and command our attention. This is paired with natural language processing (NLP) models in web search to understand what we are looking for, or might want to see, and instantly deliver it our screens.

## Contents




11 Albemarle Street, London W1S 4HH Tel: 020 3170 7420 Fax: 020 3170 7426. E-mail: info@greenash-partners.com Web: www.greenash-partners.com

## Recommender Systems

Recommender systems underlie nearly everything we do online. They control what products are shown to us when we shop, what appears at the top of our social media news feeds, and what advertising we see. They are a convenience, but also a necessity - in the Information Age there is so much content on the internet that it would be impossible to navigate without them. In essence, recommender systems come in two flavours:

**Content-based filtering** – suggest similar content or products identified based on context or labelling

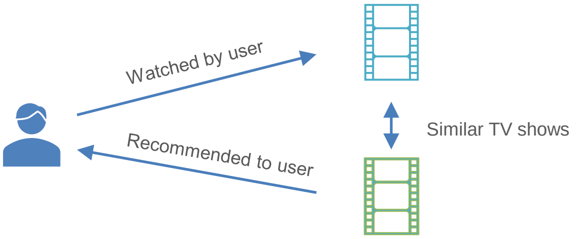


**Collaborative filtering** – utilises historical data on users and/or items to suggest new products or content

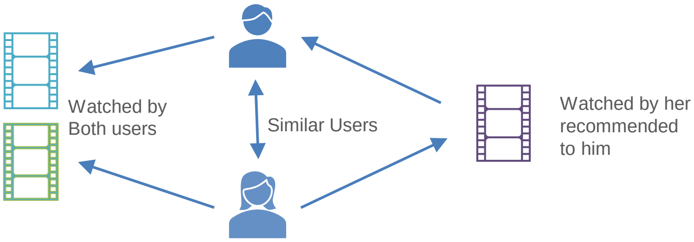


But the algorithms in use today are extremely complex, using machine learning techniques like reinforcement learning and graph-based reasoning to identify connections between large numbers of variables, and as a result they have become harder to understand and monitor.

### Graph-based Reasoning Example

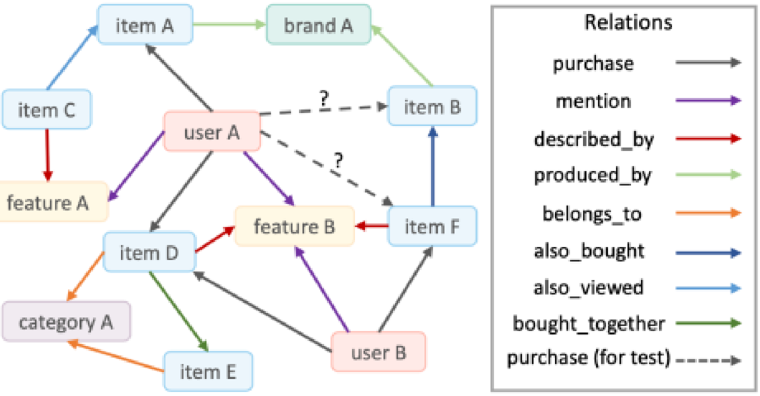


To illustrate this, here is an example of graph-based reasoning, which uses heterogeneous datasets like users, items, content, tags, and reviews to generate outputs. This approach becomes increasingly complex as it scales, growing proportionally to the number of users, items and information sources. This is mitigated using scoring functions, which weight the different variables.

*Source: Reinforcement Learning based Recommender Systems: A Survey - https://arxiv.org/pdf/2101.06286.pdf - January 2021*

Capturing attention and steering behaviour is one thing when selling products, but as we have seen it poses risks when applied to news services and politics. In these domains, success is measured by user engagement, and this is often best achieved by promoting polarising content or views. eMarketer estimates US adults spent 13 hours 21 minutes per day consuming media in 2020, of which 2 and a half hours were spent on social media (Statista estimate).

**Linear TV Viewing by Age, Indexed to 2010**

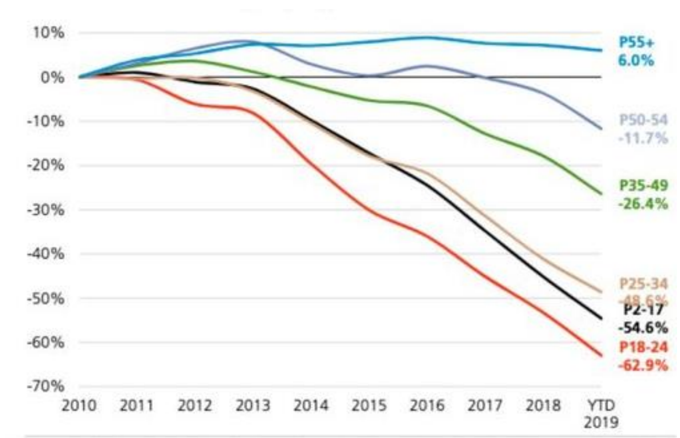


*Source: UBS analysis of Nielsen data - 2019*

**Traditional* vs Digital Media: Share of Average Time Spent in US, 2017-2023**

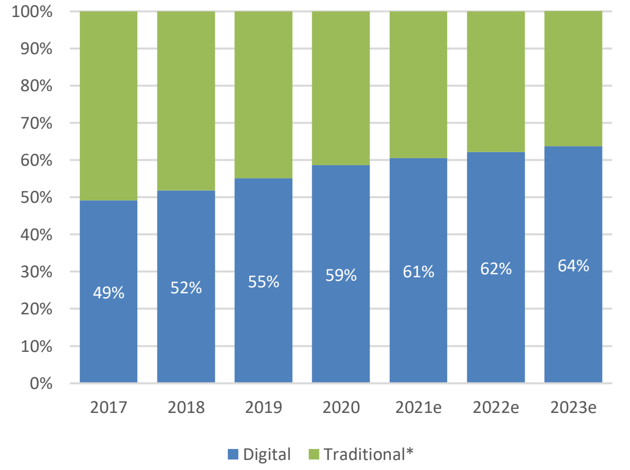


*Source: eMarketer; ages 18+; time spent with each medium includes all time spent with that medium, regardless of multitasking; *includes time spent on linear TV, newspapers, magazines and radio – April 2021*

It is hard to quantify the economic value of recommender systems, though they sit at the core of the FANGs' product offerings. As far back as 2006, Jeff Besos attributed 35% of Amazon's sales to 'cross-selling' (recommendations). Netflix has previously disclosed that "75% of what people watch is from some sort of recommendation" – indeed, they used to offer $1MM to any developer who could design an algorithm that beat their in-house system by a 10% margin. YouTube's Chief Product Officer put the percentage of time spent watching algorithmically recommended content at at 70%.


11 Albemarle Street, London W1S 4HH Tel: 020 3170 7420 Fax: 020 3170 7426. E-mail: info@greenash-partners.com Web: www.greenash-partners.com

## Natural Language Processing (NLP)

Natural language processing is a field that aims to teach computers to 'understand' natural language, with all of its nuance and context. Research in this area has been going on for decades, starting with symbolic NLP in the 1950s-1990s (rules based, manually annotated data), statistical NLP (identifying patterns through machine learning, unannotated data) in the 1990s-2000s, and neural NLP (a paradigm shift using deep learning), which has now largely replaced the statistical approach.

**Automated Speech Recognition (ASR)** is an increasingly widely used technology across industries, enabling automated transcription, as well as real time captions and subtitles (language translation). ASR is also fundamental to speech synthesis and speaker recognition, and thus voice assistants like Google Assistant, Alexa and Siri.

