EDISON

BRILLIANT KNOWLEDGE

TMT

August 2021

ENTERPRISE ADOPTION OF AI

CUTTING THROUGH THE HYPE

After several false starts over the last few decades,
artificial intelligence (AI) technology is now seeing
real adoption by businesses across multiple sectors.
Benefits include increased business efficiency,
better customer service and
innovative new products and services.
We expect usage to follow a similar pattern
to cloud adoption, with early adopters pioneering
the use of AI, and wider adoption following
as the increasing availability of AI tools and services
democratises the development of AI applications.
In our view, while AI hype is still a factor,
a disruptive trend is now well underway
and companies that do not embrace it
are likely to be at a competitive disadvantage.

AI is suited to making sense of high volumes of data

As a natural evolution from big data analytics, the use of AI techniques such as machine learning and deep learning bring an element of intelligence to the analytics process. Models can be trained to identify patterns that would take humans too long to find or that they would not think to look for. Areas in which AI can be used include image recognition, trend analysis, prediction, natural language processing, pattern recognition, anomaly detection, personalisation and discovery.

Targeted use to augment human capabilities

AI is most effective when used for narrow applications, where there is a predefined goal like fraud detection or language translation. Processes best suited to AI are those with access to large quantities of data or time-consuming, recurring manual tasks that cannot be automated with standard software. Rather than replacing humans, the best models augment human expertise, providing outputs that help inform decisions rather than make them.

Increasing enterprise adoption

Companies are increasingly considering the potential of AI to enhance their businesses, whether for internal processes to improve efficiency, or to provide better or more innovative products and services. A range of software is available to companies to develop their own applications; alternatively, multiple off-the-shelf solutions exist to address specific use cases. Companies that already have good data governance and cloud infrastructure in place are better positioned to deploy AI technology.

Limited listed exposure to pure-play AI in Europe

Companies broadly fall into two categories: AI-native and AI-enhanced. As is often the case with new technologies, AI-native companies tend to be private, so there is limited opportunity to invest directly in the theme in Europe; listed examples include expert.ai, Darktrace, Insig AI and Mirriad Advertising. Investors may also want to consider companies that are enhancing their products or services with AI as we believe this will give them a competitive edge.

Not without challenges

The use of AI faces several challenges, including the lack of explainability of deep learning models, the risk of bias, the potential for unethical use, the need for access to large quantities of data to train models and the limited supply of data scientists. Regulators are considering ways to address the issues around bias and ethics.

CONTENTS

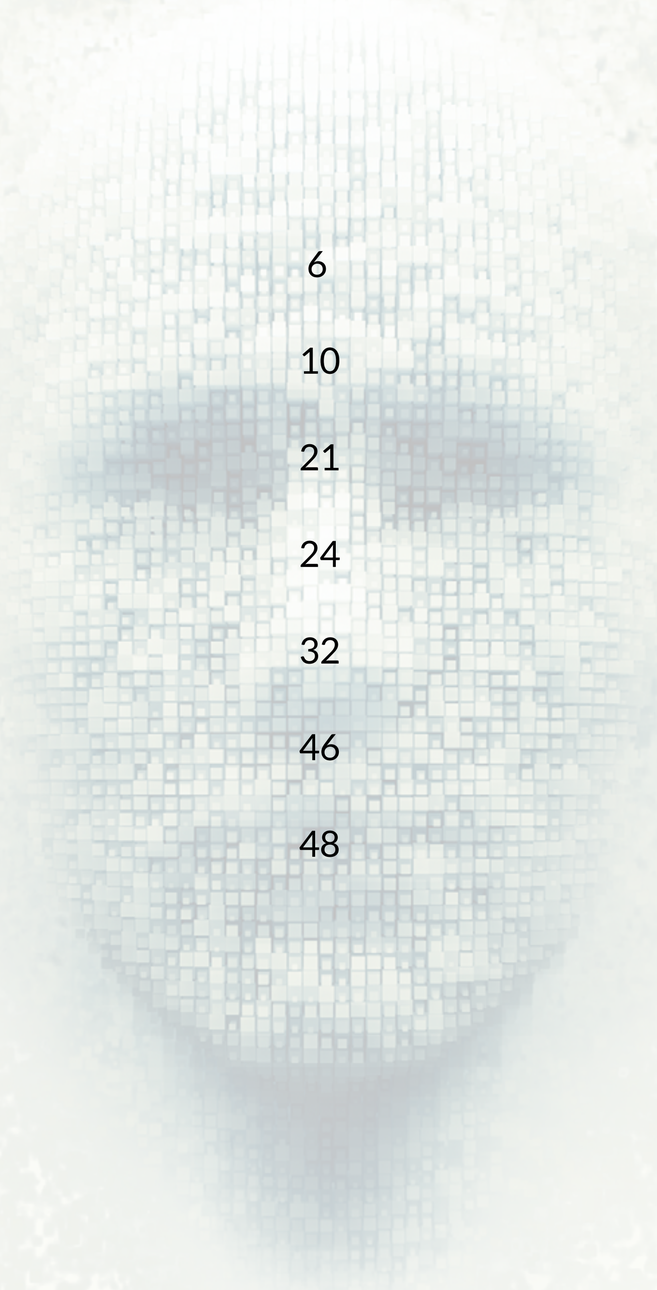

INVESTMENT SUMMARY: INCREASING ADOPTION OF AI

AI is real

After many years of false starts (or AI winters), the hype around AI technology is finally turning into reality. As technology has advanced to allow faster and cheaper processing power and storage capacity, this has provided the environment for AI pioneers to build and evolve their models. AI techniques such as machine learning and deep learning depend on the analysis of vast quantities of data. Global volumes of data have increased by a CAGR of c 40% from 2015 to 2020 and are forecast to grow another 65% by 2024, providing a further indication as to why progress in AI has accelerated in recent years (source: Statista).

The cloud service providers (Amazon, Google, Microsoft and Alibaba) and companies such as Apple and Facebook have been at an advantage, able to use the vast amounts of data at their disposal to develop and train their AI models. They have also developed a range of AI tools and services for use by their customers, democratising the development of AI applications while creating a new audience of customers to use their cloud services.

Good for pattern detection, recommendations, natural language understanding, visual analysis

AI can be used for a wide range of processes, distinguished by its ability to process both structured and unstructured data. Structured data is typically numeric, and AI adds an element of intelligence to standard software through pattern detection capabilities. Unstructured data is more qualitative and form over 90% of global data, so the ability for AI to analyse them highlights a significant advantage, given that historically computers have been limited to only manipulating numeric data.

Key examples include natural language understanding/processing for text or speech, and computer vision for image recognition (eg facial recognition, medical scans). Deep learning's strong pattern detection capability has also opened new avenues in discovery, predominantly in healthcare and science.

Practical applications abound

By concentrating on narrow applications, AI can now be used for a wide variety of tasks. It offers a range of benefits, including faster processing, better personalisation, the ability to gain insight into data-heavy processes, better decision making and improved used of in-house knowledge.

Examples of applications include fraud detection, language translation, chatbots, search, document recognition, credit scoring and face recognition.

INVESTMENT SUMMARY: INCREASING ADOPTION OF AI

Enterprise adoption is underway

Companies are increasingly looking at the potential of AI technology to enhance their businesses. It can be used for internal processes, with examples including hiring, churn analysis and back-office functions. AI can also be used to provide better or more innovative products and services. To enable this, companies across a range of sectors are building AI into their solutions. They can access tools from the cloud service providers or other software vendors and couple these with algorithms from open-source libraries to develop their own AI-powered applications. We believe that the availability of these tools is democratising the use of AI and we expect to see its adoption accelerate as companies experience the benefits. For those companies that do not have the in-house expertise or desire to develop applications themselves, there is a growing market of vertical-specific applications that provide enterprises with the advantages of AI.

Advice from early adopters

To have a good chance of success, AI projects should be based on a business use case, with the desired outcomes driving the process rather than the model. The project team needs a mix of skills, combining data scientists, software developers and domain experts. Projects can take time to design and implement and are therefore not well suited to agile development. Companies that have successfully implemented AI projects suggest that starting small to learn what works is the best approach. In fact, many found that once one project was successful, this led to many more ideas for how AI could be used elsewhere in the business.

Good data governance is crucial: a large majority of project time is typically spent preparing data compared to the time spent training the model, as poor data quality will lead to poor outcomes. While the adoption of AI can be a cause of concern for employees fearing that their jobs will disappear, in reality AI shifts the type of work done by employees. Involving staff in the process to help design the model and be involved in the training of the model and oversight once in production should produce a better outcome and better understanding throughout the business. Additionally, the use of AI should free up staff to concentrate on higher-value areas such as customer interaction or data analysis.

Active M&A and venture capital market

Corporate venture capital funding in AI companies has seen a significant uptick in recent years, with total investment growing by a CAGR of 233% to US$10.6bn between 2014 and 2019 and the number of deals increasing from 104 to 453 during the same period (source: Statista). Acquisitions by large technology companies have accelerated the development of the models created by AI start-ups, mainly by providing the vast amounts of data needed to train these architectures. Google has been most active in this space, having completed over 30 acquisitions relating to AI, including DeepMind (2014: US$500m) and Nest Labs (2014: US$3.2bn) (source: Refinitiv). Other notable names include Apple, Facebook, Microsoft and Amazon, which, in the last five years, have collectively acquired over 50 AI start-ups (source: Refinitiv).

INVESTMENT SUMMARY: INCREASING ADOPTION OF AI

In April 2021, Microsoft acquired Nuance Communications, an early pioneer of speech-recognition technology, for US$19.7bn.

How to invest in the theme

The largest players in the AI technology space are the US and Chinese mega-cap technology stocks: Amazon, Alphabet, Apple, Microsoft, Alibaba and Tencent. The US companies have been active acquirers of a large number of innovative AI start-ups as well as more mature AI-focused software companies and all invest heavily in AI-related R&D. In the US, while the cloud service providers and other large technology companies offer AI software and tools, there are very few listed AI-native companies. There is a plethora of privately held AI start-ups addressing specific vertical applications, many of which we expect to be acquired by larger established software companies to broaden their offerings. Looking at the European market, investors also have limited access to listed AI-native stocks (see Exhibit 10). Companies include expert.ai, Darktrace, Insig AI and Mirriad Advertising. It is possible to gain exposure to privately held AI companies via listed investment companies, however, we are not aware of any such companies that are solely invested in AI companies. Several privately held AI software companies are finding their way to the public markets via special purpose acquisition companies (SPACs). To date this has mainly been a US phenomenon, but with the UK's Financial Conduct Authority reviewing the rules around SPACs, this may become a more popular way to list in the UK. Another option is to consider companies that are enhancing their products or services with AI. If they are truly using AI techniques, rather than just using the term for marketing purposes, they should have a competitive advantage compared to peers that are not using the technology.

Challenges include explainability, ethics and availability of skilled staff

There are a number of challenges to the adoption of AI technology. A well-known example is the difficulty in auditing the results produced by deep-learning architectures, where it is not always possible to explain why a model came to the conclusion that it did. Research is underway to try to make AI more transparent and there are currently a number of methods that can help mitigate this issue around explainability.

Another concern is the possibility of bias in models, built in either through the design of the model or the data used, and this should be considered carefully when starting an AI project. As for technology in general, it is possible to use AI for unethical purposes; for example, using face-recognition technology to identify people from certain ethnic groups for discrimination or persecution. Regulation that balances potential harms with social opportunities will be key to ensuring that risks relating to areas like explainability and bias are mitigated, and although it may slow the pace of development, it will be critical for building user trust. Currently, the EU and the United States have started developing regulatory protocols, but as Google CEO Sundar Pichai says, 'international alignment will be critical to making global standards work'. Other regulation, such as GDPR, may act as a barrier to increased data collection.

INVESTMENT SUMMARY: INCREASING ADOPTION OF AI

As is often the case with leading-edge technology, skills are in short supply and data scientists can command high salaries. Companies are adapting to this by retraining staff, collaborating with universities, taking on graduates who are being trained up as data scientists and hiring offshore (made easier since remote working has increased during the pandemic).

Artificial general intelligence is a long way off

Artificial general intelligence (AGI) is the term given to AI that has reached human levels of intelligence, where one model can tackle a wide array of problems. To date, the greatest successes have been in 'narrow AI' applications, designed to target specific and well-defined problems with examples including chat bots or systems designed to recognise fraudulent transactions. Many scientists believe we are decades, if not more, away from developing AGI, so for now, humans are still better than machines at tasks that are complex or require emotional intelligence or creativity.

AN INTRODUCTION TO AI

Many companies today are quick to tout
their AI-powered products and services,
but what does being AI-powered actually mean?
In this report, we attempt to identify
where the use of AI is making a difference.
We explain the main types of AI technologies,
what applications they are best suited to,
and how, in particular, enterprise software companies
are adopting the technology within their own solutions.
We have interviewed a range of companies
across software and IT services
as well as companies that use AI to deliver services.
We have also recorded a series of videos
to share the views of these companies
on various AI-related topics.

AN INTRODUCTION TO AI

Defining AI

As a natural evolution from big-data analytics, the use of AI techniques such as machine learning and deep learning bring an element of intelligence to the analytics process. AI refers to any human-like intelligence exhibited by a machine, which can mimic the perception, learning, problem-solving and decision-making capabilities of the human mind (source: IBM). AI marks the next stage in software automation, where its ability to make inferences (creating processes and decisions without human input) significantly expands the scope of the mental processes that we can automate. Subsequently, we have excluded robotic process automation (RPA) from our report, as this mechanism only automates repetitive processes and does not have the ability to replicate brain power. Machine learning technology, on the other hand, has the ability to make inferences, as well as come up with new conclusions and is therefore the main area we will be focusing on. Exhibit 1 shows how the different types of AI explored in this report fit together.

