THE INVESTMENT ASSOCIATION

In collaboration with:

EY
Building a better working world

CLIFFORD CHANCE

AI AND THE INVESTMENT MANAGEMENT INDUSTRY

November 2021

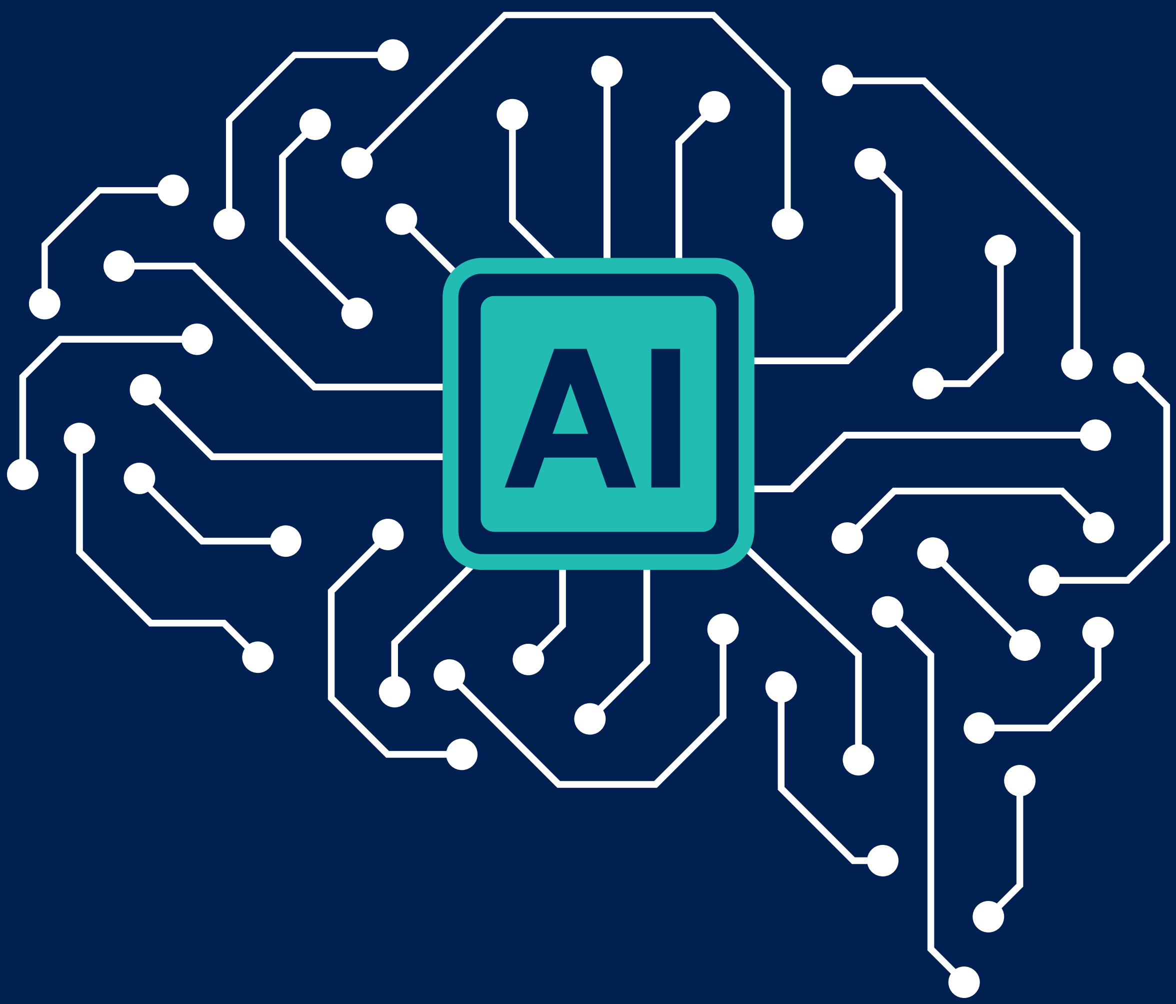

ABOUT THE INVESTMENT ASSOCIATION (IA):

The IA champions UK investment management, supporting British savers, investors and businesses. Our 270 members manage £9.4 trillion of assets and the investment management industry supports 114,000 jobs across the UK.