**Word Error Rates continue to improve (now down to 2.6 as of 2020)**

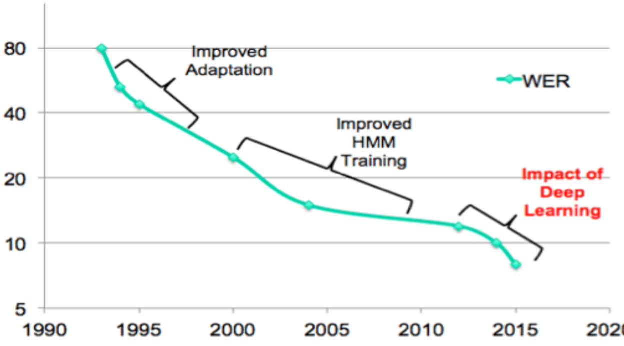


*Source: University of Southern California – April 2017; Improved HMM training relates to Hidden Markov Models which resulted in significant progress until around 2005. Performance improvements plateaued in 2005-2012. Deep learning led to sharp improvements from 2012 onwards. 2.6 value achieved on LibriSpeech dataset; source: Papers with Code, Stanford HAI - 2020*

**Forecasted Number of Digital Voice Assistants in Use – Set to exceed the global population by 2024**

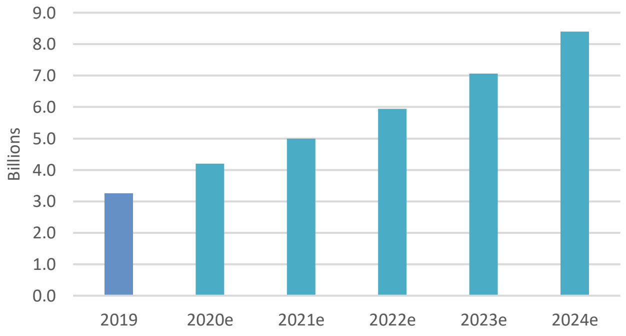


*Sources: Statista, Green Ash Partners – April 2020*

**Language Understanding Systems** - BERT (Bidirectional Encoder Representations from Transformers) is the most accurate language model in widespread use today. Originally designed by Google (and incorporated into their search engine in 2019), the model achieves human level performance in the SuperGLUE benchmark for language understanding, and now handles almost every English language query and has been expanded to 70 other languages. An adapted version of BERT was released by Microsoft in June last year, and this DeBERTa version (Decoding-enhanced BERT with disentangled attention) is currently the top of the SuperGLUE benchmark leaderboard.

Language understanding systems enable a broad suite of AI applications, including text mining, sentiment analysis, information extraction/summarisation, text generation and question answering. Progress in this area has made 'conversational AI' much more realistic and functional. While still a ways off from replacing humans, chatbots are playing an increasing role in customer services on a hybrid basis, and can competently handle repetitive, fact-based enquiries.

## SuperGLUE Benchmark for General Purpose Language Understanding Systems

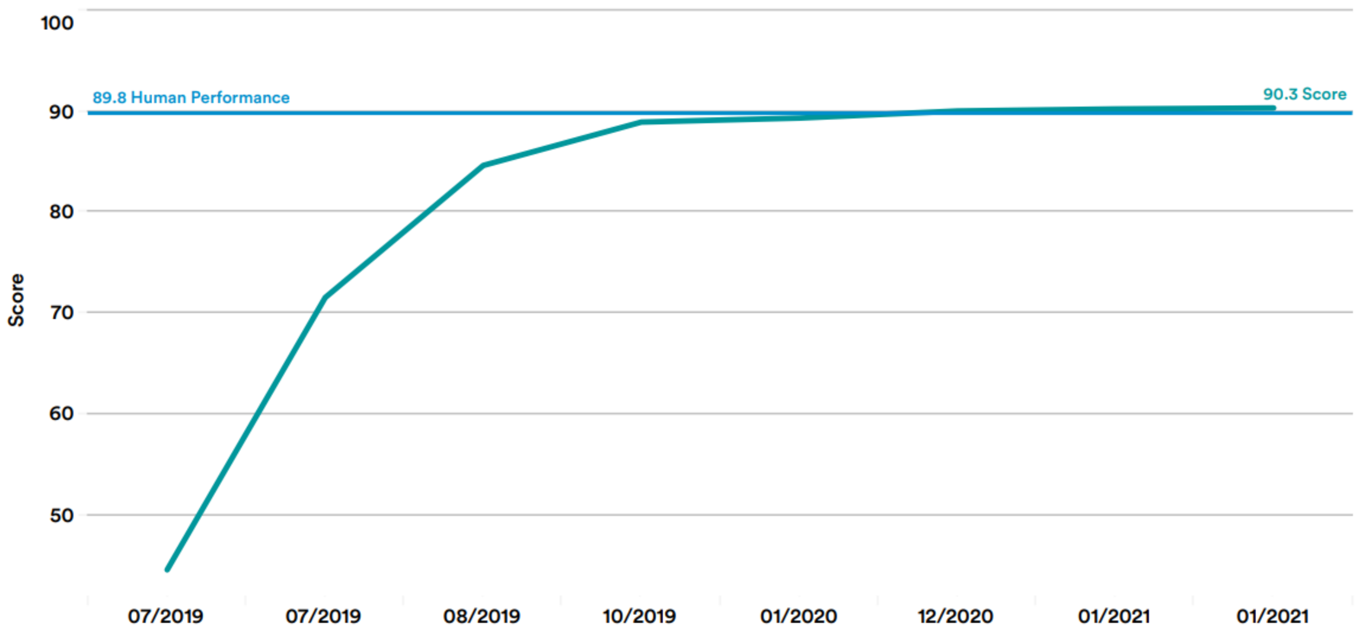


*Sources: SuperGLUE Leaderboard 2020, Stanford – 2021*

The SuperGLUE Benchmark chart above shows how quickly language understanding systems have advanced in just the last three years. Similar progress can be seen in Stanford's reading comprehension benchmark below, which measures how accurately an NLP model can provide answers to a series of questions pertaining to a paragraph of text.

## Reading Comprehension – The Stanford Question Answering Dataset (SQuAD)

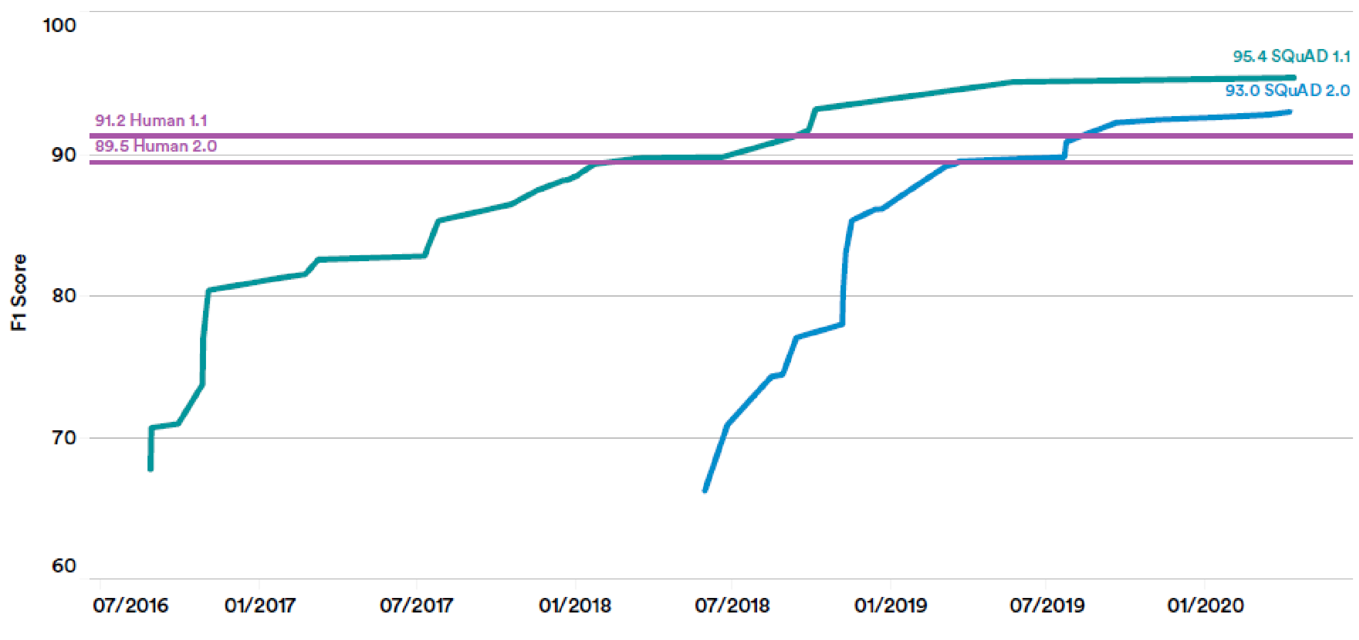


*Sources: Codelab Worksheets, 2020; Stanford – 2021*

## Computer Vision

Like NLP, the field of computer vision has existed for many years, but was accelerated with the rise of social media, Web 2.0, and the explosive growth of the datasphere.

**Image Recognition** – in 2006, AI researcher Fei Fei Li of Stanford University began work on a labelled database of images to assist with the training of AI models. ImageNet launched in 2009, and its dataset now comprises 14 million images across 20,000 classes. The mammoth task of image classification was achieved using Amazon's Mechanical Turk service, a platform offering modest compensation to millions of people around the world to complete small online tasks (initially Li tried using undergrads for $10 an hour, but realised this method would take 90 years).