Exhibit 1: Summary of the AI ecosystem

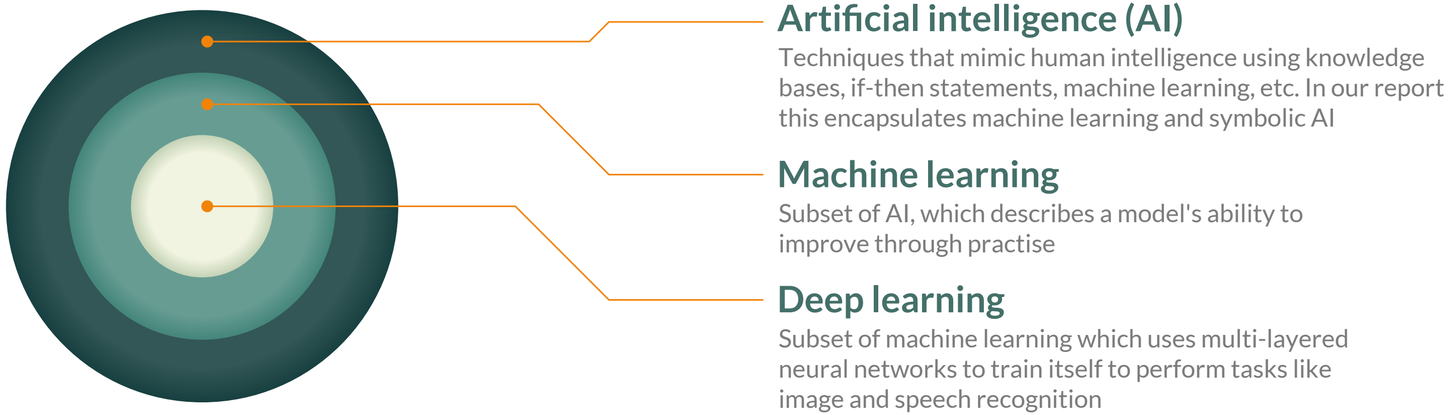

Source: Edison Investment Research

Three main types of machine learning exist: in supervised learning the model is trained using data, which is labelled or tagged before it is processed, unsupervised learning can work without the need for the user to label the data before training and reinforcement learning uses a reward-based system to maximise correct outputs.

Some commonly used algorithms for machine learning include clustering, decision trees, gradient boosting, long short-term memory, naïve Bayes and support vector machines.

AN INTRODUCTION TO AI

Neural networks are a subset of machine learning and are designed to work in a similar way to the human brain, where neurons pass and interpret information by using a nerve net. Humans process information via the five senses and subconsciously weight each component's importance, allowing for neurons in the brain to process it and create a solution based on probability. The key to a neural network is its hidden layers, which is where the model identifies patterns and weights the data by importance. Hidden layers are often referred to as a black box, as decisions made by the model are largely unknown to the user. For a fuller description, please refer to the Appendix on page 50.

Simple neural networks often only contain three or fewer hidden layers and are used to process structured/labelled data, which is quantitative, clearly defined and formatted in a way so that it is easily searchable, with examples including dates, phone numbers, etc.

Deep learning neural networks contain 10 or more hidden layers, enabling models to breakdown and interpret information at a more granular level. They can learn from and process unstructured data, which is qualitative and requires more work to process and understand, with examples including text, audio and images. Typically, the computational power of a deep learning neural network depends on the number of hidden layers and parameters used.

Symbolic AI, another type of AI, works by creating relationships between real world objects or 'symbols', such as numbers, letters, people, cars, etc, and then uses logic to search for solutions. This logic requires human intervention, separating it from machine learning, where learning happens within the architecture of the model. Please see the Appendix on page 52 for further explanation.

A brief history of AI

Alan Turing said in 1950 that 'a computer would deserve to be called intelligent if it could deceive a human into believing that it was a human', which would become the basis of his famous Turing Test, known at the time as the imitation game. Soon after, in 1956, John McCarthy coined the term 'artificial intelligence', defining it as 'the science and engineering of making intelligent machines'. However, progress at the time was inhibited by cost and memory; during the 1950s computers were only able to execute commands, with no ability to store them, and leasing a computer could cost up to US$200,000 a month[1]. These headwinds eased over time primarily explained by Moore's Law, where Gordon Moore, the co-founder of Fairchild Semiconductor and Intel, observed that the number of transistors on a computer chip was doubling every 18–24 months, while the cost of computing halved. Subsequently, it was not until the late 1990s/early 2000s that tangible progress in AI was achieved, highlighted by IBM's Deep Blue beating world chess champion and grand master Garry Kasparov in 1997. As Exhibit 2 illustrates, the most notable developments in AI have been in the last decade due to the evolution of deep learning, with one of the most recent examples being DeepMind's AlphaFold solving the historical protein folding problem in 2020.

1 https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/

AN INTRODUCTION TO AI

Exhibit 2: History of AI development

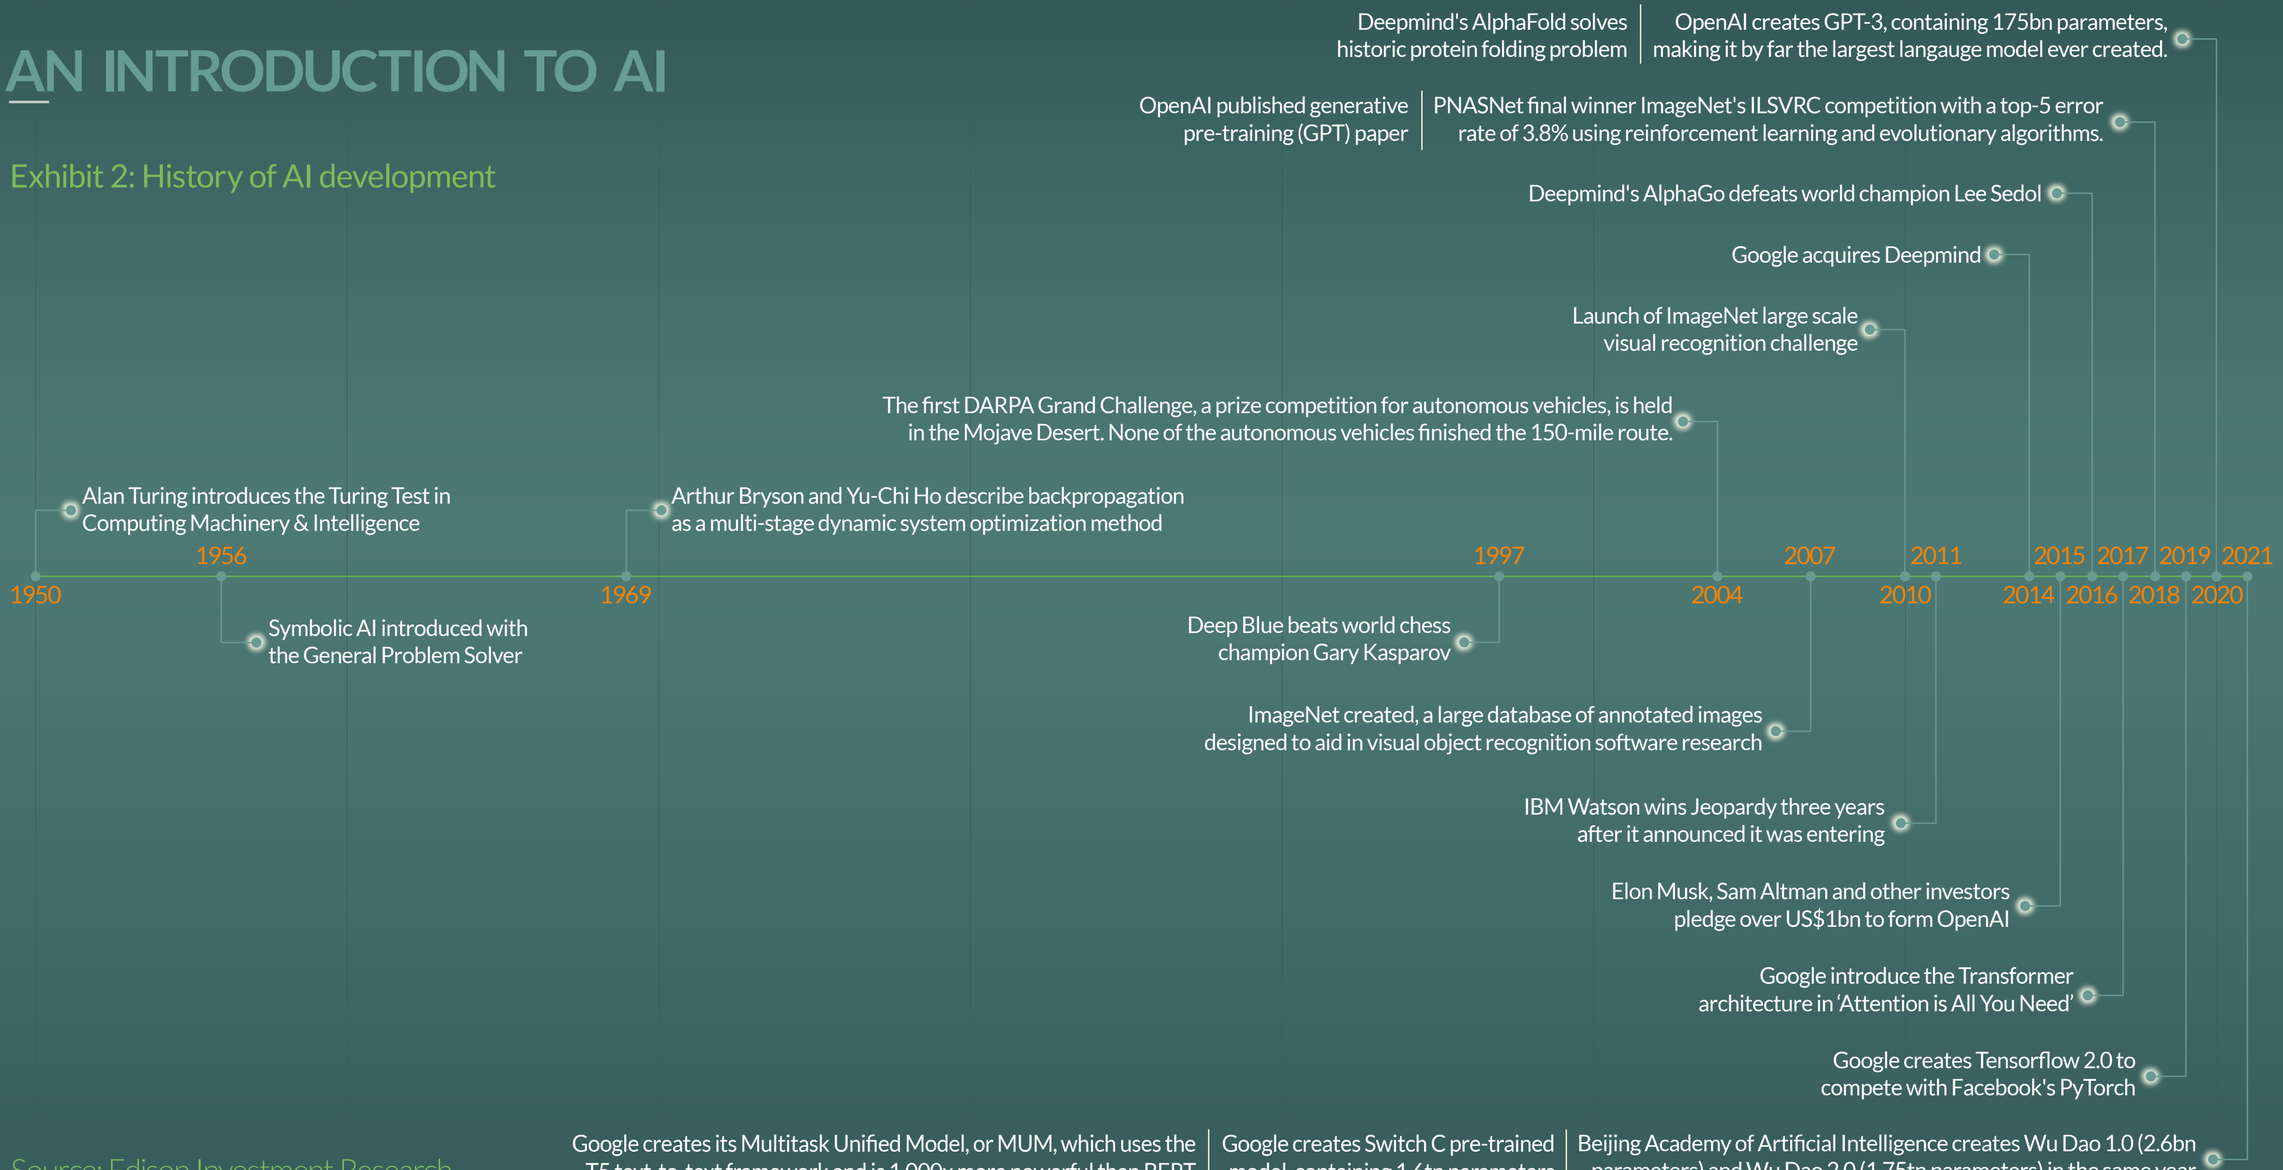

Source: Edison Investment Research

AN INTRODUCTION TO AI

Competition drives major advances in recent years

With AI and deep learning now a reality, competition among the major global players is driving AI's development, particularly in natural language processing (NLP). Companies have been rapidly surpassing one another in their ability to build larger language models with a greater number of parameters, which provides a tangible measure of performance. Notably, the Beijing Academy of Artificial Intelligence created the 1.75tn parameter model Wu Dao 2.0 in June 2021, just three months after it had released the 2.6bn parameter 1.0 version. Wu Dao 2.0 is now the largest ever AI model, exceeding Apple's 1.6tn Switch Transformer (January 2021) model by 175bn parameters and the famous GPT-3 by 10x. The progression in the size of these models over the last three years can be seen in the table opposite.

Exhibit 3: Selection of pre-trained model sizes, 2018–21

Year	Company	Model	Number of parameters
2018	OpenAI*	GPT	110m
	Google	BERT-Large	340m
2019	Facebook	RoBERTa	355m
	OpenAI	XLM	665m
		GPT-2	1.5bn
	NVIDIA	MegatronLM	8.3bn
	Google	T5-11B	11bn
2020	Microsoft	Turing-NLG	17bn
	OpenAI	GPT-3	175bn
2021	Beijing Academy of Artificial Intelligence	Wu Dao 1.0	2.6bn
	Google	Switch C	1.6tn
	Beijing Academy of Artificial Intelligence	Wu Dao 2.0	1.75tn

Source: Edison Investment Research. Note: *OpenAI was founded in 2015 as a non-profit, with a $1bn pledge from a group of founders including Elon Musk; in 2019 it became a for-profit and Microsoft invested $1bn.