Our mission is to make investment better. Better for clients, so they achieve their financial goals. Better for companies, so they get the capital they need to grow. And better for the economy, so everyone prospers.

Our purpose is to ensure investment managers are in the best possible position to:

- Build people's resilience to financial adversity
- Help people achieve their financial aspirations
- Enable people to maintain a decent standard of living as they grow older
- Contribute to economic growth through the efficient allocation of capital.

The money our members manage is in a wide variety of investment vehicles including authorised investment funds, pension funds and stocks and shares ISAs.

The UK is the second largest investment management centre in the world, after the US and manages 37% of all assets managed in Europe.

ABOUT EY:

At EY, we are focused on building a stronger, fairer and more sustainable financial services industry. The strength of EY teams lies in the proven power of EY people and technology and the way they converge to reframe the future. This is how EY professionals are helping to build long-term value for financial services clients.
www.ey.com/ukfs

ABOUT CLIFFORD CHANCE:

Clifford Chance is one of the world's pre-eminent law firms with significant depth and range of resources across five continents. As a single, fully integrated global partnership, we pride ourselves on our approachable, collegial and team-based way of working. We always strive to exceed the expectations of our clients, which include corporates from all commercial and industrial sectors, governments, regulators, trade bodies and not-for-profit organisations. We provide them with the highest-quality advice and legal insight, which combines the firm's global standards with in-depth local expertise. For more information, please see **www.cliffordchance.com** and **www.linkedin.com/company/clifford-chance-llp**

CONTENTS

1. FOREWORD

AI IS A TECHNOLOGY WHOSE TIME HAS COME. A COMBINATION OF ABUNDANT DATA AND EVER MORE POWERFUL COMPUTING ARE PROVIDING FERTILE CONDITIONS FOR THE RAPID DEVELOPMENT AND IMPLEMENTATION OF AI ACROSS INVESTMENT MANAGEMENT AND THE FINANCIAL SERVICES INDUSTRY MORE BROADLY. THE PROMISE OF AI IS TO UNLOCK ENORMOUS ANALYTICAL AND PREDICTIVE POWER THAT CAN ENHANCE PROCESSES, CREATE EFFICIENCIES, AND IMPROVE INVESTMENT RETURNS. IT HAS THE POTENTIAL TO PLAY A ROLE IN NEARLY EVERY FACET OF THE INDUSTRY AND THE WIDER ECOSYSTEM.

Embracing disruptive technologies, such as AI, is a competitive imperative for firms. Firms must either disrupt themselves from within or be disrupted from the outside. At the same time, technologies must be adopted with care, with full regard for the risks involved and awareness of regulatory obligations, and where these are yet to be fully codified, regulatory expectations.

At the Investment Association (IA), we recognise the growing importance of technology as we all become more digitally aware. Technology, whether directly or indirectly, underpins much of how investment management firms operate. In 2018, the IA started its fintech hub and accelerator, the IA Engine, with the mission to fuel the adoption of technology within investment management. In 2020, the IA established its Technology Forum to consider the opportunities, challenges and use cases of new technologies, and help define the IA's policy objectives regarding technology.

The Technology Forum's first order of business was to identify the key tech trends which will shape the future the investment management industry. Top of the list was AI. The Forum subsequently embarked on a series of deep dives, facilitated by experts from EY and Clifford Chance, setting out to separate fact from fiction, demystify the topic and unearth insights that will be of value to industry professionals.

This report leverages the value of these sessions and disseminates the lessons learned for the benefit of the wider IA membership and other interested parties. It explores how investment managers are already using AI, how it is being developed in firms, the regulatory landscape and how firms can manage the many risks of both developing and using AI.

We hope the report will provide a useful and accessible overview of the key considerations surrounding the use and development of AI for firms.

For more information on technological and associated regulatory developments that affect the industry, please visit our dedicated technology webpage: https://www.theia.org/positions/technology

Finally, I would like to thank EY and Clifford Chance for contributing their time and expertise to this project.