Starting in 2010, the ImageNet project has run an annual contest, to see whose software programs can correctly detect and classify objects and scenes. The contest uses a trimmed down dataset of around a million images and 1,000 classes. In 2012, a program called AlexNet achieved a top-5 error of 15.3% - more than 10 percentage points lower than the next best entry. AlexNet used Nvidia's CUDA platform to make use of GPUs, as well as a deep learning architecture known as a convolutional neural network. It was a pivotal moment for the industry - the AlexNet paper has since been cited 92,000 times, and AI researchers and investors began to focus intently on deep learning.

In the years since AlexNet, deep learning has driven steady progress in image recognition (as well as many other areas), and today's image recognition models are approaching human-level performance – at least in the narrow field of image classification. Algorithmic improvements have been helped along by the exponential rise in capability, and decline in cost, of GPUs. Between 2017-2020, ImageNet training costs to reach 93% accuracy have fallen 166-fold, and training times have fallen 7,000-fold (from 13 days to 2 minutes 38 seconds).

**Image Classification on ImageNet (human performance >95%)**

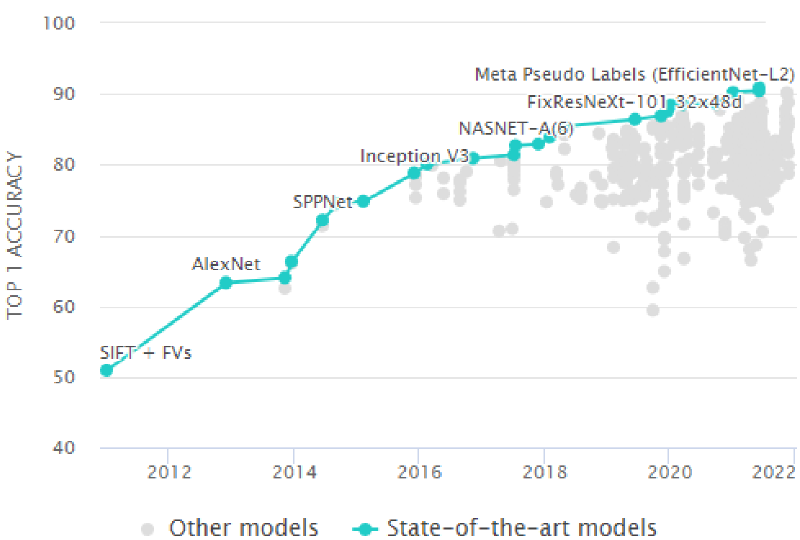


*Source: Papers with Code - 2021*

**ImageNet Training Costs (Logbase 10)**

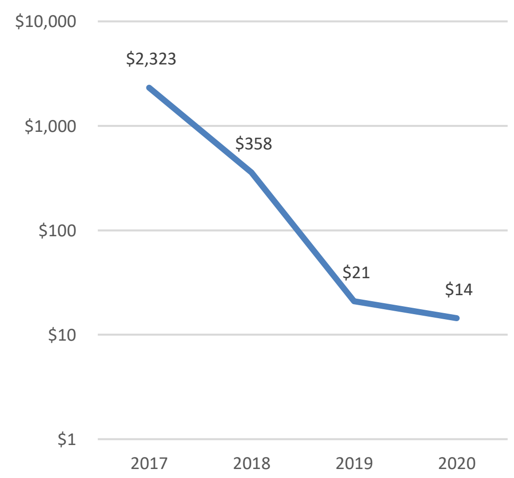


*Source: DAWNBench, Stanford, Green Ash Partners – March 2020*

**Object Detection** – Object detection can be seen as a subset of image classification, but has a deeper application in analysing videos and correctly classifying objects and people in real time. It also extends to activity recognition, which is the task of recognising activities from video clips, and things like human pose estimation, which generates 3D mesh models of a human body from a 2D image. Combining these programs can enable the creation of digital twins, improve surveillance systems both for security purposes and on factory product lines, as well as facilitate the autonomous navigation of robots in the physical world. Self- driving cars are probably the most well publicised and capitalised category of the latter (and are also an example of the industry overpromising and underdelivering).

**ActivityNet: Hardest Activities, 2019-20 – Drinking coffee remains the hardest, rock-paper-scissors showed the best YoY improvement**

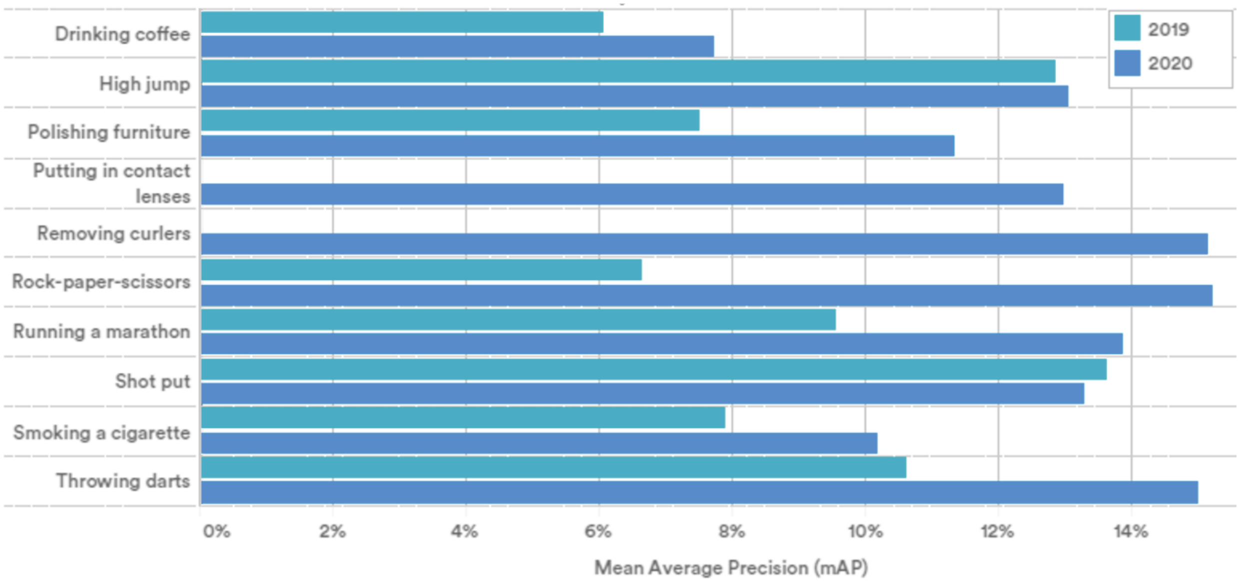


*Source: ActivityNet, Stanford - 2021*

**Object detection using a form of convolutional neural network called Mask R-CNN on the Detectron software platform offered by Facebook AI research**

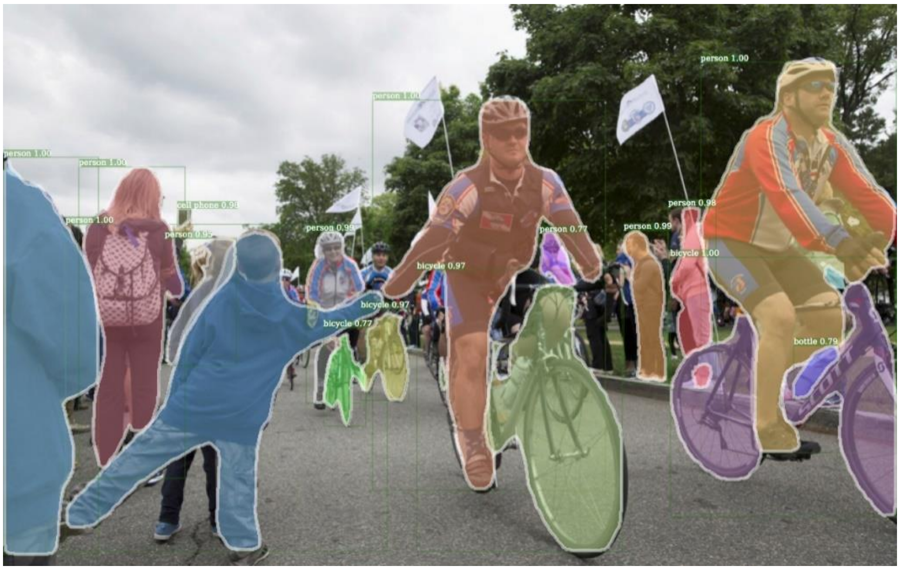

*Source: GitHub/FacebookResearch/Detectron - 2017*

**Image Generation** – AI driven image generation has taken off in recent years, since the invention of a type of neural network known as a generative adversarial network (GAN) in 2014. This approach uses a generative network, which produces candidate outputs according to its objective, and a discriminative network, which evaluates them. The generative network has the goal of 'fooling' the discriminator network, and the discriminator network tries to spot fakes. The two opposing networks improve through backpropagation.