AN INTRODUCTION TO AI

However, both the financial and environmental costs can grow exponentially when implementing more parameters to achieve linear gains in performance. A number of the companies in Exhibit 3 are therefore seeking to reduce the number of parameters, while still achieving similar levels of performance. Google's DistilBERT, for example, has only 66m parameters, while retaining 97% of BERT-Large's NLP capabilities and is up to 60% faster.

Computer vision is another area of deep learning technology that has benefited from competition. ImageNet is a database of 14m images that have been labelled according to the WordNet hierarchy. It is available for free for researchers for non-commercial use. ImageNet's Large Scale Visual Recognition Challenge (ILSVRC), an annual event that ran from 2010 to 2018, challenged researchers to develop models that improved the ability of machines to classify images. This resulted in the top-5 error rates of winning image classification models declining from c 25% in 2010 to around 3% in 2018. Top-5 error rates refer to the probability that any of the model's top five predictions matches the answer. In comparison, a 2017 study by Arizona State University tested a human's performance on the ImageNet Dataset and found that performance was restricted to a top-5 error rate of 5.1%. The same study, however, also indicated that a deep neural network's performance was only superior for high-quality images and was still materially lower than a human's performance on distorted images.

Pioneers now: Cloud service providers (CSPs) plus others

AI's integration with cloud computing is one of the most notable developments in AI over the last 10 years and has been a key investment area for many of world's leading tech firms. The large cloud service providers (eg Google Cloud, Microsoft Azure, Amazon AWS, Alibaba Cloud) have invested heavily in developing their own AI expertise, and this has resulted in the development of AI tools and services that can be used with their cloud services. Platforms allow users to either build models from scratch or implement pre-existing architectures from other frameworks, such as Facebook's PyTorch or Google's TensorFlow. These are the two leading frameworks and can be used across a range of machine learning (ML) tasks, but have a particular focus on the training and inference of deep learning neural networks. The CSPs may also offer applications for specific use cases, with details provided in Exhibit 4.

2 WordNet: a large lexical database developed by Princeton University, originally developed in English but now available in more than 200 languages.

AN INTRODUCTION TO AI

Exhibit 4: AI technology available from cloud service providers

Company	Platform	Model optimising solutions	NLP	Bots and virtual agents	Text analytics	Translation	Speech to text/ text to speech	Computer vision	Intelligent search	Consumer personalisation	Business forecasting	Fraud prevention	Specific vertical applications
Amazon	SageMaker	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Google	VertexAI	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✓
Microsoft	Azure	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗
IBM	Watson	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Alibaba	Cloud Intelligence Brain	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	✗	✓
Tencent	Tencent Cloud	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✓	✗

Source: Edison Investment Research

AN INTRODUCTION TO AI

The availability of these tools and easy access to on-demand cloud processing and storage has allowed a broad spectrum of companies to experiment with AI technologies without the need for initial significant investment in software or hardware. This gains importance when considering the rising number of companies that are accelerating the development of their digital infrastructures following COVID-19. Additionally, training AI models has historically been slow due to the amount of computational power required, providing scope for hardware technologies to build more efficient computers that can better cope with the demands of AI. NVIDIA, Graphcore and Intel create processors that can drastically speed up AI model training. For example, Graphcore's Intelligence Processing Units were able to train Google's BERT-Large two and half times faster than comparable processors.

Active M&A and venture capital market

A key element of the platform development for the US CSPs has been the acquisitions of both innovative AI start-ups and more mature AI-focused software companies. Interestingly, these CSPs have exhibited similar patterns over the past six years. Many of the earlier acquisitions in this period were focused on integrating the CSP's leading AI capabilities into innovative solutions, which aimed to tackle antiquated approaches in areas such as cyber security or data analysis, as well as evolve rapidly growing areas like IoT. More recently, companies developing new AI architectures are being acquired to evolve the CSP's in-house capabilities, ranging from improving democratisation to driving efficiency and effectiveness. See Exhibit 9 in the Appendix for further detail on acquisitions.

Alphabet has been the most active in this space, making close to 20 AI-related acquisitions for a total cost of at least c US$7bn since 2014 (however, this could be significantly higher as details regarding costs are rarely disclosed). Notably, Nest Labs – an AI Internet of Things (IoT) company – has been its largest and was bought in 2014 for US$3.2bn, illustrating Google's intent to improve automation in the smart-home market. DeepMind, acquired in 2014 for US$500m, is another of its most significant acquisitions and has been widely covered in the press for both its solving of the historic protein folding problem and key use in the Drugs for Neglected Diseases Initiative. Most recently, in May 2021, Alphabet acquired Provino Technologies, a start-up developing network-on-chip (NoC) systems, for an undisclosed amount; its hardware technology should help drive efficiency and lower power usage of Google's AI models. Microsoft's acquisition of Nuance Communications for US$19.7bn in April 2021 was its second largest in history, highlighting its ambition to grow its Microsoft Cloud for Healthcare product, which should benefit from the growing tailwinds of the telehealth market. Microsoft also made five AI-related acquisitions in 2018, the highest out of all the major CSPs, all of which were either focused around simplifying and democratising Azure's AI capabilities or in NLP/virtual assistant technology to improve its Cortana product. In June 2021, Amazon filed an 8-K form revealing its intentions to acquire 20% of autonomous trucking company Plus in a deal worth up to US$150m. Apple publicly reported 25 AI-related acquisitions between 2016 and 2020, the greatest recorded figure globally.

AN INTRODUCTION TO AI

A number of high-profile venture capital firms have also been active in this space, with US$70bn invested in AI in 2019 according to Stanford's AI index. Softbank's US$108bn Vision Fund 2 provides a notable example of venture capital activity, with the majority of its funds earmarked for AI investments. Despite a slow start for the fund due to the overall group's declining performance, it has been actively leading sizeable funding rounds of AI start-ups in recent months:

- In June 2021, Eightfold AI, a deep learning talent platform, raised US$220m in a Series E funding round, led by Vision Fund 2. Since November 2020, the company's valuation has doubled to US$2.1bn.
- In April 2021, retail-focused computer vision company Trax raised US$640m in its latest funding round, led by Vision Fund 2 and Blackrock.
- In April 2021, Indian social commerce start-up Meesho raised US$300m in its Series E funding round led by Vision Fund 2, leading to a $2.1bn valuation.
- In April 2021, Franklin Templeton and Vision Fund 2 provided an extra US$210m of funding to OneTrust's Series C round. OneTrust is a privacy, security, and governance platform, which raised at total of US$510m in the round at a US$5.5bn valuation.

Also notable is NVIDIA's proposed acquisition of Arm from SoftBank for a total consideration of US$40bn, positioning NVIDIA as a clear market leader in graphics processing units (GPUs) and AI hardware technology. SoftBank originally bought Arm in 2016 for US$31bn in order to expand its presence in the IoT market.

What can AI be used for?

AI has multiple applications. We describe below the main areas in which it is gaining adoption:

Pattern recognition

- **Cyber security** continually learns from live data to identify patterns of normal behaviour, making it easier to detect anomalies and breaches. This enables it to identify new attack vectors unlike standard methods that rely on historical examples.
- **Fraud** is becoming more frequent and costly, while changing behaviours can make fraud harder to identify when using standard software. In 2020, an average of 3m phishing emails were sent every minute and business losses stemming from data breaches reached $3tn in 2020, which is expected to rise to $5tn by 2024. AI-enhanced software can more effectively learn from larger volumes of real-time data compared to more traditional products.
- **Recommendations and personalisation**. By analysing the behaviour of groups of similar people, AI can be used to provide personalised information or recommendations. For example, Netflix suggests programmes that might be of interest based on your viewing history, Spotify produces a 'release radar' of new songs that are of a similar style to music you have previously listened to, and Amazon suggests products you may want to buy based on the products in your basket.

AN INTRODUCTION TO AI

Natural language processing

- **Chatbots** simulate human conversations through voice or text commands, with well-known examples including Amazon's Alexa and Apple's Siri. Developing emotional intelligence is the latest development in chatbot AI, acknowledging that users favour human-like interaction. Huawei's emotionally interactive voice assistant is a notable example.
- **Language translation** has been significantly improved in recent years, with the latest developments stemming from contextual language understanding.
- **Autocomplete** allows for effective brainstorming, where pre-written prompts can be used as inputs and networks then use previous experiences to produce new ideas. This will be increasingly utilised in financial services, insurance underwriting and law, where models can process large volumes of text as an input and then provide a summary of the key information or standardise inconsistent narratives, for example insurance claims.
- **Search engines** are implementing deep learning to better understand the meaning of search requests, resulting in more relevant search results.
- **Information retrieval**: AI can be used to scan through vast amounts of text to extract key and relevant information, which can be done far faster and more accurately than any human. It can also be used to compare documents to identify whether text has changed between the two.

Discovery

- **Healthcare**: deep learning catalysed the development of COVID-19 vaccinations, reducing the development time to a matter of months rather than years and is currently being used to accelerate drug discovery times.
- **Science**: AlphaFold, a programme developed by Google's DeepMind, highlighted AI's ability to solve some of science's greatest challenges when it determined a protein's 3D shape from its amino-acid sequence. Elsewhere, AI is an effective tool for detecting the probabilities of extreme weather events and can be used in space exploration.

Object recognition

- **Facial recognition** software relies on AI's ability to learn the specific geometries of a person's face and is now used by several companies, with examples including Apple's iPhone and Facebook's DeepFace programme.
- In **healthcare**, computer vision is used for more accurate detection of early stages of certain cancers, where, for example, a machine developed by the Center for Research in Computer Vision was able to detect tiny specs of cancer at a 95% accuracy versus 65% for typical radiologists. It can also be used to compensate for shortages of expert radiologists or ophthalmologists, highlighting scans that could require further consideration by humans.
- **Autonomous vehicles** rely on a range of deep learning technologies, but it will be a long time until this technology is fully commercially available. For full automation, vehicles will have to be able to comprehend a large number of different sensory inputs, including radar, lidar and computer vision, and take action in real-time while in some cases having to consider competing priorities, for example in the event of a potential collision with another car.

AI IN THE ENTERPRISE

AI IN THE ENTERPRISE

AI is used in a wide variety of areas including consumer products and services, enterprise software and services, industrial automation, healthcare and transportation. In this report, we focus on the use of AI in the enterprise, exploring how companies develop AI-enhanced products and services. At the heart of this is software that has been developed to take advantage of AI technology. As with most major technology trends (the shift from mainframe to PC, internet, mobile, cloud), an ecosystem builds up that serves the needs of users: infrastructure (hardware and software), development platforms, tools and vertical specific applications. We have broadly grouped AI software into three categories:

- **Platforms:** provide the software to enable users to develop their own AI-based applications. Platforms are available from the cloud service providers (see Exhibit 4 on page 17 for more detail). Other companies providing platforms include IBM (Watson), expert.ai (for natural language processing), H20.ai and C3.ai.
- **AI-native applications:** software developed with AI at its core. It does not include many listed companies yet; as AI is a relatively young technology, many AI-native companies are still early-stage and privately held.
- **AI-enhanced applications:** software that includes AI technology to enhance performance. This category contains more mature software companies that have started to incorporate AI into their solutions where appropriate.

Companies can access tools from the cloud service providers or other software vendors and couple these with algorithms from open-source libraries to develop their own AI-powered applications. We believe that the availability of these tools is democratising the use of AI and we expect to see its adoption accelerate as companies experience the benefits. For those companies that do not have the in-house expertise or desire to develop applications themselves, there is a growing market of vertical-specific applications that provide enterprises with the advantages of AI.

Exhibit 5 below shows a non-exhaustive list of companies that are active in each of the three categories. We have included private companies, as particularly in the AI-native space, most companies are private.