Pauline Hawkes-Bunyan
Director, Business: Risk, Culture & Resilience
at The Investment Association

2. BACKGROUND

The industrial revolution heralded huge increases in productivity as for the first time human and animal power was augmented by machine power. We are now in the early days of a similar information revolution, as machine cognitive power via AI begins to augment human cognitive power, enabling massive transformation and innovation possibilities. As Jeff Bezos memorably said, "We are at the beginning of a golden age of AI ... we've only scratched the surface of what is possible." This promise is fuelling record investment into the AI sector, over US$20bn in Q2 2021 alone.[1]

Certainly, there is now widespread agreement on the potential benefits for investment managers. Eighty-two percent of investment management companies surveyed in 2020 listed AI to be of either 'very high' or 'high' strategic importance in two years.[2]

STRATEGIC IMPORTANCE OF AI IN INVESTMENT MANAGEMENT

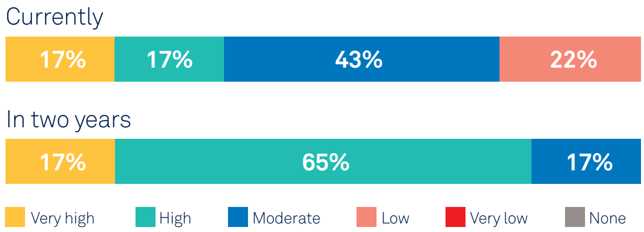

KEY ACTIONS

We have identified eight key actions that boards and senior managers should take in order to realise the potential benefits of AI whilst protecting themselves from the risks.

1. Incorporate AI into your innovation agenda

Many commentators think AI will be the single most transformative technology; firms need to ensure that it is incorporated in their innovation efforts.

2. Perform an AI landscape review to identify the use cases that would be most beneficial to you

There are many potential applications of AI through the value chain. It is necessary to analyse what is possible and what is most beneficial for your firm given your situation. Famously, Netflix applied an ensemble of AI techniques to improve the quality of its movie recommendations and increase the amount of time subscribers spent watching Netflix – its key KPI. What is most important to your firm?

3. Test via a proof of value

A proof of value significantly reduces the risk of wasted time and money by ensuring a focus on the quick delivery of business value whilst learning from the exercise.

4. Identify how best to leverage the broad ecosystem to deliver new AI-based capabilities

There is a very rich ecosystem that can potentially significantly accelerate your AI efforts whilst reducing cost and risk. In particular, there are a large number of fintechs of different types that have helped other investment managers; Other than fintechs, academia, service providers and established vendors can all potentially be leveraged.

5. Ensure the organisation has the right talent and skills to deploy and manage AI

With this understanding of how you can partner with the broad ecosystem, map the internal capabilities required and ensure you begin to build out the required internal talent.

6. Create a robust governance structure with a culture of transparent and ethical use of AI

Create a robust governance structure with a culture of transparent and ethical use of AI embedded in the organisational leadership, including the company's boards, general counsels, senior data, compliance, and risk and policy teams overseeing AI risk management.

7. Implement regular and frequent testing and monitoring of your AI solutions

Regular testing and monitoring of AI solutions, far beyond the development stage of the solution, should involve a wide range of stakeholders and span multiple areas across the business.

8. Implement regular audits of your AI estate

These regular audits should assess where and how AI is used in the organisation, AI functionality, the limits and boundaries set for the use of AI, and the apportionment of contractual liability, as well as ensure consistency with existing policies.

[1] "Artificial Intelligence In Numbers Q2 2021," CB Insights, July 2021.
[2] "Transforming Paradigms: A Global AI in Financial Services Survey," World Economic Forum, January 2020.

3. WHAT WE MEAN BY AI

There has long been a challenge to define AI; for this report, we will consider AI to be viewed as a set of technologies and techniques used to complement traditional human attributes, such as intelligence, analytical ability and other capabilities. AI, machine learning (ML) and modern data techniques have been greatly enabled by recent advances in computer processing, power and speed, and advances in AI depending, in turn, on advances in data techniques.