GANs have improved rapidly, and they can now easily fool humans with still images. Video content generation is also improving, giving rise to realistic 'deepfakes', though identifying these is still (just) within human capabilities. Here is a YouTube clip with some examples.

**GAN Progress on Face Generation**

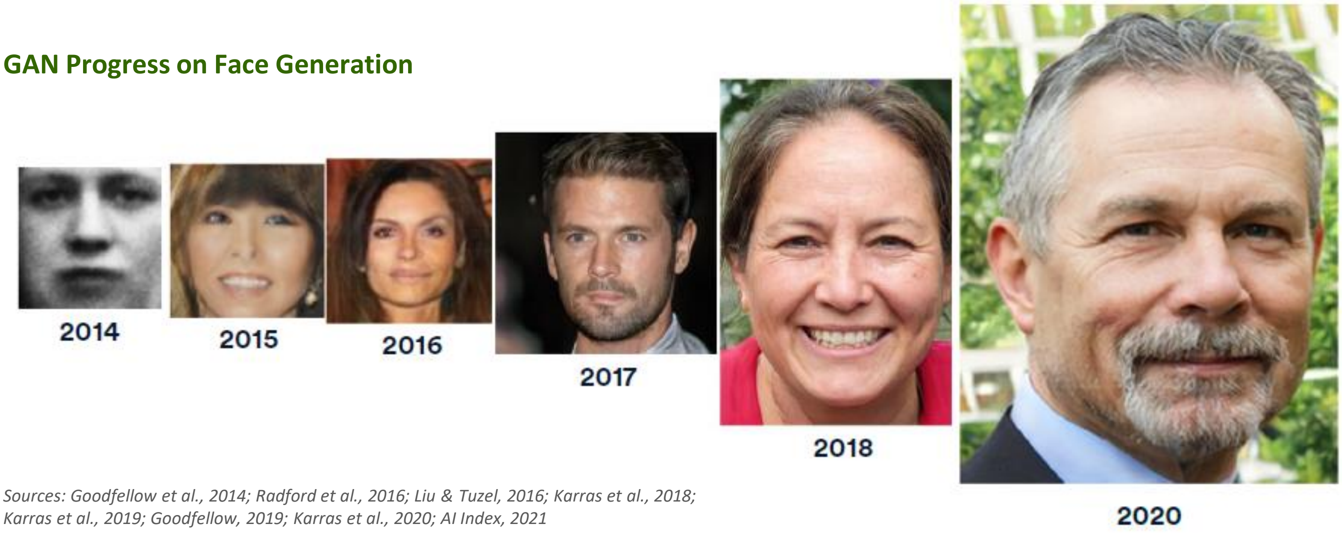


*Sources: Goodfellow et al., 2014; Radford et al., 2016; Liu & Tuzel, 2016; Karras et al., 2018; Karras et al., 2019; Goodfellow, 2019; Karras et al., 2020; AI Index, 2021*

As the capabilities of GANs expand, so to do their applications:

**Creative industries** - the ability to synthesise photorealistic images and video has many applications across advertising, videogames, film, architecture, design and art. GANs are accessible to creators with no background in coding, which has led to all kinds of creativity. Readers might have used a filter that can make oneself look older or younger, or an app that can superimpose another face on a film clip or music video. In fashion and marketing, fictitious models can be generated to show off clothing and other products without the need to organise expensive and time consuming photoshoots.


**Healthcare** – GANs can also be applied to medical imaging, due to their proficiency in upscaling resolution and noise reduction. This could improve the quality of magnetic resonance imaging, X-rays and low-dose scanning in computer tomography, though there is a much higher bar for accuracy in this kind of application due to the more serious consequences of errors. Trials are ongoing, but at the moment all outputs in this space undergo evaluation by multiple experts to ensure there are no distortions that could lead to an incorrect diagnosis.




11 Albemarle Street, London W1S 4HH Tel: 020 3170 7420 Fax: 020 3170 7426. E-mail: info@greenash-partners.com Web: www.greenash-partners.com

## The Impact of Deep Learning

"AI is one of the most important things humanity is working on. It is more profound than...electricity or fire"
*– Alphabet CEO Sundar Pichai, February 2018*

"AI could potentially deliver additional global economic activity of around $13 trillion by 2030, or about +16% higher cumulative GDP compared with today."
*- McKinsey & Company, September 2018*

As alluded to in previous sections, deep learning, a subset of machine learning, is responsible for the material acceleration of AI performance, and its expansion to new domains over the course of the last few years. This progress also accelerated interest and investment in the space, from both 'Big Tech' players and venture capital.

**AI Startup Funding Worldwide 2011-2021**

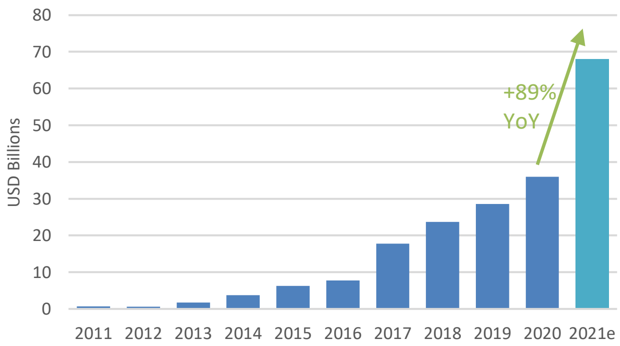


*Sources: Venture Scanner, Statista, CB Insights, Green Ash Partners; July 2021*

**Combined R&D Spend by Amazon, Alphabet & Facebook**

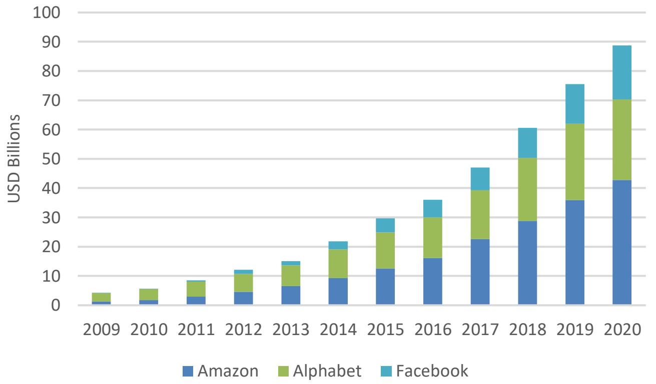


*Sources: Bloomberg, Green Ash Partners – December 2021*

AI continues to benefit from advances in chip architecture that support the massive parallelism demanded by ever larger neural networks. This is mostly driven by improvements in graphics processing units (GPUs), but has also been aided by the development of an application-specific integrated circuit (ASIC) called a Tensor Processing Unit (TPU). TPUs were developed by Google in 2015 (and reportedly co-designed Broadcom). Google's TPUs, which are currently on their 4th generation, have been publicly available on Google Cloud since 2018. Broadcom's ASIC revenues have grown from $50 million in 2015 to a run rate of over $1 billion. JPMorgan estimate the compute acceleration segment of the ASIC market, which also includes Data Processing Units (DPUs), is growing at a forward CAGR of +25-30%.

**AI Hardware Market Worldwide 2018-2025**

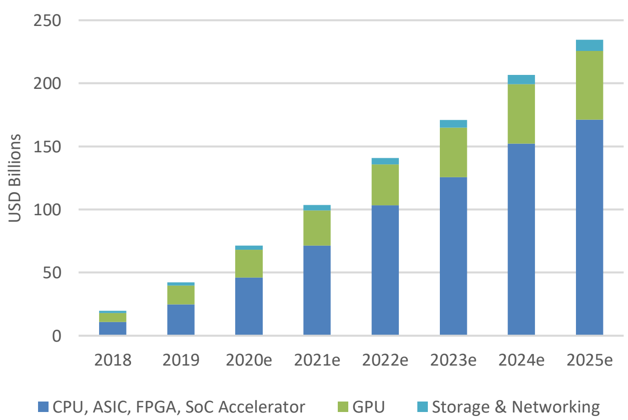


*Sources: Statista, Tractica, Green Ash Partners – April 2019*

The availability of high performance computation, large datasets, and more sophisticated deep learning architectures has created an emergent technology, which Stanford's Human-Centered Artificial Intelligence department (Stanford HAI) has labelled a *foundation model*. These have arisen from a deep learning approach used for NLP called a transformer, and are characterised as models that can be trained on broad, multi-modal data, at scale, and which are adaptable to a wide range of downstream tasks. A key feature is that of transfer learning – the ability to take knowledge learned from one task (e.g. object recognition) and apply it to another (e.g. activity recognition). As shown below, a foundation model is trained on various types of data, identifying statistical relationships between the different datasets. From there, it can be fine-tuned – adapting it for different tasks.