AI IN THE ENTERPRISE

Exhibit 5: A selection of AI software providers by application

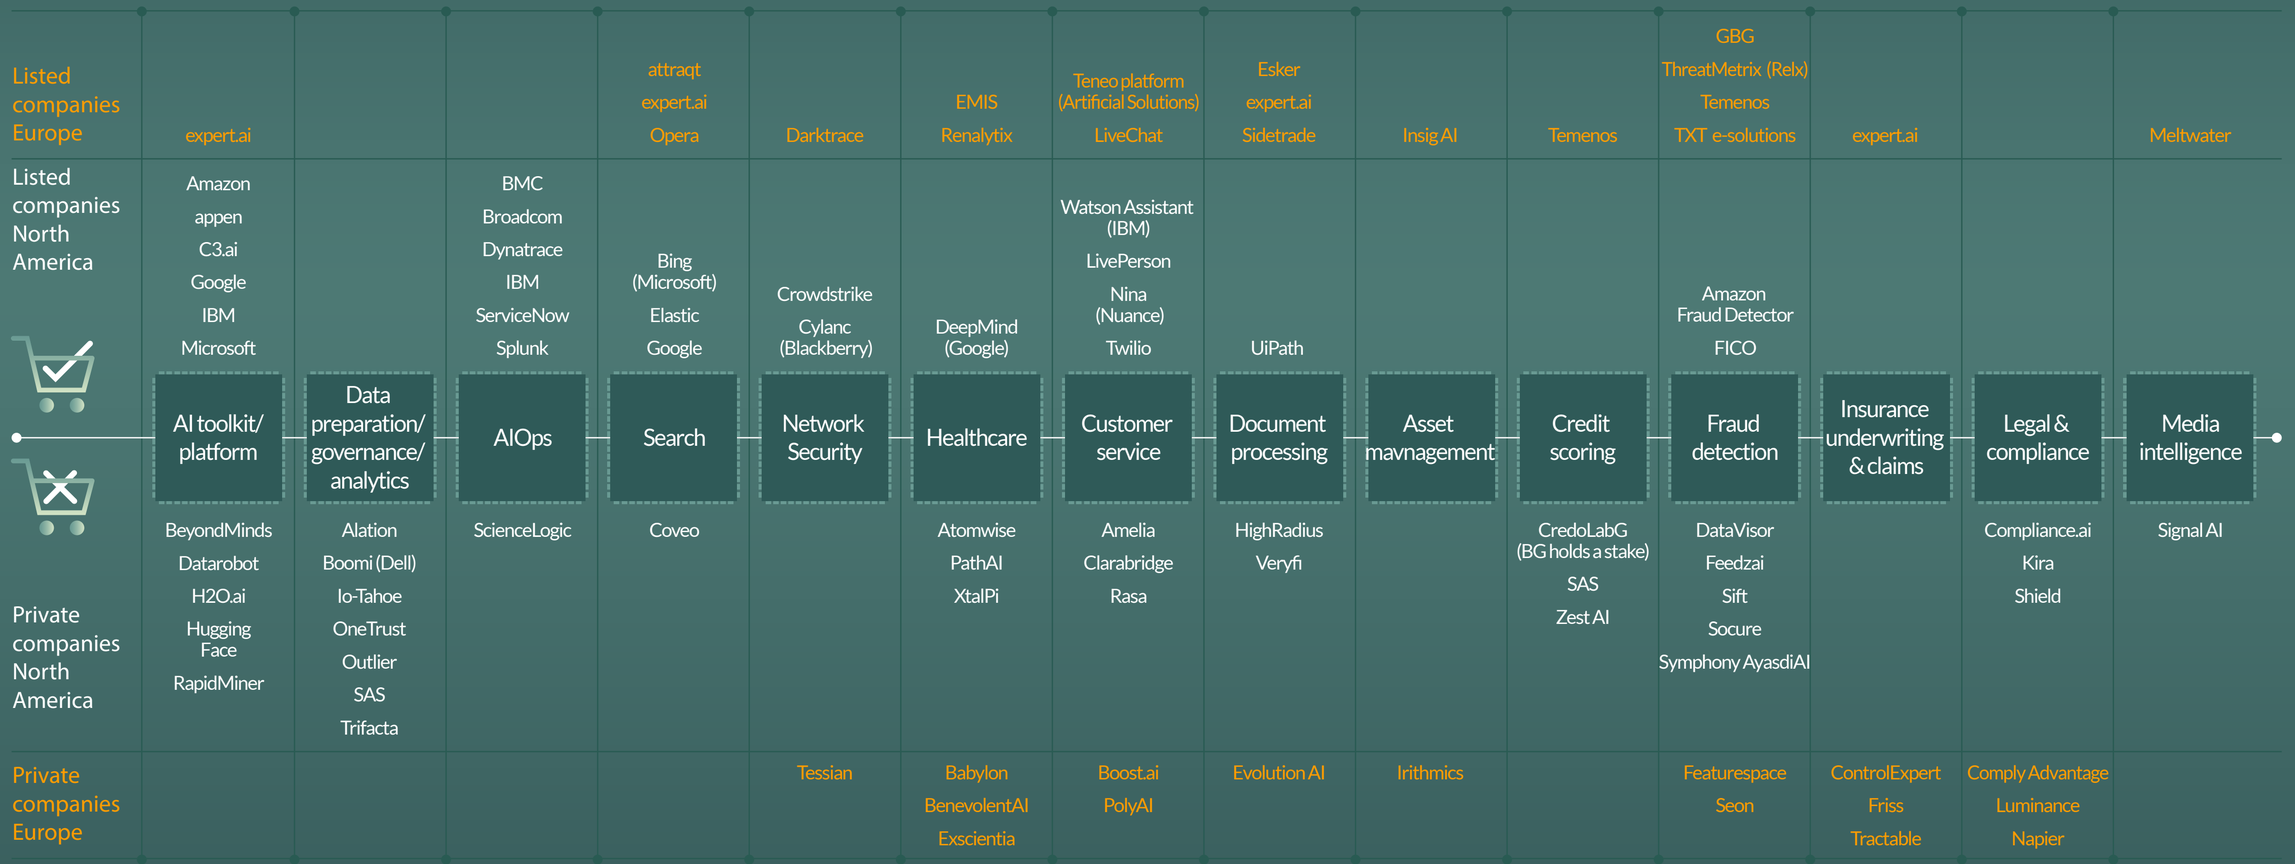

Source: Edison Investment Research

LEARNINGS

Following our meetings with the companies mentioned in this report,
we have identified several key points that management teams
should consider before implementing AI in their businesses.
We identify the growing list of areas that are well suited to using AI,
including those that generate high volumes of data
and those with processes that include time consuming,
recurring manual tasks that cannot be automated with standard software.
Following this, we indicate what companies need to do
if they are to implement AI effectively.
Additionally, we have analysed the potential challenges
to AI adoption and the impact it could have on the workforce.
We also include video interviews with eight companies
discussing these topics.

LEARNINGS

We have spoken to companies across the three groups described above and provide a summary of our findings below.

Where is it beneficial to have an AI offering?

Areas of a business that are well suited to using AI technology include those that generate high volumes of data and those with processes that include time-consuming, recurring manual tasks that cannot be automated with standard software.

The use of AI makes less sense where a business is not transactional, where there are limited historical data or where there is a high labour component that cannot be automated, for example consulting, caring or creative businesses. It is more complex to implement in highly regulated businesses – deep learning in particular is not always explainable so the technology may not meet regulatory requirements. It is also more difficult to use in critical processes where decisions can be the difference between life and death. In both cases, the models would need to be designed to ensure compliance and safety, with human involvement in final decisions. AI is not recommended for dealing with complex customer queries or problems or where there is a need to develop longer-term relationships with customers.

However, even if AI cannot add value to a company's products or services, it could still be used to improve a company's internal processes. For example, compliance AI software for highly regulated industries provides visibility of changes to regulation that specifically effect the company.

CLICK TO HEAR MORE ABOUT THIS

'Where is it most and least applicable?'

Is AI an opportunity or threat to a business?

Used correctly, AI has the potential to support better and/or new functionality for products and services and to improve a company's internal processes, giving the company a competitive edge. The diagram below shows the benefits a company can gain through the use of AI.

LEARNINGS

Exhibit 6: Benefits of using AI

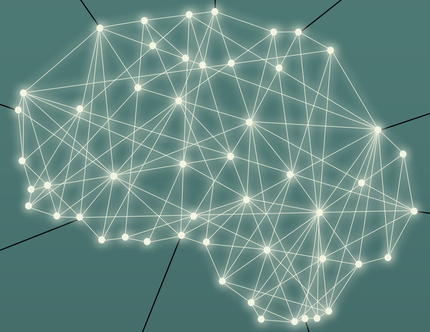

More accurate results
Eg better pricing of risk in insurance underwriting, fewer false positives in fraud detection, identification of cyber threats

More accurate search functionality
Employees can find information faster, saving time; customers can find relevant products more easily, improving conversion

Faster processing
Eg claims handling, routine customer service tasks

Better use of in-house knowledge
Eg identify trends in data that are too complex or time consuming for a human to process; make predictions; improve forecasting; identify patients at risk of certain diseases earlier, allowing preventative treatment

Less manual intervention in processes
Eg call routing, email routing, document recognition

Standardise processes
Take away scope for variability in how different people record or define things

Gain insight into data-heavy processes
Eg accessing and making sense of unstructured data held in different documents across a business

Better decision making
Through improved and faster insight into data

Better personalisation
Eg retail merchandising

CLICK TO HEAR MORE ABOUT THIS
'How does AI add value to your business?'

Source: Edison Investment Research

LEARNINGS

The main threat, or risk, to a business is that time and money could be wasted on unsuccessful AI projects. In certain verticals, the use of AI in software is becoming standard; those not using it run the risk of being left behind. And in the longer term, like many disruptive technologies, we expect the use of AI to become more widespread, posing a threat to those who are slow to adopt it.

Who is well placed to implement AI effectively?

Companies with access to large volumes of data are in a good position to develop their own AI models. Those that already take data analytics seriously should have good infrastructure and processes in place for collecting, cleaning, analysing and storing data, so will be in an even better position. Companies that have already taken steps to integrate AI into their product development or services will clearly have a head start on the competition.

What do companies need to do?

- **Good data governance:** collect the right data, label them correctly at the time if possible, have systems that make it easy to find data. In the example above, healthcare data included in patient records tends to be coded in a standard fashion, in EMIS's case using SNOMED codes. This makes it more likely that the data being used are accurately labelled, which reduces the preparation time before using the neural network and should increase the accuracy of the output.
- **Lead with the application, not the model:** when using AI, the biggest challenge is not building the model, it is working out what you want it to do. Companies need to define what data they have, what they can do with it and what they want to know before building the model. The project should be driven by the desired outcomes, not the model.
- **Start small, learn what works:** companies can extend the use of AI to new areas or larger production environments. MAPFRE noted the company tends to start with a proof of concept, followed by a small pilot with a single or limited number of use cases, and if this is successful, the project is scaled up for the use case in question. Several companies said that once they had successfully developed an AI application in one part of the business, this triggered many more ideas for use in other parts of the business, often with more ideas than the company currently has capability to execute.
- **Allow sufficient time to get projects into production:** adopting AI is not necessarily suited to agile development processes used in other areas of software development. It can take time (at least six months) to ascertain whether a project will lead to a working solution, but once this has been proven, the project can be connected to the scrum teams to roll out into the product. The preparation phase (finding, cleaning and labelling the data) can be the most time consuming. Then the model needs to be trained and tested before it is moved into production. Even once it is in production, it will need to be supervised and tweaked as circumstances change.

LEARNINGS

- **Consider all stakeholders:** projects have the best chance of success when they include both data scientists and the people within the business who will be using the application. Cross-functional teams comprising people from all relevant aspects of the business are more likely to achieve positive outcomes from using AI.
- **Adaptable workforce:** media hype about AI taking over human jobs has created fear of the technology among the workforce. In reality, AI is the latest in a long line of technology developments that could ultimately make certain jobs obsolete but at the same time create demand for new or altered jobs. The use of AI to replace narrowly defined processes tends not to completely remove the need for human involvement, but instead shifts the job to one of overseeing the technology and dealing with more complex situations that the technology cannot cope with.

Challenges to adoption of AI

There are a number of challenges in using AI, which we discuss in more detail below.

Explainability crucial for critical applications
Deep learning is effectively a black box as it is not easy to trace how decisions are made by the model. This can have implications in areas that require traceability and auditability in the decision-making process, such as healthcare, insurance underwriting and hiring. TXT is the coordinator of a €6m EU project, XMANAI, which has 15 participants and is exploring how to improve the explainability of AI so that customers can trust it. The project started in October 2020 and lasts 42 months. Participating companies include leading manufacturers such as Ford, Whirlpool EMEA and CNH Industrial and leading research institutions such as the Fraunhofer Institute and Politecnico di Milano. The project is exploring how to turn deep learning from a black box to a glass box, so that it is possible to see what path a neural network took to come to its decision.

Ethics a consideration, particularly to reduce bias
AI algorithms are created by humans and usually trained on historic data. This creates the opportunity for bias to be built into AI applications. For example, Amazon started to use AI to analyse CVs as part of its hiring process. Unfortunately, because historically staff hired to undertake technical roles were predominantly male, the AI used this as one of the factors for success when reviewing CVs. This skewed the selection of people called to interview to perpetuate the same gender mix as had been hired historically. Once Amazon realised the technology had unwittingly introduced gender bias, it stopped using it as a hiring tool.

As AI uses large quantities of historic data, unless directly addressed, trained models are likely to perpetuate the trends of the past. AI models will also incorporate the views of the people constructing them, which may include unconscious bias. Some facial recognition models have struggled to correctly identify people with darker skin – it is unclear whether this is due to unconscious bias in the construction of the model or if the training data did not include a sufficient range of skin tones to learn accurately.

LEARNINGS

In some regions where it has been used by police forces, the inaccurate identification of black and Asian males as suspects (ie false positives) has exacerbated tensions between the police and ethnic minorities. Amazon recently permanently banned its face recognition technology being used by US police.

As well as the risks of bias, the application of AI technology needs to be considered carefully to ensure it is not being used in an unethical way. A worrying example of this is the use of facial recognition technology to identify Uighur people in China, either to detect content filmed and uploaded by them (for takedown) or to detain them. Alibaba Cloud, Huawei, Megvii and SenseTime have all been cited as developing facial recognition software that can target ethnic minorities, including Uighurs, and in some cases, automatically alert authorities, although Alibaba has since stopped its cloud division from doing this.

Regulators seek to address the risks

Regulation is gaining importance as the powerful capabilities and potential harms of AI are being recognised. The above example of the use of facial recognition technology provides one of the leading worries people have over their fundamental rights, not just in China, but globally. Recent regulation has therefore been focused on building consumer trust, while aiming to encourage ethical innovation and competitiveness.

Europe has been at the fore when it comes to AI regulation, with the EU's work having already influenced several international discussions. Key examples include the OECD's ethical principles for AI25 and the UN Secretary-General's 2020 Roadmap on Digital Cooperation. An example of the EU's approach can be seen in its White Paper on Artificial Intelligence, which provides an interesting approach to different risk levels. AI that is considered 'high-risk', either in its intended use or sector (eg healthcare or transport), will have a greater number of rules to adhere to. By proportioning regulation, it provides greater freedom for 'low risk' areas to accelerate AI innovation.

Data protection regulation, such as GDPR, is also a significant factor. Europe again leads in this area, especially relative to countries that have led AI development , such as the US and China. This is a potential reason as to why the tech companies in these regions have been able to accelerate AI development given their heightened access to a variety of data. This also highlights the clear potential for regulatory arbitrage without further international alignment. The UK, for example, could potentially exploit this to the detriment Europe given the greater scope it has with its own data protection regulation following Brexit.

Additionally, without international alignment, the costs of entering new regions could become excessive. As indicated by Darktrace in its prospectus, entering new markets like China or Russia could lead to an 'increase in the costs of providing the platform' to comply with regulations.

Power consumption an issue

The drastic improvement in the capabilities of AI technology has required exceptionally large computational power, necessitating substantial financial and energy resources. This is based on the idea that the accuracy of a model depends on the number of parameters and data processed, supported by the success of models such as GPT-3 and BERT-Large.

LEARNINGS

Looking at the environmental cost, research done by the University of Massachusetts showed that training one NLP model could emit more than 626,000 pounds of carbon dioxide, almost five times the amount that a car emits in its lifetime, due to the energy required to power the necessary hardware for weeks or months at a time. Further, tuning a model, an exhaustive trial and error method to optimise a network's design, can drastically increase both the environmental and financial costs for little performance benefit. A study by Google et al. showed that the costs of one training run were $1 per 1,000 parameters. Open AI's GPT-3, for example, contains 175bn parameters, indicating that the training costs for just one run could have exceeded $10m.

Green AI, a growing space, aims to improve the efficiency of deep neural networks whilst achieving the same results as standard AI (red AI), reducing potential financial and environmental costs. Hybrid models, which use both neural networks and symbolic AI, can also reduce costs.