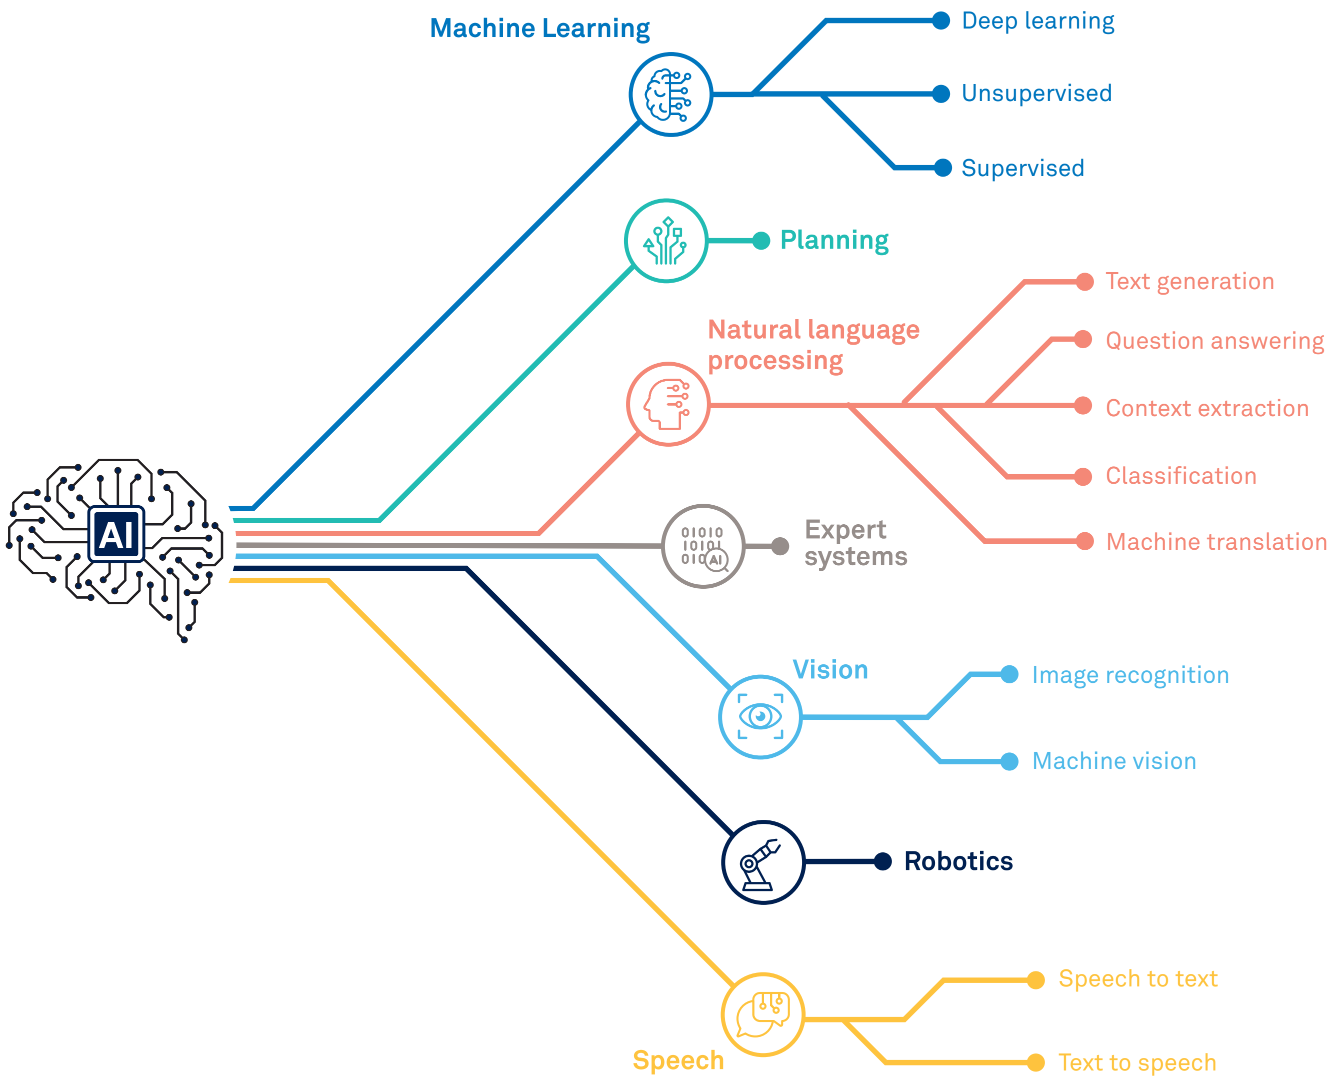

There is a broad spectrum of AI that focuses on a range of problem statements in every industry. Whilst there are many different techniques within AI, the most relevant one comes down to the problem that is trying to be solved. In this report, we will focus on the techniques listed below:

- Machine learning (ML): "ML is the science of getting computers to act without being explicitly programmed," says Andrew Ng (adjunct professor of Computer Science at Stanford University and acknowledged AI pioneer). Instead, the algorithm is 'trained' on a set of representative data called a 'training set.' This training set enables the algorithm to 'learn.' Larger sets of training data lead to better ML results. The advent of the internet and its vast stores of data have facilitated access to much larger datasets for training and, when combined with more powerful modern compute power, resulted in much more successful ML outcomes. The three main approaches within ML are:
  - Supervised learning: This subset of ML uses labelled data sets to try and classify data or make predictions, such as classifying spam from your normal email inbox, or for predicting when there may be a spike in inflation.
  - Unsupervised learning: In contrast to supervised, unsupervised learning uses unlabelled data and tries to draw patterns from the data itself. Essentially, it is made to find correlations and similarities in the data that you did not know you were looking for.
  - Deep learning: Deep learning is a subset of ML and involves the algorithm training and learning from experiences. A practical example could be virtual assistants. The successful games-playing algorithms (e.g., AlphaGo that beat the world Go champion) typically play themselves many, many times and improve by learning from those experiences.
- Natural language processing (NLP): This includes extraction and classification of natural language to understand the context of speech and text, so that emotion and state of mind can be inferred. One example is reading a client query about a trade and classifying it into a queue. Another one that has been used in many successful investment strategies is identifying CEO and CFO sentiments from earnings call transcripts.
- Document intelligence: This technique involves digitising, extracting and classifying machine-readable and image-based documents for processing using image recognition techniques and computer vision.