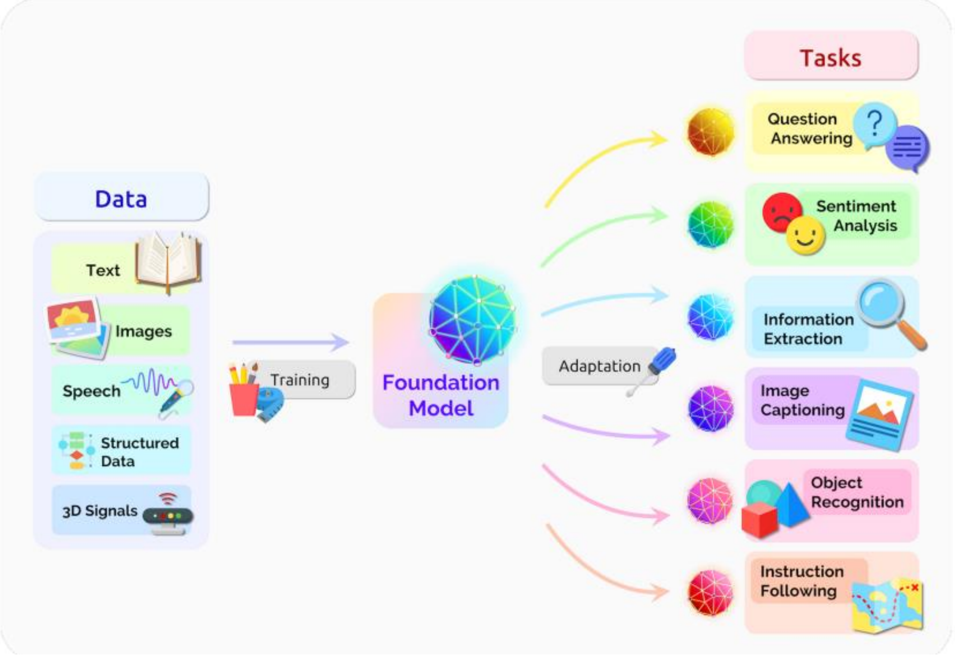


*Source: Stanford Center for Research on Foundation Models (CRFM) - 2021*

One of the most widely publicised foundation models is GPT-3 (Generative Pre-trained Transformer 3), which was created by OpenAI in 2020. It was a successor (unsurprisingly) to GPT-2 which was released in 2019; this is notable for the fact that in just a year GPT-3 was created with a hundred-fold higher parameter count and a thousand-fold larger training dataset relative to its predecessor. GPT-3's 45 terabyte training dataset equates to nearly a trillion words, and the model was trained with 175 billion parameters. As well as improving accuracy, this scale led to emergent capabilities that were absent from GPT-2. Some examples of this are:

- Writing newspaper articles – here is an article written by GPT-3 on why AI means us no harm
- Coding – GPT-3 can write code in a dozen programming languages, including CSS, JSX and Python. Microsoft have recently adapted GPT-3 to launch a low code/no code platform called Power FX which allows users to write programs using natural language. This feature is also making its way into spreadsheet software like Excel and Google Sheets, which will enable users to add formulas using natural language
- Conversational AI – GPT-3 can convincingly mimic stylistic language, allowing simulated conversations with literary or historical figures. Experiments range from the puerile (generating Donald Trump tweets) to the pathos-inducing story of a man simulating his fiancé after she passed away, to help him grieve
- As is the case with nearly all leading-edge technologies, models like GPT-3 will be embraced by the videogame industry, where it could be an invaluable tool simulating dialogue for non-player characters (NPCs), and even designing characters, objects and environments without human assistance.


11 Albemarle Street, London W1S 4HH Tel: 020 3170 7420 Fax: 020 3170 7426. E-mail: info@greenash-partners.com Web: www.greenash-partners.com

GPT-3's early applications show impressive natural language capabilities, however perhaps more significant is the improvement in zero-shot, one-shot and few-shot learning accuracy that comes with increased scale. This suggests continued improvement in task accuracy with minimal context can be expected as models continue to get larger.

**GPT-3: Average Performance Across 42 Accuracy-Orientated Benchmarks**

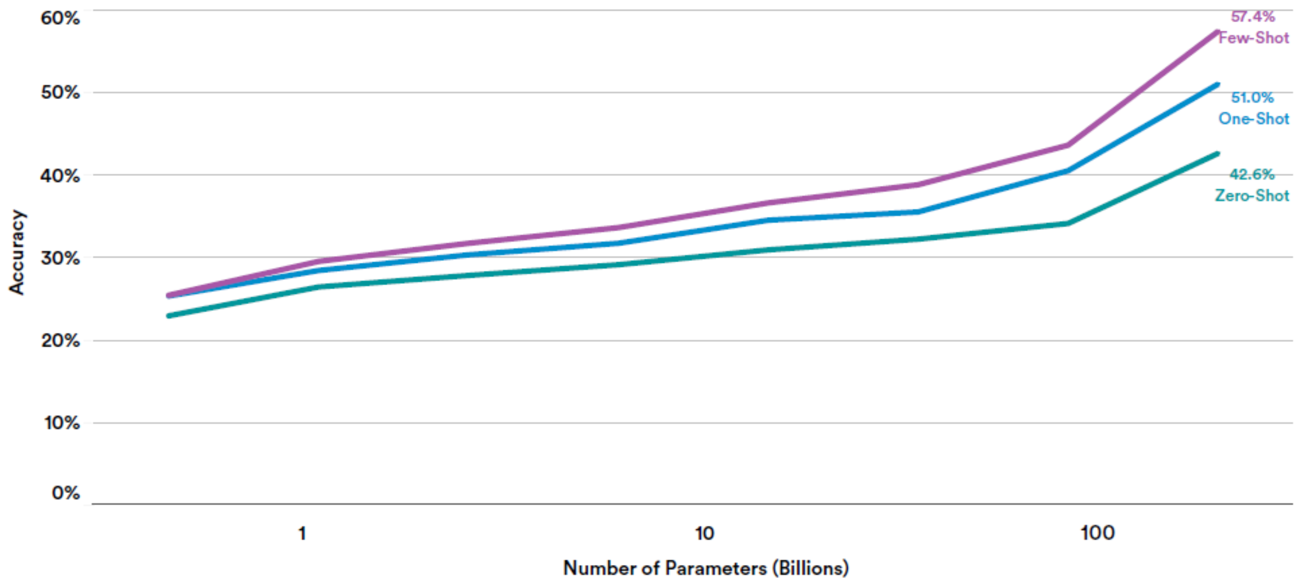


*Source: OpenAI (Brown et al.); Stanford - 2020*

A year after the release of GPT-3, a 175 billion parameter model, Microsoft and Nvidia teamed up to release MT-NLP – a 530 billion parameter model. In January 2022, researchers at Google Brain published a paper on a 1.6 trillion parameter model called a Switch Transformer, which uses a 'Mixture of Experts (MoE) routing algorithm to materially reduce computational costs while still achieving state-of-the-art accuracy in language modelling and machine translation (7x speed-up versus T5, Google's largest language model in 2020). AlexNet 2012 had just 60 million parameters, and since then compute used in the largest AI training runs has increased by more than 300,000x (3.4 month doubling period)[1]