Availability of skills

All of the companies we spoke to highlighted the difficulty in hiring sufficient numbers of data scientists; as they are in high demand, they can be very expensive. But just as importantly, most companies talked about the need for data scientists who can work with domain experts, or even be domain experts themselves. It is the combination of data science expertise and knowledge of the business that is most likely to result in successful AI projects. As one company described it, someone with big data skills, AI expertise and domain knowledge (including the ability to talk to all the relevant stakeholders) is essentially a 'unicorn' and very difficult to find. Hiring people with a consulting background is one way to deal with the third requirement. By putting together cross-functional teams, AI projects should be able to bring together all the necessary expertise for success.

Several companies talked about their work with local universities (for example, Esker with Université Claude Bernard Lyon 1, GB Group with Universitat de Barcelona, TXT with University of Pisa and Politecnico di Torino) where relationships with PhD students can be mutually beneficial: the company has access to up-and-coming talent and the students get access to data on which they can train and test their models. Others talked about retraining software engineers and hiring graduates straight from university to train them. TXT works on various EU projects that can provide a good opportunity for students to work on their theses.

LiveChat noted that although all of its R&D takes place in Poland, the shift to remote working during the COVID-19 pandemic has widened the potential pool of hires so location is no longer a restriction. As with other areas of software development, hiring off-shore teams can be a way to access skills; for example, Mirriad has a team in India to supplement its European R&D team.

CLICK TO HEAR MORE ABOUT THIS

'What will the increasing use of AI mean for the workforce in terms of business and technology skills?'

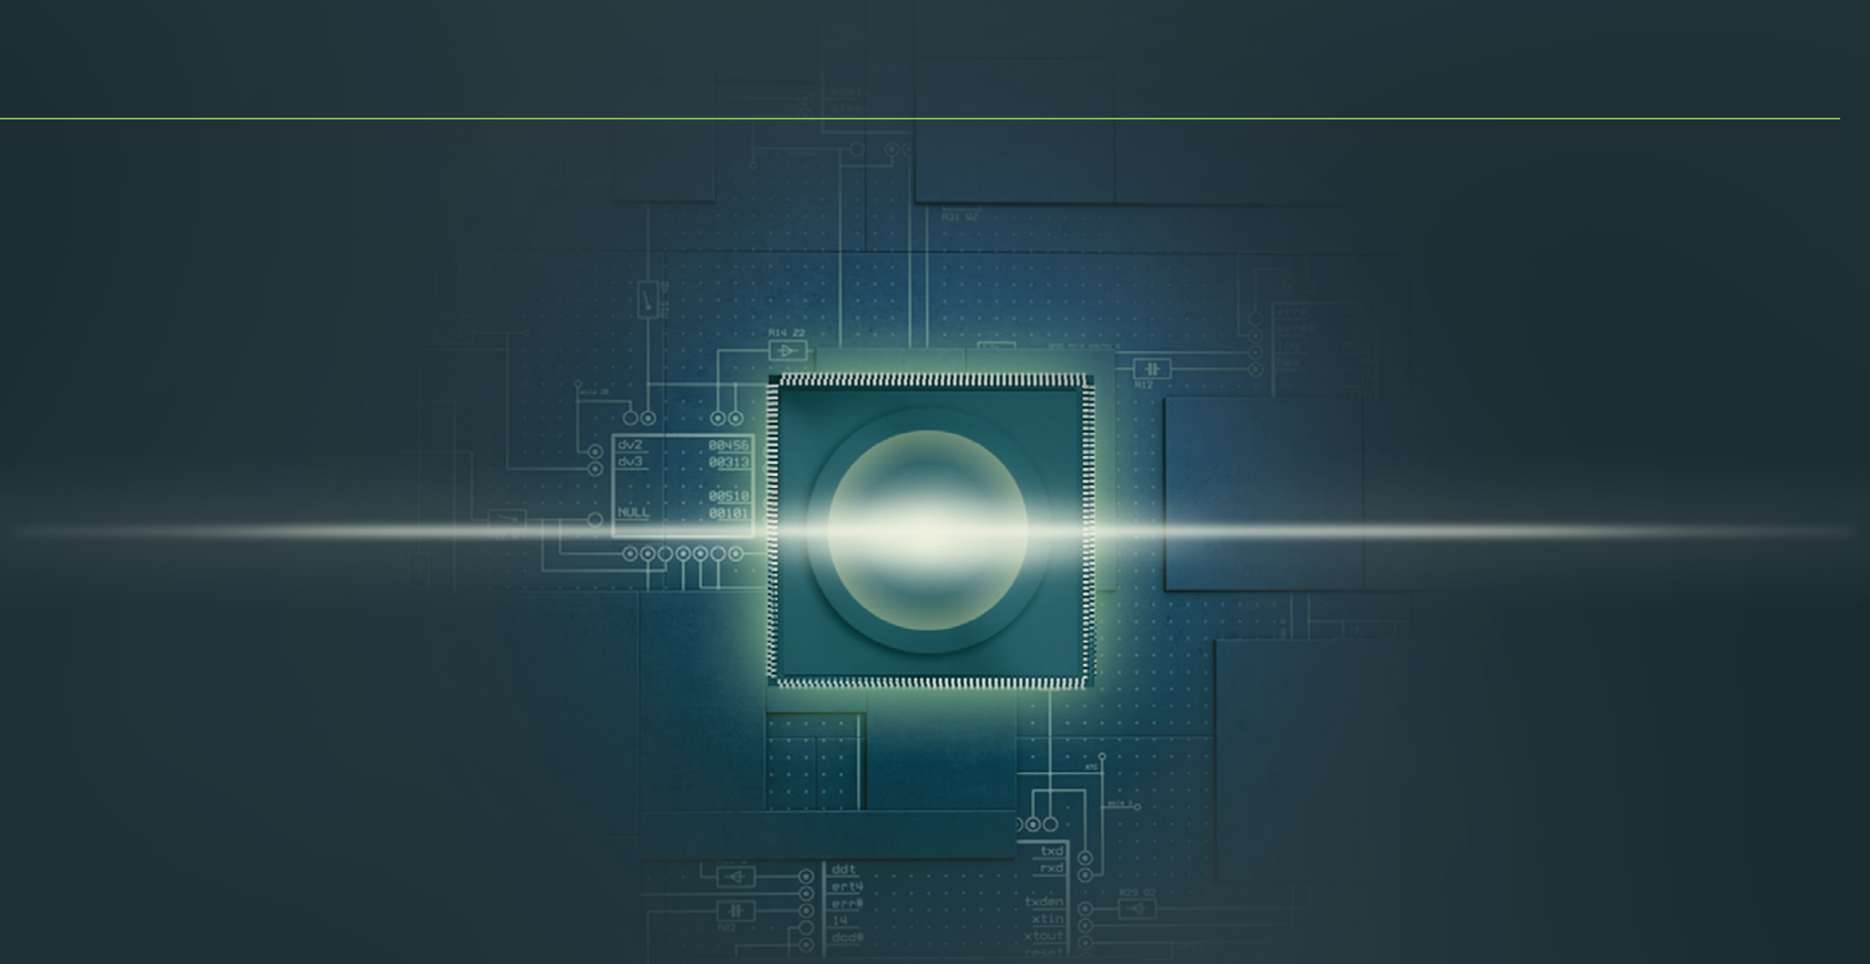

APPLICATIONS IN MORE DETAIL

We spoke to 11 companies actively developing and/or using AI technology
and describe in more detail a range of applications where AI can be used.
Verticals such as data services, financial services and insurance
have clear benefits to implementing ML technologies,
as they are typically industries characterised by high volumes of data.
Healthcare is an industry that has frequently appeared
in the news for its use of AI during the pandemic;
below, we show how AI can be used for identifying at-risk patients.
Additionally, we explore AI's innovative use in applications
for advertising, back-office processes, manufacturing and retail.
We discuss how natural language processing supports
multiple applications, with software available for specific use cases,
or for those keen to develop in-house applications,
as an end-to-end platform.

APPLICATIONS IN MORE DETAIL

Advertising

Mirriad Advertising (MIRI.L, MMDDF.PK) has developed proprietary AI over the past 10 years to support its in-content advertising solution. Its platform connects advertisers with content owners and distributors/broadcasters, identifying the most suitable locations within TV, film and music videos for relevant advertising to be displayed. Following this, Mirriad's technology dynamically inserts products and innovative signage formats into content post-production, for example overlaying an advert on the side of a bus or replacing a drink with another brand. Mirriad combines multiple narrow AI applications to achieve this. The platform works as follows:

- The ingestion process identifies brand and product insertion opportunities and achieves this by combining two different patented AI technologies in two layers. The analysis uses low resolution footage to boost speed/efficiency, while reducing cloud/networking resource loading. If the scene is selected for advertising, an API is used to download a high-quality version for production purposes.
- The first layer identifies physical insertion possibilities, with examples including surfaces, blank spaces and objects such as buildings or moving vehicles.
- The second layer identifies insertion opportunities in scenes, using contextual intelligence to narrow down the possibilities. It uses multiple algorithms to determine scenes that evoke a certain emotion or display a particular backdrop to match with specific brand categories or brands. Mirriad has undertaken research to show which emotions work best with certain product types. For example, insertions with cars work best when the scene is intense, while food and beverages work best with emotively sad scenes.
- The next stage is insertion, which uses native in-video advertising (NIVA) technology. Once the artwork has been created, the software modifies the scene and sends a high-resolution version to the content management platform. NIVA's key competitive advantage is that it can create multiple versions of the same scene, allowing the product or advert to fit the preferences of a geography, audience base or time of day. Mirriad refers to this as 'Versioning' with its customers. The use of AI to automate as much of both stages as possible is critical in reducing the otherwise heavy opex that is needed to scale the business.

The algorithm's ability to understand emotion, context and continuity across scenes means the process is highly automated, providing a scalable platform. Human intervention is only required to make sure the models are working correctly at each stage, with inventory reviewed before being made available to brand managers and content creators for approval before campaigns are activated.

Mirriad's AI strategy is to use existing neural network technology, adapted and continually trained to higher levels of accuracy. Notably, when asked if its business model could work without AI, management stated that it would require an unworkably large team of people just to analyse and record critical details, before taking into account the difficulty in consolidating the data for the standardisation of the process.

APPLICATIONS IN MORE DETAIL

Areas under development include live content and virtual reality and augmented reality. While at an advanced stage of development in both areas, each will naturally create new R&D challenges for the platform. Management noted plans for strategic M&A to help tackle these and other opportunities, with the end goal of creating a fully programmatic offering to serve Mirriad's US$135bn addressable market.

Document processing automation

Esker (ALESK.PA, ESKEF.PQ) is a developer of document processing automation software. Its software is used in the order-to-cash and procure-to-pay cycles to automate the receipt or sending of invoices and orders. The company aims to provide software that simplifies a customer's life, identifying where people spend their time on back-office processes and developing technology to automate the more mundane parts of the process. Its technology is designed to be invoice or sales order-agnostic, so that suppliers do not have to send their invoices in a customer-specific format and likewise customers do not have submit sales orders in a supplier-specific format. This complicates the processing of the documents for the recipient, and this is where Esker has developed software to help. When documents are received, often a human has to read the data and input it into the company's enterprise resource planning (ERP) system. Esker's software has been designed to automate this process as far as possible.

To extract information from documents, Esker has traditionally used auto-learning, a machine learning-based approach. Where it has seen an invoice from a particular supplier before, the software knows where to find the necessary data. When a new type of invoice is received (either from a new supplier or because an existing supplier has amended the structure of an invoice), it will present this to the user with an initial attempt at extracting the data included in the invoice; the user then corrects it if necessary and the feedback teaches the software how any subsequent invoices of this type are populated.

Started with invoice batch splitting

Esker initially experimented with deep learning (neural network) for splitting batches of invoices. The company took inspiration from a Google neural network that had previously been used to identify scenes in a film. When a customer receives paper-based or faxed invoices, they scan them and batch them together in a single pdf. A human is usually required to understand where the invoices from one supplier stop and those from a new supplier start, and where individual invoices start and stop. Esker trained a model to identify how to split these invoices. After successfully implementing deep learning for this application, the company decided to consider other uses within its software and decided that it would be useful to apply deep learning to the document recognition process to speed up the learning process for new or amended invoice and order types.

Now in use for data extraction

Using a long short-term memory (LSTM) algorithm Esker developed the Synergy neural network, training it on 10-years' worth of invoices.

APPLICATIONS IN MORE DETAIL

This currently has a 60% accuracy rate with unknown document types and reverts to auto-learning when the deep learning is not able to detect the data. Over time, the company expects this accuracy rate to increase. When the user of the software is happy that a certain invoice type is being correctly identified by the software and the data extracted correctly, they can switch to touchless mode so that invoices of this type are no longer presented to the user for verification. Clearly, using software that reduces the amount of human interaction while increasing data accuracy should improve the efficiency of back-office operations. It should also improve customer service as it is more likely that sales orders will be correctly recorded and having better oversight over purchase invoices can open up opportunities for companies to access early payment discounts from suppliers.

Sharing the benefits

Building on from the auto-recognition of invoices, Esker has developed a network that enables its customers to benefit from the learning of other customers. Invoice formats are shared in a global database across Esker's customer base. This is based on a version of facial recognition: it treats the invoice as an image, like a face, so it does not need to use optical character recognition (OCR) software.

Adapting for anomaly detection

Sales order processing is more difficult than invoice processing as every line in a sales order must be extracted, and the penalties for reading a sales order incorrectly are more severe. For example, reading a quantity incorrectly could lead to the company producing 100 units of something rather than 10, or alternatively, reading the product wrong could result in the manufacture of an unwanted product. Esker has developed an AI-based anomaly detection feature. This assesses whether the quantity or amount ordered makes sense based on recent history. If the system thinks it looks odd, it will highlight the amount in red for the user to check. This gives the user the opportunity to approve the order, or to double-check with the customer, so that the correct order is fed into the ERP system.

Future activity: Resource planning, invoice approval process, cost centre prediction

The company is looking at using AI in a number of other areas:

- For analytics and forecasting for order management. From analysing previous orders (say over the last two to three years), the system can detect peaks and troughs in demand and use these to predict when a customer might need to hire more staff to cope (or fewer). This is likely to be added as an additional feature on the dashboard rather than as a separate product.
- For the invoice approval process. By analysing every step in the process, it should be possible to identify where delays are being caused. This is important because late or failed payment can result in financial penalties, and in the worst case, the supplier might halt further shipments.
- To predict general ledger account and cost centres for an invoice. This information is not part of a supplier's invoice and is usually added manually by an accounts payable clerk. Esker's algorithm predicts this information based on previous manually entered information.