- Advanced analytics: This technique includes using a combination of ML and visualisation techniques to analyse complex data sets, e.g., multi-asset trade positions or holistic view of a client profile.
- Process automation: This technique involves a combination of tools used to automate (usually) middle- to back-office batch processes, using techniques such as robotic process automation (RPA) and workflow. It is typically applied on manual and repetitive processes.
- Communications mining: This technique includes using natural language models to understand unstructured data embedded in forms of organisation communications e.g., chats, emails, and customer relationship management (CRM) systems.
- Process mining: Using data science techniques, process mining consumes event data from across an organisation, usually in the underlying databases to provide action and insights on what people are really doing to address operational inefficiencies.

WHY HAS THERE BEEN A RAPID GROWTH IN THE PRACTICAL APPLICATION OF AI?

The exponential increase in the availability of data, computing power and storage — particularly cloud — and the increased sophistication of tools to analyse complex sets of data have led to the rapid growth of the use of AI within broad financial services, and investment management.

- Rise of data: Data is the new currency; new datasets continue to be formed from applications, outputs from business processes, social media, the internet and many more. It is estimated that as much as 90% of the data within financial services is unstructured which has always been challenging to analyse, but due to advancements in the way data is managed by organisations and the availability of data processing tools, this data can now be utilised to create competitive advantage.

- Developments in power, storage and cloud: The rise of cloud computing and compute power has allowed large-scale data processing and storage to be significantly cheaper, thereby reducing the barriers to entry for people to act as a provider or consumer of AI and data services. In turn, the flexibility these solutions offer has allowed organisations to innovate their business models at a greater pace and velocity. Investment managers are increasing their focus on leveraging cloud capabilities in response to the COVID-19 pandemic to compete in a remote world.