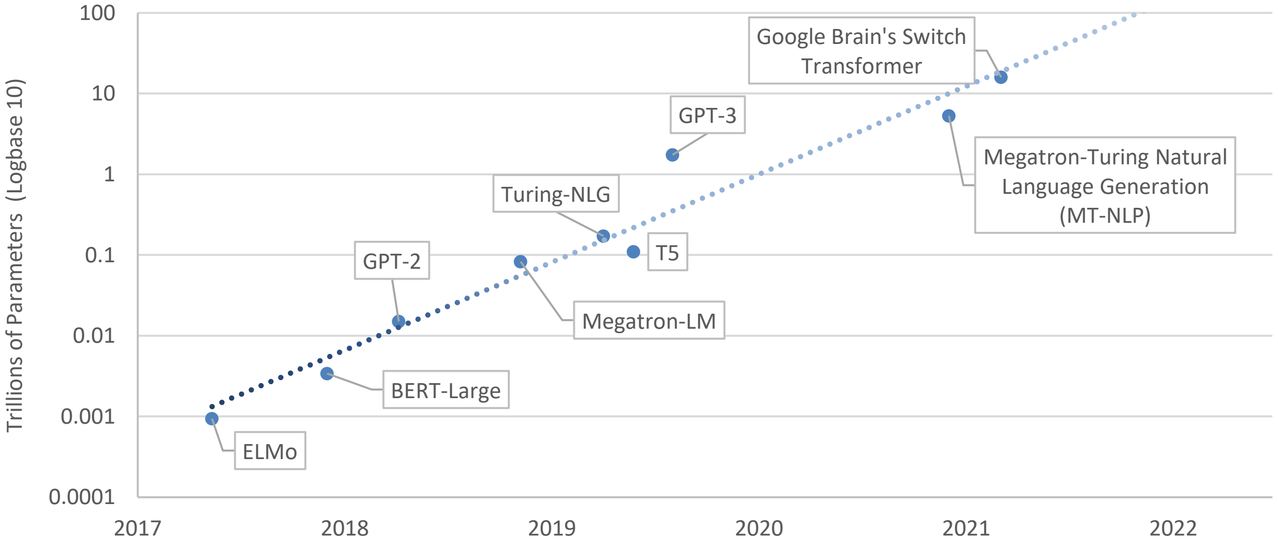


*Source: Green Ash Partners - 2021*

[1] *OpenAI (https://openai.com/blog/ai-and-compute/)*

## What Is Next For AI?

It is important to note that while deep learning has greatly broadened the applications for AI, the technology is still at its core a statistical engine, albeit a sophisticated one. NLP models like GPT-3 do not ‘understand’ language, and no deep learning approach in use today is likely on a path to eventually create Artificial General Intelligence (AGI). This is good news for philosophers like Nick Bostrom, who wrote *Superintelligence: Paths Dangers, Strategies* in 2014, when the early successes in deep learning were creating peak euphoria/hysteria in the press (a year later Elon Musk called AI humanity’s greatest threat). Since then, industry experts like Gary Marcus and Ernest Davis have gone the other way, arguing in *Rebooting AI* (2019) that deep learning creates ‘idiot savants’, and many tasks in which AI putatively exceeds ‘human-level’ capabilities are far narrower than many real world applications. If the first view is a bit sensationalist, the second is overly reductive – there are ‘narrow’ tasks that are excessively time consuming or repetitive, where AI can and does add real value. There are also tasks that involve such huge volumes of data that it is impossible for any single person to approach them.

Without AI, the personalisation of our online experience would be greatly diminished: browsing the internet would be marred by irrelevant advertising, content and products; finding information would be harder due a more rigid and unnatural web search; eCommerce would have greater friction at the checkout; and fulfilment/supply chain logistics would be less efficient, so shopping online would be more expensive.

While the first major roll out of AI has focused on the consumer, perhaps most intriguingly, transformer models have recently demonstrated capabilities in fundamental science and mathematics, suggesting deep learning could accelerate research in fields like biology, chemistry and physics. Healthcare is ripe for disruption by AI, and parallel advances in robotics are in the early stages of making factories and warehouses ‘smart’.

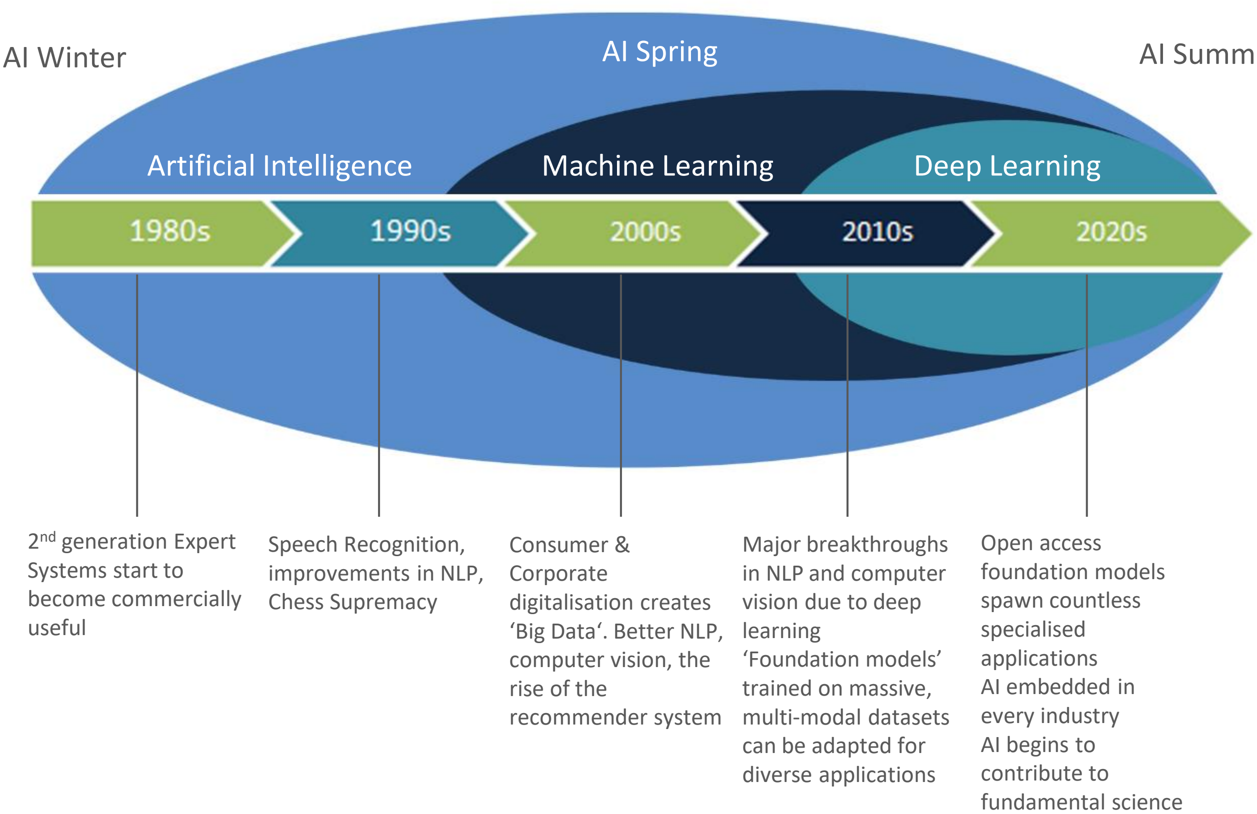

By its literal definition, it looks like Moore's Law may come to an end in the next few years, as the miniaturisation of transistors approaches the limits of atomic physics. That said, computation capability may actually accelerate this decade due to ongoing innovation. From the perspective of AI at least, TPU and GPU architectures have leapfrogged the more basic measure of transistor density as the main drivers of performance gains, and algorithms continue to become more efficient, able to do more with less (Google's DeepMind just announced a language model called RETRO that can match the performance of neural networks 25 times its size). In the latter half of the decade, we may see noisy intermediate scale quantum computers reach a size that becomes useful in machine learning tasks.

In On  the Horizon #1 - Quantum Computing, we outlined a number of transformational use cases for quantum computers that are outside of the reach of even the most powerful classical computer. But smaller quantum computers should be able to make meaningful contributions well before full quantum supremacy. In November, Dell announced a partnership with IonQ to build a hybrid classical-quantum platform, with the goal of leveraging the strengths of each to solve problems.

This is significant because the largest machine learning models require a lot of time and money to train, which puts them out of reach of academia and smaller companies. While not formally disclosed, it is estimated it would have cost $4.6 million to train GPT-3 on a Tesla V100 cloud instance from Nvidia (and would have taken a single GPU 355 years). OpenAI used a 10,000 GPU supercomputer built exclusively for them by Microsoft, on which researchers estimate would have taken 34 days of training time. As well as time and money, datacentres are extremely energy intensive, responsible for an estimated 1.1-1.4% of global energy use (IEA estimate). Adding quantum processors (QPUs) to machine learning system could theoretically reduce training times from days and months to minutes, reducing energy use and computation costs by a similar order of magnitude.

We expect AI capabilities to improve exponentially in the 2020s, leading to a broad range of new use cases across industries:

**The Metaverse -** We wrote about the Metaverse and Web 3.0 in On the Horizon #2 – The Metaverse; AI has a major role to play as the software infrastructure that stiches together the diverse aspects of the Metaverse experience, whether it be in relation to commerce, searching for information, or socialising/entertainment. Wherever we are on the spectrum of immersive virtual reality, augmented or mixed reality, or the screen-based experience that predominates today, we will need AI to regulate seamless transitions from online to offline. Expanding the concept to Nvidia's 'Omniverse' pulls in areas like Industry 4.0, smart cities, smart grids, smart networks/edge computing and automated mobility, all of which will need reliable and competent AI. Investment dollars in this theme will be huge – in December 2021, Intel's Raja Koduri wrote in a blog post that realising the vision of the Metaverse would need "a 1,000-times increase in computational efficiency from today's state of the art" to power "persistent and immersive computing, at scale and accessible by billions of humans in real time".