APPLICATIONS IN MORE DETAIL

It has achieved a 90% accuracy paving the way to a touchless service for invoices that do not have a related purchase order, for example electricity or phone bills.

Data services

As well as the fraud business described below, **GB Group (GBGP.L)** provides identity verification and location services and uses AI in both.

Identity verification
A person's identity can be checked either based on data attributes (eg appearance on electoral role, mobile phone number) or by document verification (eg checking passports or driving licences). GB Group uses AI in both areas. The US identity verification business operates a give/get commercial model. Customers give access to their data on an anonymised basis. Using machine learning, this data is analysed, and if GB Group spots something unusual, the customer is alerted, helping to reduce origination fraud. For document verification, the system needs to recognise the document (for example, is it a passport? If so, from which country?) and then extract the necessary information. GB Group does this using a combination of machine learning and open-source OCR technology. Once the data has been extracted, the system checks for digital tampering to ensure the document is genuine. Again, GB Group uses machine learning as part of this process. The company has the ability to use facial recognition technology but chooses not to, due to concern over the potential for built-in bias, instead using liveness testing to check that an image has not been faked.

Location intelligence
GB Group supplies address data on a global basis and aims to be able to provide a customer's address in five keystrokes or less (rather than by typing the post code). This is easier in some countries than others, as not all countries have structured addressing nor reliable data sources. While GB Group typically uses data from multiple sources, verified or enhanced manually if necessary, it also uses machine learning techniques (such as gradient boosting) to make poor data more reliable.

Financial services

As financial services is an industry that generates a high volume of transaction data, AI has been adopted for a number of different applications.

Fraud detection and compliance
Both **GB Group** and **TXT e-solutions (TXTS.MI)** use AI techniques in their fraud management solutions.

GB Group's fraud solutions have been designed to meet the needs of its global customers, which generate vast amounts of data in different languages and dialects, including different character sets such as Chinese and Cyrillic. These customers also operate in countries with different levels of maturity with respect to government schemes and standards. GB Group offers a rules-based fraud detection solution with a pluggable machine learning capability to create proprietary algorithms for different scenarios. This can increase detection accuracy by applying indicators of new fraud patterns to existing fraud detection rules. Customers tend to be larger banks that have the resources to use the machine learning engine.

APPLICATIONS IN MORE DETAIL

TXT e-solutions' Fintech division includes TXT Risk Solutions, which was acquired as a start-up in 2018. TXT Risk Solutions' main product is Faraday. The Faraday platform is used by financial institutions (FI) to ensure they are meeting compliance requirements, covering areas such as anti-money laundering (AML), countering terrorism financing (CTF) and corruption. Faraday analyses transactions using machine learning to understand which transactions are suspect. It works on the principle 'see the set and not only the single', using not only transaction data but also looking at the network around it to put the transaction in context. For example, the software can check whether the recipient of a transaction is related to someone who is a politically exposed person. The software makes use of FI customer data and publicly available databases from the government and the police.

Faraday will highlight a list of transactions that it has detected as potentially suspicious, leaving the FI to investigate each one and say whether it was suspicious or not, creating a feedback loop so the software continues to learn. Importantly, Faraday is used as a detection tool based on probability, leaving staff at the FI to make executive decisions. In some cases, the FI cannot prove a transaction is suspicious but may still believe it is; Faraday is able to use this information as part of the feedback loop.

Asset management

Insig AI (CTNA.L), which listed on AIM in May 2021, develops a range of machine learning based products for the asset management industry, with the aim of optimising investment performance by enhancing insight and analysis. The company originally started as a consultancy in 2017, born from the belief that machine learning technology was perfectly suited to address the problem that many asset managers are still dependant on antiquated technologies and outdated methodologies. However, after seeing how many of these underlying issues could be addressed by a core set of modular applications, management switched its focus into developing proprietary software in late 2019. This approach is far more scalable, providing accessibility to smaller asset managers that cannot afford the substantial costs of using a consultancy and risk falling behind if they do not implement AI technologies. The company now has five core products, which fall under three main areas: data, portfolio and ESG.

Data products are designed to transform and move both private company and third-party data into a cloud-based infrastructure by using data wrangling techniques. This is a necessary first step before AI can be fully optimised and serves as a feeder for the portfolio and environmental, social and governance (ESG) products. The three main products are detailed below:

- **Data** extracts third-party data from vendors such as Bloomberg, Capital IQ, etc, as well as from the client's private data (including unstructured data such as transcripts or annual reports), and stores it in a cloud-based infrastructure in a structured format. The ability to analyse both public and private data in tandem is a notable differentiator.
- **Exceleton** is a framework for transforming Excel data into a usable, structured format (ie Python code) that can be used with AI models. A common issue with excel is that the ability to circulate spread-sheets between users can quickly lead to multiple,

APPLICATIONS IN MORE DETAIL

inconsistent variations of the same document; this is not possible with a cloud-based infrastructure.
- **Docs** integrates leading OCR technologies (eg Google Vision and Amazon Textract) to extract data from a variety of file types. Typically these are pdf documents, where text extraction tends to be manual and an expensive use of human resource.

Its **Portfolio product** provides a highly visual element to an asset manager's portfolio, clearly illustrating its exposure to sectors, geographies and risks. When forming an investment strategy, asset managers can only consider a limited number of factors to achieve alpha. By implementing AI, the number of factors can increase tenfold, and the software will then list which of these factors are contributing most to growth. By providing a more comprehensive and easier to understand visual of the portfolio, decision-making accelerates, allowing asset managers significantly more time to generate innovative new ideas.

The **ESG product** is its latest development, due for release in Q321, and aims to detail the ESG case for a particular investment. Management has seen that an aggregated score from large data vendors such as MSCI or S&P is no longer enough to justify an ESG case, so the product addresses this problem by using NLP classifiers to surface evidence of disclosures and validate whether companies are living up to the stated goals and disclosure requirements. Further, some companies' ESG reporting can be too detailed and frequent for investors to fully appreciate; the product can therefore help cut through this noise.

Investor relations

Irithmics is a private company that uses AI to analyse stock market trading data to try to calculate what the market is anticipating for a given stock. The company uses neural networks, originally designed for epidemiology, to recognise patterns in the market. Data inputs are all from publicly available sources, updated every day, and include share registers, shareholder notifications and funds' performances. The output of the model predicts when behaviour is likely to change within a 90-day window (the forecast period can be longer but tends to become less accurate).

This information can be used in several ways:
- **Companies** can use analysis on market sensitivity to news and data to understand market movement, investor appetite and the impact of volatility. Irithmics provides data to the London Stock Exchange (LSE) for its Issuer Services. Companies listed on the LSE can take advantage of this information to help them to identify the best times, within regulatory constraints, to issue information to the market and to engage with investors.
- **Asset managers** can get enhanced portfolio insights into how the changing views, strategies and expectations of other market participants could affect their portfolio.
- **Financial PR professionals** can gain greater visibility over investor views, outlook and strategy, and can use this to guide clients with data-driven insights into behavioural changes and the impact of news and events.

APPLICATIONS IN MORE DETAIL

Healthcare

AI has multiple uses in healthcare, including image recognition (for identifying abnormalities on scans, useful for cancer diagnosis and also for ophthalmology), drug discovery (including predicting the properties of potential compounds and generating ideas for novel compounds) and furthering understanding of illnesses. During the pandemic, AI has been used to identify more quickly who is at higher risk of serious illness if they contract COVID-19, and to understand what treatments might be effective in reducing the severity of the disease.

We spoke to **EMIS (EMISG.L)**, a UK software supplier to primary and acute care. The company, working with a life sciences partner, is using deep learning to attempt to identify which patients are at risk of developing atrial fibrillation (AF). AF is a heart condition that causes an irregular and often abnormally fast heart rate, and if left untreated, can increase the risk of stroke and heart failure. EMIS has trained a neural network developed by a third party; this is able to analyse a vast amount of data and identify patterns, linking different conditions and factors that correlate with a patient developing AF. The model calculates the risk of the patient suffering from AF; over a certain threshold it triggers the recommendation that the patient is screened for AF. In practice, doctors have said they do not want to know the percentage score, just whether the patient should be screened or not. This is an example of 'human in the loop': the software produces a recommendation but it is for doctors to decide whether and how to screen patients. The company is planning to add this functionality to its existing GP software at no extra cost, building on the risk tools it has previously helped to develop (eg QRISK3, QCancer), which are based on linear rather than deep learning models. While the goal of this project is not to undertake research into the causes of AF, and in any case the AI is not able to predict causation, the identification of patterns that might not be spotted by human researchers could be of interest to researchers investigating the disease.

EMIS uses clinical data from patient records that have been anonymised. From a data governance perspective, EMIS and its partner have to obtain approval from the data controllers to use the data and must specify exactly what it is being used for. It cannot then be used for another application. It is possible that researchers with the correct permissions could use the model EMIS has developed to advance the science around AF.

As deep learning is effectively a black box, in approving the use of the output of the model the Medicines and Healthcare products Regulatory Agency wanted to know what training data had been used and what the model would be used for but did not seek to analyse the workings of the model.

Insurance

The insurance industry is characterised by high volumes of data, both structured and unstructured. AI is being adopted in multiple areas to improve efficiency in claims processing, to improve risk management in the underwriting process and to provide better customer service. We spoke to Joan Cuscó, the head of innovation at MAPFRE, to understand how the company uses AI within its business.

APPLICATIONS IN MORE DETAIL

MAPFRE (MAP.MC) is a Spain-headquartered global insurance group operating in over 100 countries and serving more than 30 million customers worldwide. The company has teams of data scientists in its main countries of operation as well as working with third-party consultancies and start-ups for specialist expertise. Areas in which it is using AI include:

- **Claims processing:** MAPFRE uses AI to speed up the process of dealing with car insurance claims, reducing the process from days to minutes. When a customer makes a claim on their car insurance, they are asked to upload information and photos of the damage to the mobile website. The system analyses the images and, using a neural network that has been trained on thousands of images, determines where damage has occurred, advises on whether it should be replaced or repaired, and calculates repair costs. MAPFRE also works with Shift Technology for AI-powered claims automation. The solution uses AI to instantly identify those claims that can be indemnified immediately, separating them from those that require further intervention on the part of insurance professionals. The technology reads the insured's policy and analyses documentation, images and other forms of structured and unstructured data to make the right decision about the claim at every point in the process and delivers a rapid response based on objective criteria. The result is a quantum leap in speed: client response times are now measured in seconds instead of days.
- **Underwriting**: MAPFRE partners with Kovrr, a private Israeli company, to better measure and price risk for its cybersecurity reinsurance product. Kovrr uses data analytics and machine learning to create customised cyber risk models for insurers and re-insurers. MAPFRE also works with ControlExpert to automate the process for customers to take out insurance on second-hand cars. The customer takes photos of the car and sends these to ControlExpert, where it can detect dents, scratches and small body defects, allowing a policy to be arranged in a matter of minutes. For home inspection, MAPFRE works with Flyreel, using AI to identify and document the unique risk of each property.
- **Business development:** the company uses AI to help salespeople to prioritise leads and commercial opportunities. An initial pilot that used historical data to calculate potential lifetime values for customers resulted in a 30% increase in sales. This has since been rolled out more widely.
- **Customer service**: the company receives more than three million emails a year. MAPFRE uses an AWS solution to automate three main use cases, reducing processing time from 1.7 days to 30 seconds, and improving the customer experience. The company also uses chatbots and voice virtual agents in its contact centres, where for basic interactions and certain sensitive topics, customers may prefer not to speak to a live agent.

Manufacturing

TXT e-solutions is an international end-to-end provider of IT consultancy, software services and solutions, supporting the digital transformation of customers' products and core processes. TXT serves several markets: its A&A division works with the aerospace, aviation, automotive, transportation and manufacturing sectors. Its Fintech division services the financial services market. In the A&A division, the company has a range of customer types, some already experienced in the use of AI, and others aware that it could be useful but without the expertise to apply the technology themselves.

APPLICATIONS IN MORE DETAIL

This is where TXT is able to step in and help customers to adapt AI technology to their specific needs. We discuss below how TXT uses AI in two areas of manufacturing:

- **Predictive maintenance:** TXT works with several manufacturing customers, helping them analyse IoT data to monitor factory performance and extract predictions. This has led to work with one manufacturer to analyse the best time to schedule maintenance of a particular part of a machine. Using machine learning to analyse the performance of the machine and the typical time to failure of certain components, combined with information from the manufacturing execution system (MES), the customer is able to schedule maintenance so that it causes the minimum a mount of disruption to the manufacturing process while avoiding the risk of machine breakdown.
- **Product design:** TXT worked with a customer on the design of a new product. TXT used machine learning to analyse the vibrations induced in the prototype to provide the designer with the information to define the best configuration. As well as helping the customer to design the product more efficiently, as the new product will need to be flown by a test pilot, it also helps reduce the risk in the test flight phase.

Natural language processing

Natural language processing enables companies to access and make sense of the unstructured data they hold. It can also be used to automate areas within customer services, such as providing call routing and basic processing (such as card activation) for inbound customer calls, providing chatbots in place of live customer agents and enhancing frequently asked questions by analysing live calls.

Expert.ai (EXAI.MI) has been developing AI-based natural language understanding and natural language processing (NLU/NLP) software for more than 20 years. Its NL platform offers a combination of symbolic AI and machine learning techniques to design, develop and deploy enterprise grade business applications leveraging data embedded in language. Symbolic AI is an AI technique based on high-level, 'human readable' representations of problems, logic and knowledge. The approach is based on the use of knowledge graphs, which contain for a given language, information about every word and its relationship to the other words in the language. The use of knowledge graphs means that the software is able to understand context and meaning, resulting in more accurate understanding of unstructured data, and the software can be trained via machine learning to add domain or customer-specific language. Expert.ai has developed vertical-specific applications covering multiple industries, with insurance, banking and finance and publishing and media the most commonly sold solutions.

In 2020, the company decided to launch an end-to-end NL SaaS platform, leveraging a set of different AI techniques to power language understanding in any application or process across any domain. The platform addresses the complete life cycle of NLU/NLP challenges including, for example:

- Reading, understanding and extracting meaningful data from text-based documents at speed and scale. For example, in the insurance industry this could include extracting the key information from policies, renewals, claims, medical reports and

APPLICATIONS IN MORE DETAIL

risk analysis reports. In pharma and life sciences, this could include reviewing scientific literature for key information.
- Detecting discrepancies between documents. This is useful for comparing different versions of the same document or spotting where contractual terms may vary across different documents.
- Categorising content so that it can be more easily discovered and navigated.
- Digital customer interaction including chatbots, email management and self-service systems.
- Helping to automate internal workflows by extracting data from complex documents to feed RPA bots.

The platform, launched in 2021, is simple to use so that any type of user, from business analyst to NLP expert, can design and develop their own NL-based applications in the fastest and easiest way without needing to be an expert in machine learning or linguistics. As well as including expert.ai's core NL capabilities, the platform integrates with proprietary or third-party machine-learning algorithms (and, in the future, open-source libraries) to provide a comprehensive service that allows users to work with their preferred tools. The platform supports users across the full process, from concept to full production, and enables them to monitor results to keep the application stable once in production.

The hybrid AI/NL platform can be used in different combinations to achieve the optimum outcome. For example, many enterprises do not have the volumes of data that are required to train machine learning models from scratch but do have sufficient data to augment expert's knowledge graphs, where only hundreds of documents are required to reach a high level of accuracy. The company highlights that its technology is explainable, providing full transparency into how it works, allowing users to see how results are arrived at, which is crucial in regulated industries and essential to bypass the 'black box' problem endemic in machine learning. By using knowledge graphs that have been developed over many years, rather than solely basing its technology on machine learning trained on high volumes of data, expert's technology has a lower carbon footprint.

Livechat (LVCP.WA) is a communication software provider listed on the Warsaw Stock Exchange. Its main product, LiveChat, is a SaaS solution that connects customers with live customer service agents in real-time. Several years ago, the brand started developing a chatbot that could provide automated responses to customers on chat. In 2018, the company released the first version of Its ChatBot platform. The solution now has over 2,000 users, compared to more than 32,000 customers of LiveChat software. The company has designed the chatbot to be simple to set up and to automate the work of customer service departments. The company developed the Visual Builder within ChatBot so that customers can build AI-based chatbots without having experience with AI and without needing to code. The customer can drag and drop conversational elements and test them in real time to design engaging chatbot stories. ChatBot integrates with Facebook Messenger, LiveChat, Slack, Zapier and WordPress. It provides ChatBot's Chat Widget that allows customers to add a chatbot to their website without coding.

APPLICATIONS IN MORE DETAIL

The ChatBot platform uses a combination of NLP and machine learning that has been developed in-house. The solution applies two matching systems to match the user's questions with pre-defined chatbot responses. Machine learning algorithms check entire sentences while keywords put a fine point on matching. ChatBot also saves all unmatched phrases; the customer can use the in-built training tool or the ChatBot API to train the model to recognise these phrases in the future. This way, the customer can improve their chatbot's performance over time and deliver a better user experience. The ChatBot team has started working on adding the possibility to use archives from its channel integrations, for example LiveChat. This could help customers to build chatbots faster and base their stories on the most frequent customer queries.

ChatBot contains ready-to-use system entities, which help extract information like emails, names or phone numbers from ongoing chats. If necessary, the customer can create its own user entities that are lists and synonyms of customer products, services, etc. The customer can later use those data when building the chatbot.

The ChatBot and LiveChat integration allows for transfer to a live agent if necessary. Additionally, the company is working on making ChatBot an omnichannel solution, so the end user can start the conversation on one channel and move to another mid-way through. The company recommends customers are selective about the areas of customer service that chatbots are used for, as they cannot replace all of the assistance provided by human agents. It suggests companies start small with repeatable tasks and, once these have been mastered, they can expand to more processes. LiveChat wants to progressively use AI technology in both its product suite and internally. The company has already integrated ChatBot with KnowledgeBase, another product from its suite. KnowledgeBase allows customers to offer self-service support for end users. The integration helps ChatBot predict the answers to customer questions based on KnowledgeBase's database. So far, the company has run the tests of the integration internally, sees value in combining the products and plans to develop this integration for commercial release.

ChatBot does not use voice technology. It has no plans to develop its own proprietary technology but has not ruled out using third-party software. For internal use, the company is using AI to analyse reasons for customer churn so it can manage it proactively.

Retail

The shift to online retail was well underway before COVID-19 but has been accelerated by the pandemic. Many online retailers have switched their focus from attracting potential customers to their websites to converting them into customers and retaining them. Consumers visit e-commerce sites for different reasons, for example, to find a specific item, browse for ideas within a certain category, or take advantage of a promotion or sale. Retailers need to have websites that work for each of these use cases, making it easy for customers to find the products they want, as well as presenting customers with products that they might be interested in. AIM-listed **attraqt (ATQT.L)** is a developer of software for e-commerce search, personalisation and merchandising, and is using AI to improve its offering. In fact, one of its strategic priorities is to 'infuse AI everywhere'.

APPLICATIONS IN MORE DETAIL

Historically, retailers have used a rules-based approach to personalisation and merchandising, ranking products according to chosen criteria, such as best sellers, highest margin, stock availability. However, as retailers expand channels, geographies and product catalogues, this does not scale efficiently. AI can be used to deal with the scaling problem, by analysing large quantities of data and finding patterns that no human could.

- **Personalisation:** the retailer has data on what products people are buying and can combine this with information about a consumer to treat them as an individual. The retailer is able to see what a consumer does once they enter their website, that is, their search, browsing and product selection history for each session. Using this, they can create a profile for that customer and check how often they visit (it typically takes three visits from a consumer before they buy). By looking at the behaviour of many users to detect patterns, the retailer can then combine this with the customer-specific data to bring back the most relevant results for that user and make product recommendations. attraqt offers three variants: prescriptive personalisation, product recommendations based on known information on the shopper; adaptive personalisation, based on real-time online behaviour of the shopper; and customised personalisation, based on the wisdom of the crowd. The retailer can also use the software for one-on-one marketing. For example, if integrated with the e-commerce platform, if a user abandons their basket, the retailer can send them an email afterwards to encourage the purchase.
- **Search:** attraqt uses AI to improve its search capabilities. The traditional search method directly matches words used in the search with the retailer's product catalogue. When a customer types words that are not in the catalogue, the search software needs to go through various processes, such as fixing typos or finding synonyms to try to find a match, and may still not return a useful result. Using AI, the search function can understand the relationships between words. It can take one to two weeks to train the algorithm (which can be done in the background), using product catalogue feeds and Google analytics data. The more datapoints fed into the algorithm pre-production, the more accurate it becomes. This process will need to be repeated for each language used. Once in production, the algorithm attempts to understand the intent of customers, learning from every interaction. So far, 13 customers have deployed attraqt's AI-enhanced search.
- **BYOA:** the company's view is that the most effective retailers combine expert merchandisers and data science teams. In some cases, retailers have developed their own algorithms or use third party algorithms. To remain relevant to these retailers, attraqt offers an algorithm orchestration (AO) solution that supports bring your own algorithm (BYOA). Using the AO solution, customers can connect attraqt/proprietary/third-party algorithms, review the performance of each and select the best ones to use.

HOW TO INVEST

HOW TO INVEST IN THE THEME

The largest players in the AI technology space are the US and Chinese mega-cap technology stocks Amazon, Alphabet, Apple, Microsoft, Alibaba and Tencent. The US companies have been active acquirers of a large number of innovative AI start-ups as well as more mature AI-focused software companies (see Exhibit 9 in the Appendix) and all invest heavily in AI-related R&D. They are using AI to improve their product offerings (eg Google using the BERT technology to improve its search functionality) as well as to attract customers to their cloud services via the availability of AI platforms and applications.

In the US, while the cloud service providers and other large technology companies offer AI software and tools (eg Apple, Facebook, IBM, Salesforce), there are very few listed AI-native companies. There is a plethora of privately held AI start-ups addressing specific vertical applications, many of which we expect to be acquired by larger established software companies to broaden their offerings.

Looking at the European market, investors also have limited access to listed AI-native stocks (see Exhibit 5). Companies include expert.ai, Darktrace, Insig AI and Mirriad.

It is possible to gain exposure to privately held AI companies via listed investment companies. For example, Draper Esprit holds a c 5% stake in Graphcore. However, we are not aware of any such companies that are solely invested in AI companies.

A number of privately held AI software companies are finding their way to the public markets via SPACs. Babylon Health, which provides telehealth services in the UK, is listing on Nasdaq via Alkuri Global Acquisition at a valuation of c $4.2bn. To date this has mainly been a US phenomenon, but with the UK's FCA reviewing the rules around SPACs, this may become a more popular way to list in the UK.

Another option is to consider companies that are enhancing their products or services with AI. If they are truly using AI techniques, rather than just using the term for marketing purposes, the scalability of the technology should give them a competitive advantage compared to peers who are not using the technology. Questions to ask company management to ascertain whether and how a business is using AI technology include:

- Where are you using AI within the business?
- What kind of AI are you using? Machine learning, neural networks, deep learning, symbolic?
- Where do you source your algorithms?
- Do you use any third-party tools, software applications or platforms?
- How well established is your data governance and cloud infrastructure? Do you need to improve this before you can fully exploit AI?
- How long did it take to train the model(s)?
- What oversight do you have of the AI models? How do you know they are providing you with accurate answers?
- How do you ensure bias is not built into the models?
- Do you have plans to expand the use of AI within the business?
- If you are not using it, why not?

CLICK TO HEAR MORE ABOUT THIS

'How can an investor tell if AI is real? Does it matter?'

APPENDIX

APPENDIX

What are neural networks?

Neural networks work like a human brain, where neurons (the circles in Exhibit 7) pass and interpret data between each other and predict an outcome based on weighted probability. The type of data used can vary greatly but are generally defined as structured or unstructured.

- Structured data is quantitative, clearly defined and formatted in a way so that it is easily searchable, with examples including dates, phone numbers etc.
- Unstructured data is qualitative, which requires more work to process and understand, with examples including text, audio and images.

Exhibit 7: Neural network diagram

Simple neural network

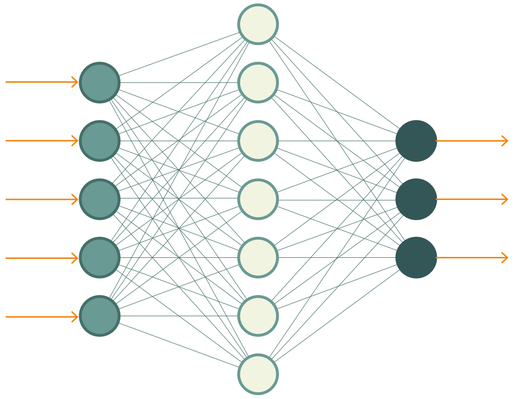

Deep learning neural network

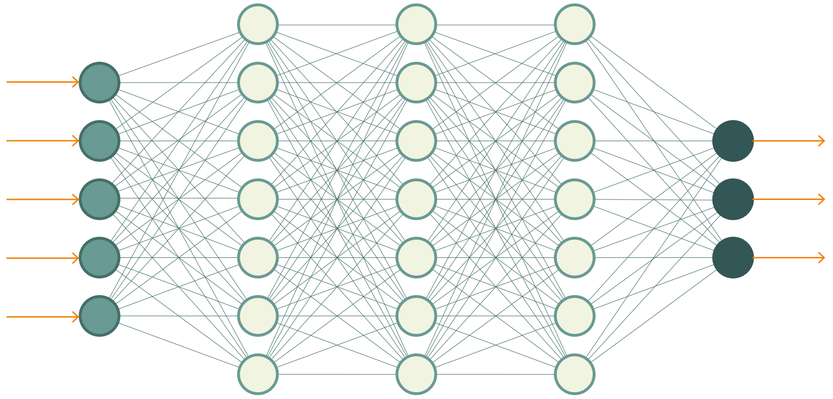

Source: Edison Investment Research

APPENDIX

Networks are split into three layers: input, hidden and output.

The input layer receives outside data, transforms it, and then parses it on to a network's hidden layer. Neurons in the hidden layer apply weights to the data and then directs it through an activation function, which becomes activated only when it reaches a specified threshold. Hidden layers are often referred to a black box as the decisions made by the model are largely unknown to the user. The model predicts the final value in the output layer based on weighted probability.

Simple networks contain three or fewer hidden layers and are normally used to analyse structured data, as it is quantitative and requires less work for a model to predict an outcome. Deep learning neural networks contain multiple hidden layers, allowing for the analysis of unstructured data. Each hidden layer function is specialised to produce a defined output; for example, functions that are used to identify human ears and noses can be used in conjunction with subsequent layers to identify a face in a picture.

Users can use a variety of techniques to train the model, which can be broken down into three main types:

- Supervised learning: specified input and output using labelled training data.
- Unsupervised learning: recognises patterns from previous experiences with no need to pre-label data.
- Reinforcement learning: trial and error with rewards and punishments.

In a process called backpropagation, the network will learn what weightings and measures to assign to each parameter until the user's desired results are achieved. An out-of-sample dataset will then be used by the user to validate and tweak the model.

APPENDIX

What is symbolic AI?

Symbolic AI works by converting any real-world object, such as numbers, letters, people, cars etc, into symbols and then uses logic to search for solutions. Symbols work together through relations, which can be either adjectives or verbs that describe how symbols interact together. For example, in the sentence 'Max drives a red car', 'Max' and 'car' would be the symbols of the sentence and 'drives' and 'red' would be the relation.

Symbols and relations are combined through logical connectives, 'and'/'or'/'not', which create propositions that a computer can use to test the validity of a statement using propositional logic, a type of Boolean logic. A few key concepts are highlighted below and Exhibit 8 demonstrates how logic is implemented.

- Symbols/relations: true values will be given a value of 1 and false values a value of 0.
- Logical connectives: 'And' values are treated as a multiplication and 'or'/'no' values treated as an addition.
- Proposition: 0 values are false, any value greater than 0 will be treated as true.

Exhibit 8: Examples of symbolic AI's use of truth tables

Symbol		Relation		Relation		Symbol		Proposition
Max		drives		red		car		Max drives a car and the car is red
1	x	1	x	1	x	1	=	1
	and		and		and			true
Max		drives		red		car		Max drives a car or the car is red
1	x	1	+	1	x	1	=	2
	and		or		and			true
Max		drives		blue		car		Max drives a car and the car is blue
1	x	1	x	0	x	1	=	0
	and		and		and			false
Max		drives		blue		car		Max drives a car or the car is not blue
1	x	1	x	0	+	1	=	1
	and		and		no			true

Source: Edison Investment Research

APPENDIX

If/then statements, known as implications, combine propositions together and are important for understanding how symbolic AI can begin to tackle real world problems. For example, 'if a credit card was reported stolen' and 'if it was used in an unusual location' then it is 'most likely a fraudulent transaction'. This is a very simple example but does illustrate one of the most common fraud types that can be tackled by AI.