**Healthcare** – healthcare expenditure equates to 18% of US GDP and is a 9% weight in the US CPI basket. So far it has neither benefitted from the deflationary force of exponential technologies, nor has it become personalised in the way our online experience has. Ageing demographics make this a crucial target for disruption and improvement, though risks are high, given the potentially dire consequences if things go wrong.

- **Medical Imaging** – In 2018, a CNN designed by Google beat most dermatologists at identifying dermoscopic melanomas (from a group of 58 dermatologists, including 30 experts). There are only 12,000 dermatologists in the US, serving 330 million Americans, so AI can play an important role in increasing their screening bandwidth and improving accuracy. Melanomas have a 99% five year survival rate if detected early, versus 14% if detected in late stages. One in five Americans develop skin cancer in their lifetime (5.4 million new cases per year, at a cost of $8BN). Neural networks' prowess in image classification is moving on to the field of radiology - again, the goal is to provide a tool to improve accuracy and efficiency rather than to automate radiology. There are 800 million X-rays

undertaken in the US per annum, with false positive rates of 2% and false negatives of 25%. 31% of US radiologists have experienced a malpractice claim[1].

- **Genomics & Proteomics** – One unexpected feature of transformer models is the transferability of their proficiency with natural language to the language of life. In November 2020, the 14th Critical Assessment of Techniques for Protein Structure Prediction contest (CASP) had their AlexNet moment, as AlphaFold 2, a model designed by Google's DeepMind, achieved a paradigm shift in accuracy for modelling the decades-old protein folding problem. They followed up this achievement by open sourcing their program, and with a paper published in Nature in July 2021.

**Median Free-Modelling Accuracy**

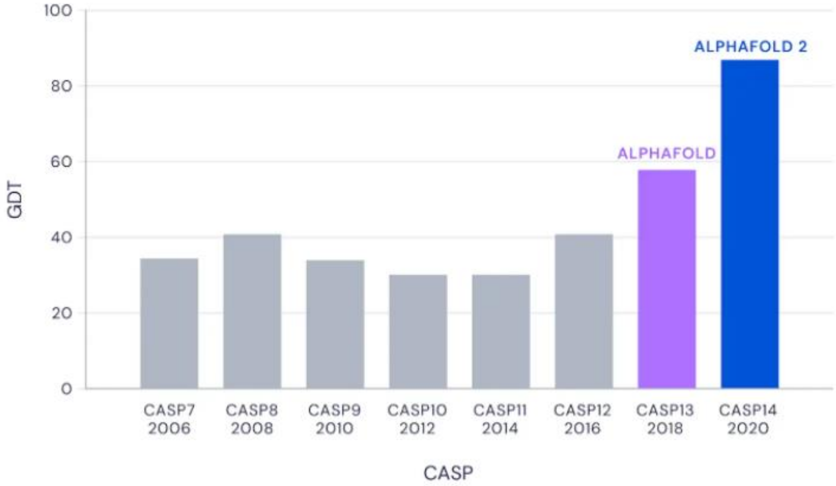


*Source: DeepMind*

Since AlphaFold 2, DeepMind have adapted a transformer model to predict gene expression, a major advancement for the field of genomics, which shows great promise in unlocking the secrets of longevity and banishing disease. We have recent, impactful examples of the power of AI in medicine from the COVID pandemic, where the development of vaccines and monoclonal antibody treatments were accelerated by next generation gene sequencing and AI.

- **Personalised Medicine** – Advances in genomics open up the possibility of personalised medicine. Doctors will have the data to better understand an individual's susceptibility to different illnesses and cancers, give better advice on diet and exercise, and diagnose diseases earlier, saving healthcare systems and patients vast expense. Even something as basic as patient records can be revolutionised by better natural language processing and understanding models, freeing up a doctor's time to tend to patients by automatically transcribing notes, and sifting through the disparate, unstructured data contained a patient's medical file. AI could even help make better doctors. Tens of thousands of medical papers are published every year – far more than single person could ever read. A sufficiently competent AI could monitor new publications, highlighting relevant advances in a given field.

- **AI Driven Drug Discovery** - Tufts Center of Drug Development Research and Deloitte both released studies in 2016 estimating a cost of around $1.5BN to bring a new drug to market, with a failure rate of about one in ten, and an average development time of 12 years (in the US). IQVIA estimate the global pharmaceutical market is worth $1.3 trillion annually, so there is a massive opportunity for any company that can reduce costs and shorten timelines for the industry. This is reflected in Pitchbook

[1] *Statistics in this bullet point come from Deep Medicine, by Eric Topol; see Further Reading*

data, showing increasing interest with the space, as well as in public markets, which have seen at least six AI-first biotech companies IPO in the last two years (Schrodinger, AbCellera, Recursion Pharmaceuticals, Relay Therapeutics, Absci Corp and Exscientia). These companies were all start ups founded in the 2010s, and have numerous partnerships with Big Pharma across hundreds of indications. If AI can demonstrate similar capabilities in drug discovery as in its proven domains, we may see the FANGs of the 2020s arise from this sector.

**Total VC Deals in AI Drug Discovery Over Time**

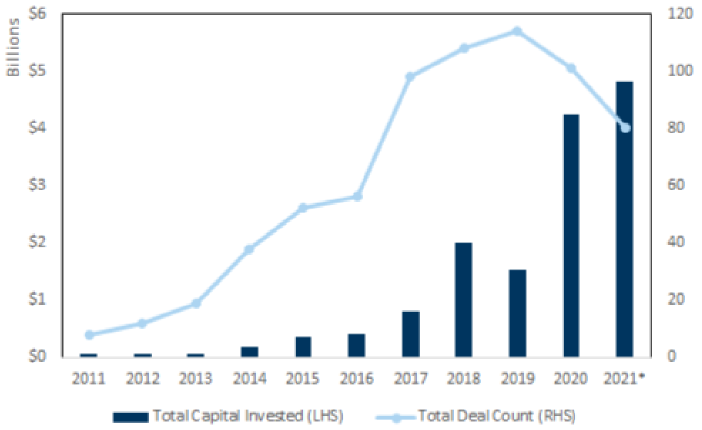


*Source: Pitchbook Data, Goldman Sachs Research*

**Professional Services** – Things like automated transcription are already very useful in many professions, including ours at Green Ash, given the number of quarterly earnings calls we follow. Sentiment analysis and keyword searches are used by some investment strategies as a new class of 'alternative data' in the hunt for a competitive edge. Better natural understanding in NLP models would open the way for the next major leg of productivity gains, enabling the summarisation of legal documents and many other paperwork or process intensive areas. Companies like UiPath and Automation Anywhere aim to use Robotic Process Automation (RPA) to reduce the billions of hours office workers spent on repetitive daily tasks. In 2018, the Trump Administration toyed with the idea of moving financial account filings to semi-annual from quarterly to reduce the regulatory burden on public companies – automating this reporting cycle might achieve the same goal, while preserving transparency, which is an increasingly important consideration for investors.

**Environmental, Social, and Governance (ESG)** – Following COP26, there is real momentum behind ESG investing, and regulatory bodies and independent agencies around the world are working together to come up with standards that can be incorporated into analysis of companies and countries. Monitoring compliance will become the next challenge, and the space industry may become central to this effort. Recently listed Planet Labs has a constellation of 200 satellites, that have photographed the entire surface of the Earth every day for seven years. They can create visual time series' of natural disasters like floods or droughts, and even monitor methane emissions from the oil & gas industry. Their satellites have uncovered human rights abuses, secret missile silos, and troop movements around the world. The company's >7 petabyte archive is the perfect dataset for CNNs and other machine learning approaches in image classification, and will eventually help index the Earth – making its 50 million square kilometre surface as searchable as the Internet.

**Fundamental Science** - GANs have been used to improve image quality in astrophotography and can even simulate gravitational lensing for dark matter research. At the other end of the scale from cosmology, they are also being applied to particle physics, for example to model realistic radiation patterns from high energy particle collisions.

[1] *Statistics in this bullet point come from Deep Medicine, by Eric Topol; see Further Reading*

## In Summary

Breakthroughs like deep learning in the 2010s have set the conditions for 'AI Summer' in the 2020s. More recently, a collegiate culture of open-sourcing models and publishing research is expanding the field and fostering accelerated innovation. Algorithms are getting more efficient, hyperscale datacentres more powerful, and we may see new technologies like quantum computing lead to a step change in the capabilities of AI in the coming years.