We have only used examples using a couple of propositions so far; however, modern symbolic AI systems can process billions of implications (if/then statements) every second, indicating their transformative potential for solving real-word problems.

Knowledge bases contain a large array of propositions and are used for the logic of AI systems. That said, it would be impossible to populate a knowledge base with every single possible proposition, so programmes can mitigate this by solving logic problems. New propositions created by a model's solving of these problems are called inference.

Symbolic AI is now used in a range of industries, including supermarkets, banks and insurance, and are commonly known as 'expert systems' because they can replace experts like loan officers or insurance agents.

APPENDIX

Selected M&A by cloud service providers

Exhibit 9: Selection of AI M&A activity from the three major US CSPs

Year	Acquiror	Acquired company	Value (US$m)	Description
2014	Google	Deep Mind	3,200	Produces pre-programmed convolutional neural networks. Originally designed to play video games, latest models like Alphafold tackled problem of protein folding and finding new drugs for neglected diseases.
		Nest Labs	560	Acquired to create Google's home automation and IoT hardware, now its Google Nest product.
2015	Amazon	Sefaba Translation Solutions	N/A	Company automated the translation of different types of digital content and was bought to be integrated into the AWS cloud platform.
2016	Google	Eyefluence	N/A	Eye-tracking interface start-up to support Google's Daydream virtual reality platform, with team including ML expertise.
		Moodstocks	N/A	Paris-based start-up that developed ML technology to support image recognition features on smartphones.
	Microsoft	Genee	N/A	AI powered scheduling service to accelerate user intelligent experiences in Office 365.
		TouchType	C 250	London-based AI start-up that makes predictive keyboards for smartphones.
2017	Amazon	Body Labs	C 50-70	Uses AI, specifically computer vision, to create true-to-life 3D body models to support various b2b applications.
		Graphiq	C 50	Semantic search and data visualisation company that was acquired to enhance Alexa.
	Google	AIMatter	N/A	Photo editing app that uses AI to provide photo editing to selfie photos.
		Halli Labs	N/A	Young Indian AI start-up focused on building deep learning and machine learning systems to address what it describes as 'old problems'.
		Kaggle	N/A	Started in 2010 to offer ML competitions and now a public data platform that provides a workbench for data science and AI education.
	Microsoft	Hexadite	c $100m	Israeli cybersecurity start-up that uses AI to identify and protect against attacks.
		Maluuba	N/A	Maluuba's expertise in deep learning and reinforcement learning for question-answering and decision-making systems used to advance Microsoft's strategy to democratise AI.

Source: Refinitiv, Edison Investment Research

APPENDIX

Selected M&A by cloud service providers, cont.

Exhibit 9: Selection of AI M&A activity from the three major US CSPs

Year	Acquiror	Acquired company	Value (US$m)	Description
2018	Google	ABEJA	N/A	An investment, rather than acquisition, but notable as the Japanese AI start-up are collaborating with Google to create AI/ML solutions across various sectors, including retail and manufacturing.
	Microsoft	Bonsai	N/A	Bonsai has developed a novel approach using machine teaching that abstracts the low-level mechanics of machine learning, where technology will be implemented into 'brains' of Microsoft's AI programme.
		Lobe Artificial Intelligence	N/A	Offers visual tools that enable users to automatically build custom deep learning models and begin training.
		XOXCO	N/A	Known for its conversational AI and bot development capabilities, which will be used to advance Microsoft's conversational AI technology.
		Semantic Machines	N/A	Specialists in areas such as NLP and speech recognition, bought to enhance a number of products including Cortana.
2019	Amazon	Canvas Technology	N/A	Warehouse robotics start-up, bought to develop Amazon Robotics' portfolio of fulfilment centre machines.
	Google	Superpod	C 60	Automated Q&A engine, strengthens Google Assistant's ability to accurately answer questions.
2020	Amazon	Zoox	>1,200	Self-driving start-up, where technology will be used for Amazon's entry into the robo-taxi market.
	Google	AppSheet	N/A	No-code app development platform, supported by machine learning capabilities including computer vision and NLP.
	Microsoft	ADRM Software	N/A	Company provides large-scale industry data models, which help drive AI automation.
2021	Amazon	Plus	c 500	Acquired minority share (c 20%), indicating intended use of driverless trucking technology.
	Google	Provino Technologies	N/A	Acquired to enhance Google's network-on-chip capabilities and to increase speed of training machine learning models.
	Microsoft	Nuance Communications	19,700	Combining capabilities to improve cloud and AI offering, particularly for NLP in healthcare.

Source: Refinitiv, Edison Investment Research

APPENDIX

Exhibit 10: Companies in this report

	Ticker	Quoted currency	Market cap m	Market cap $m	Revenue CY	EBIT margin CY (%)	EV/Sales CY	P/E CY	AI platform	AI-native	AI-enhanced
attraqt	ATQT.L	GBp	86	119	23.6	N/A	3.4	N/A			×
EMIS	EMISG.L	GBp	817	1,135	1,64.2	23.6	4.7	24.3			×
Esker	ALESK.PA	EUR	1,559	1,850	129.7	12.4	11.8	114.1			×
expert.ai	EXAI.MI	EUR	142	169	33.1	-41.3	3.7	N/A	×		
GB Group	GBGP.L	GBp	1,709	2,376	211.9	22.7	8.0	44.5			×
Irithmics	Private	N/A								×	
Insig AI	CTNA.L	GBp	73	102	N/A	N/A	N/A	N/A		×	
LiveChat	LVCP.WA	PLN	3,313	862	216.7	61.5	15.0	N/A			×
MAPFRE	MAP.MC	EUR	5,356	6,358	21,174.3	5.8	0.5	7.6			×
Mirriad	MIRI.L	GBp	100	138	3.1	-400.0	20.7	N/A		×	
TXT e-solutions	TXTS.MI	EUR	104	124	86.4	9.1	1.1	16.9			×
Alphabet	GOOGL.OQ	USD	1,799,758	1,799,758	250,983	29.1	6.7	26.6	×		
Alibaba	BABA.N	USD	530,561	5,30,561	927,066	17.9	3.5	20.3	×		
Amazon	AMZN.OQ	USD	1,685,226	1,685,226	478,967	6.5	3.5	62.3	×		
Apple	AAPL.O	USD	2,411,090	2,411,090	366,672	29.7	6.7	26.1	×		
C3.ai	AI.N	USD	5,207	5,207	245	-47.6	16.8	N/A	×		
Facebook	FB.OQ	USD	1,004,568	1,004,568	118,421	40.2	8.0	25.8	×		
IBM	IBM.N	USD	126,345	126,345	75,129	16.5	2.3	13.1	×		
Microsoft	MSFT.O	USD	2,141,068	2,141,068	191,565	41.1	10.9	32.6	×		
Tencent	0700.HK	HKD	4,599,442	591,797	587,688	29.4	6.7	27.1	×		

Source: Edison Investment Research, Refinitiv (as at 2 August, 2021)

APPENDIX

Videos in this report

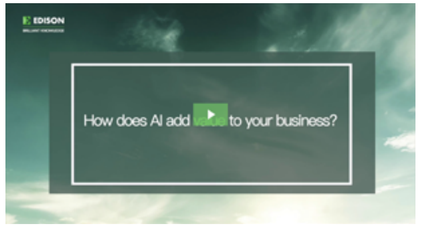

‘How does AI add value to your business?’

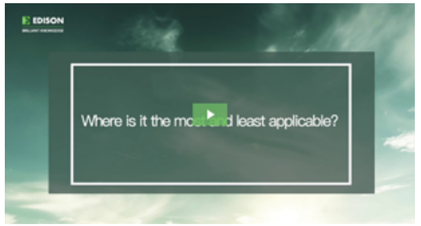

‘Where is it most and least applicable?’

‘What challenges have you faced when implementing an AI project?’

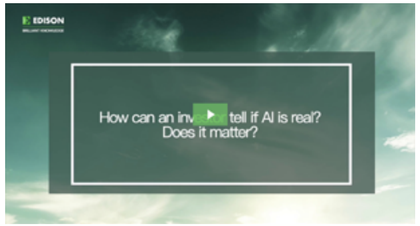

‘How can an investor tell if AI is real? Does it matter?’

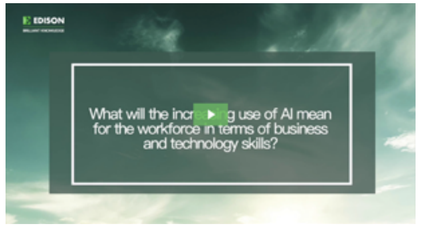

‘What will the increasing use of AI mean for the workforce in terms of business and technology skills?’

Companies interviewed:

EMIS	Irithmics
Esker	LiveChat
expert.ai	MAPFRE
Insig AI	Mirriad Advertising

EDISON

BRILLIANT KNOWLEDGE

REACH OUT

If you'd like to find out how Edison can help you:

Katherine Thompson
Director, TMT research
tmt@edisongroup.com

Max Hayes
TMT analyst
tmt@edisongroup.com

Amsterdam

Athens

Cape Town

Frankfurt

Hong Kong

London

Milan

New York

Paris

Singapore

Sydney

Tel Aviv

Warsaw

GENERAL DISCLAIMER AND COPYRIGHT

This report has been prepared and issued by Edison. Edison Investment Research standard fees are £49,500 pa for the production and broad dissemination of a detailed note (Outlook) following by regular (typically quarterly) update notes. Fees are paid upfront in cash without recourse. Edison may seek additional fees for the provision of roadshows and related IR services for the client but does not get remunerated for any investment banking services. We never take payment in stock, options or warrants for any of our services.

Accuracy of content: All information used in the publication of this report has been compiled from publicly available sources that are believed to be reliable, however we do not guarantee the accuracy or completeness of this report and have not sought for this information to be independently verified. Opinions contained in this report represent those of the research department of Edison at the time of publication. Forward-looking information or statements in this report contain information that is based on assumptions, forecasts of future results, estimates of amounts not yet determinable, and therefore involve known and unknown risks, uncertainties and other factors which may cause the actual results, performance or achievements of their subject matter to be materially different from current expectations.

Exclusion of Liability: To the fullest extent allowed by law, Edison shall not be liable for any direct, indirect or consequential losses, loss of profits, damages, costs or expenses incurred or suffered by you arising out or in connection with the access to, use of or reliance on any information contained on this note.

No personalised advice: The information that we provide should not be construed in any manner whatsoever as, personalised advice. Also, the information provided by us should not be construed by any subscriber or prospective subscriber as Edison's solicitation to effect, or attempt to effect, any transaction in a security. The securities described in the report may not be eligible for sale in all jurisdictions or to certain categories of investors.

Investment in securities mentioned: Edison has a restrictive policy relating to personal dealing and conflicts of interest. Edison Group does not conduct any investment business and, accordingly, does not itself hold any positions in the securities mentioned in this report. However, the respective directors, officers, employees and contractors of Edison may have a position in any or related securities mentioned in this report, subject to Edison's policies on personal dealing and conflicts of interest.

Australia

Edison Investment Research Pty Ltd (Edison AU) is the Australian subsidiary of Edison. Edison AU is a Corporate Authorised Representative (1252501) of Crown Wealth Group Pty Ltd who holds an Australian Financial Services Licence (Number: 494274). This research is issued in Australia by Edison AU and any access to it, is intended only for "wholesale clients" within the meaning of the Corporations Act 2001 of Australia. Any advice given by Edison AU is general advice only and does not take into account your personal circumstances, needs or objectives. You should, before acting on this advice, consider the appropriateness of the advice, having regard to your objectives, financial situation and needs. If our advice relates to the acquisition, or possible acquisition, of a particular financial product you should read any relevant Product Disclosure Statement or like instrument.

New Zealand

The research in this document is intended for New Zealand resident professional financial advisers or brokers (for use in their roles as financial advisers or brokers) and habitual investors who are "wholesale clients" for the purpose of the Financial Advisers Act 2008 (FAA) (as described in sections 5(c) (1)(a), (b) and (c) of the FAA). This is not a solicitation or inducement to buy, sell, subscribe, or underwrite any securities mentioned or in the topic of this document. For the purpose of the FAA, the content of this report is of a general nature, is intended as a source of general information only and is not intended to constitute a recommendation or opinion in relation to acquiring or disposing (including refraining from acquiring or disposing) of securities. The distribution of this document is not a "personalised service" and, to the extent that it contains any financial advice, is intended only as a "class service" provided by Edison within the meaning of the FAA (i.e. without taking into account the particular financial situation or goals of any person). As such, it should not be relied upon in making an investment decision.

United Kingdom

This document is prepared and provided by Edison for information purposes only and should not be construed as an offer or solicitation for investment in any securities mentioned or in the topic of this document. A marketing communication under FCA Rules, this document has not been prepared in accordance with the legal requirements designed to promote the independence of investment research and is not subject to any prohibition on dealing ahead of the dissemination of investment research.

This Communication is being distributed in the United Kingdom and is directed only at (i) persons having professional experience in matters relating to investments, i.e. investment professionals within the meaning of Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005, as amended (the "FPO") (ii) high net-worth companies, unincorporated associations or other bodies within the meaning of Article 49 of the FPO and (iii) persons to whom it is otherwise lawful to distribute it. The investment or investment activity to which this document relates is available only to such persons. It is not intended that this document be distributed or passed on, directly or indirectly, to any other class of persons and in any event and under no circumstances should persons of any other description rely on or act upon the contents of this document.

This Communication is being supplied to you solely for your information and may not be reproduced by, further distributed to or published in whole or in part by, any other person.

United States

Edison relies upon the "publishers' exclusion" from the definition of investment adviser under Section 202(a)(11) of the Investment Advisers Act of 1940 and corresponding state securities laws. This report is a bona fide publication of general and regular circulation offering impersonal investment-related advice, not tailored to a specific investment portfolio or the needs of current and/or prospective subscribers. As such, Edison does not offer or provide personal advice and the research provided is for informational purposes only. No mention of a particular security in this report constitutes a recommendation to buy, sell or hold that or any security, or that any particular security, portfolio of securities, transaction or investment strategy is suitable for any specific person.