When we first started focusing on AI from an investment point of view back in 2013, the only real options for AI-first companies in the public markets were the Big Tech stocks. Since then, the opportunity set has expanded significantly, especially in the last few years, following a busy IPO market and proliferation of SPAC listings in 2020 and 2021. Between announced SPAC deals and IPO listings, investors can gain exposure to pure plays in quantum computing (IonQ, Rigetti and soon Quantinuum), self-driving systems (Aurora, Embark, TuSimple), industrial and agricultural automation (Symbotic, AppHarvest), robotic process automation (UiPath, Automation Anywhere to follow), conversational AI (Soundhound), and AI-powered cybersecurity (Darktrace, SentinelOne). There have been listings in entreprise AI (C3.ai, Palantir), and of course financial services have also been a big focus, from AI-powered credit lending (Upstart) to insurance (Lemonade) We have already mentioned half a dozen AI-first IPOs in healthcare, and in addition to these, several prominent SPAC targets have been in this sector, including AI-enabled companies like BenevolentAI, 23andMe, Babylon and Quantum-Si.

The list above is not exhaustive, and we offer no investment opinion on the companies mentioned, however it is intended to highlight the broadening of innovation beyond the handful of large companies that have dominated the AI space over the last decade. We expect this Cambrian explosion in AI related listings to continue in the coming years, as the many start ups waiting in the wings mature and go public. It is an exciting time to be a technology investor when so many industries are on the cusp of profound change.

## Further Reading


Super Intelligence: Paths, Dangers, Strategies - Nick Bostrom (2014)
Homo Deus: A Brief History of Tomorrow – Yuval Noah Harari (2015)
The Technological Singularity – Murray Shanahan (2015)
The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World – Pedro Domingos (2015)
Weapons of Math Destruction – Cathy O'Neil (2016)
Life 3.0: Being Human in the Age of Artificial Intelligence – Max Tegmark (2017)
AI Superpowers: China, Silicon Valley and the New World Order – Kai-Fu Lee (2018)
Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again – Eric Topol (2019)
Rebooting AI: Building Artificial Intelligence We Can Trust – Gary Marcus and Ernest Davis (2019)
State of AI Report 2020 – Nathan Benaich and Ian Hogarth
AI 2041: Ten Visions for Our Future – Kai-Fu Lee and Chen Quifan (2021)
Gathering Strength, Gathering Storms: AI100 Study Panel Report – Stanford University (2021)
On the Opportunities and Risks of Foundation Models – Center for Research on Foundation Models (CRFM), Stanford University (2021)
State of AI Report 2021 – Nathan Benaich and Ian Hogarth (2021)


*[1] Statistics in this bullet point come from Deep Medicine, by Eric Topol; see Further Reading*

## AI Taxonomy

**Artificial Intelligence (AI)** – AI is an umbrella term that has been around since the 1950s. At its most simplistic, It can be applied to any machine-based system that can process a set of inputs to achieve a pre-defined goal.

**Artificial General Intelligence (AGI)** – The holy grail of AI research is to create a system that can match or exceed human intelligence. Also known as 'strong' AI, such a system would be able to apply itself to the full range of problem-solving capabilities we have as humans, such as the ability to infer causal relationships in the real world with imperfect knowledge through 'common sense' reasoning. Futurists like Ray Kurzweil and Max Tegmark view AGI as an inevitability as current technologies continue to improve; others argue it will be impossible for such a system to emerge through the incremental development of current approaches – strong AI and, hypothetically, machine consciousness would require as-yet-unidentified breakthroughs in both hardware and software.

**Expert Systems** – Expert systems aim to emulate the decision-making process of a human expert. These rely on highly structured datasets, and are rigidly programmed with if-then type rules. They are composed of a knowledge base, an inference engine, and a user interface. Expert systems were around as far back as the 1970s-80s, and were one of the first commercial applications for AI. Strengths: reliable performance with routine tasks; weaknesses: narrow, non-transferable to other tasks.

**Machine Learning** - Machine learning is a statistical approach to AI that seeks to identify patterns in large datasets and draw inferences without requiring explicit instructions.

**Deep Learning** - Deep learning is a subset of machine learning that uses 'neural networks' to attempt to simulate the behaviour of the human brain. A key advantage of neural networks is their ability to train on unstructured data, like text or images. The algorithm then uses techniques such as gradient descent and backpropagation to fine tune its outputs, to improve its accuracy.

**Types of Learning** - Both Machine Learning and Deep Learning approaches may use Supervised Learning (using manually labelled datasets to make predictions), Unsupervised Learning (detecting patterns in unstructured data), and reinforcement learning (self-optimising outputs through a feedback/reward mechanism).

**Foundation Models** – A paradigm shift; as defined by Stanford HAI: "models trained on broad data (generally using self-supervision at scale) that can be adapted to a wide range of downstream tasks. These models, which are based on standard ideas in transfer learning and recent advances in deep learning and computer systems applied at a very large scale, demonstrate surprising emergent capabilities and substantially improve performance on a wide range of downstream tasks."

LEGAL DISCLOSURES

FOR EU INVESTORS

The information contained in this document is issued by Green Ash Partners LLP (hereinafter "Green Ash"), at 11 Albemarle Street, Mayfair, London, W1S 4HH who is authorised and regulated by the Financial Conduct Authority: Firm Reference Number (FRN) – 500315. This disclaimer is not intended to exclude or restrict any liability under the rules of the Financial Conduct Authority ("the FCA Rules") or FSMA.

This presentation reflects the opinion of Green Ash as of the date of issue. This document is not intended for distribution, publication, or use in any jurisdiction where such distribution, publication, or use would be unlawful, nor is it directed to any person or entity to which it would be unlawful to direct such a document. This presentation is for information purposes only and does not constitute an offer or a recommendation to purchase or sell any security. It does not constitute investment research or a research recommendation and is not intended for distribution to the public or a large number of persons. The opinions herein do not take into account individual clients' circumstances, objectives, or needs. Before entering into any transaction, each client is urged to consider the suitability of the transaction to his particular circumstances and to independently review, with professional advisors as necessary, the specific risks incurred, in particular at the financial, regulatory, and tax levels.

All examples of financial strategies/investments set out in this document are for illustrative purposes only and do not represent future performance. The information and analysis contained herein have been based on sources believed to be reliable. However, Green Ash does not guarantee their timeliness, accuracy, or completeness, nor does it accept any liability for any loss or damage resulting from their use. All information and opinions as well as the prices indicated are subject to change without notice. Past performance is no guarantee of current or future returns and you may consequently get back less than he invested. From time to time the partners of Green Ash Partners LLP may enter into personal transactions in the securities and strategies discussed in this presentation. The firm has a personal account dealing policy that manages such conflicts and ensures the fair treatment of its clients.

This document may not be reproduced (in whole or in part), transmitted, modified, or used for any public or commercial purpose without the prior written permission of Green Ash.

FOR SWISS INVESTORS

This document is issued by Green Ash Partners LLP. This document is exclusively intended for qualified investors within the meaning of article 10 paragraph 3, 3bis, 3ter and 4 of the Swiss Collective Investment Schemes Act ("CISA"). The information and data presented in this document are not to be considered as an offer or solicitation to buy, sell or subscribe to any securities or financial instruments. Information, opinions and estimates expressed in this document reflect a judgment at its original date of publication and are subject to change without notice. Green Ash Partners LLP has not taken any steps to ensure that the securities referred to in this document are suitable for any particular investor and this document is not to be relied upon in substitution for the exercise of independent judgment. Tax treatment depends on the individual circumstances of each investor and may be subject to change in the future. You should obtain specific professional advice before making any investment decision. The value and income of any of the securities or financial instruments mentioned in this document can fall as well as rise and, as a consequence, investors may receive back less than originally invested. Risk factors are not intended to be reproduced in full in this document. Past performance is no guarantee nor a reliable indicator of future results. This material is not intended to be a substitute to the full, legal documentation and to any information which investors must obtain from their financial intermediaries acting in relation to their investment in the funds mentioned in this document. Please note that none of the management company, the registrar and transfer agent, the central administration or the custodian of the relevant fund has independently verified any information contained herein and no party makes any representation or warranty as to the accuracy, completeness, or reliability of such information. Further information about Woodman SICAV - Green Ash Global High Yield Fund, its prospectus, its KIIDs and its latest annual and semi-annual report may be obtained free of charge, in English language, from MultiConcept Fund Management S.A., Luxembourg, the CH Legal Representative, Credit Suisse Funds AG, the appointed distributors or online at www.greenash-partners.com and www.credit-suisse.com/ Multiconcept.