- Growth of AI providers, tools and techniques: There has been a significant rise in the number of new providers, tools, and methods of analysing data using ML, deep learning and artificial intelligence. This extraordinary growth has increased the speed, accuracy and volume of data that can be analysed, giving an investment manager access to innovative new capabilities. The market is also maturing and being embedded across enterprise systems or becoming more accessible to non-technical users fuelling sustained adoption.

The combination of these factors has provided the ideal conditions to tackle more diverse and complex problem statements in the investment management industry, as we will explain in section 4.

4. AI USE CASES THROUGH THE INVESTMENT MANAGEMENT INDUSTRY

Just like machine power has been added to almost every type of device (e.g., the electric toothbrush), in the future, machine cognitive power via AI will be applied to almost every process. Right now, we are still in the early stages of that progression but already, we have now seen well over forty different applications of AI through the value chain, all generating business. One CEO commented, "If computing power and data generation keep growing at the current rate, then ML could be involved in 99% of investment management in 25 years." Below, we highlight the characteristic goals of the AI use cases in different parts of the value chain as well as some illustrative examples.

THE INVESTMENT PROCESS

Information is power and since time immemorial, there has been a race for the power that information advantage can grant. Venetian merchants used Galileo's telescopes not to study the stars, but rather to study the cargoes of approaching ships a few hours before their competitors. Unsurprisingly, the most vigorous application of AI at investment managers is attempting to generate the holy grail of new insights for the investment process and hence improving the investment outcomes for clients.

There is now widespread agreement on the potential benefits in this area. For example, a significant majority of firms expect that AI will have either a 'very high' or 'high' long-term impact on investment returns.[3]

EXPECTED LONG-TERM IMPACT OF AI ON INVESTMENT RETURNS

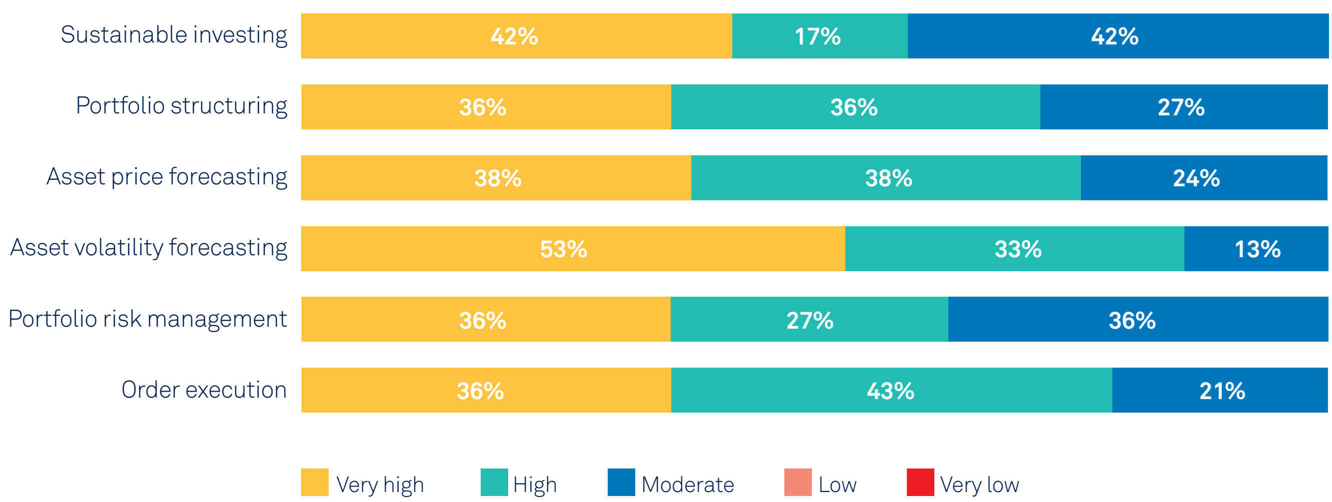

[3] "Transforming Paradigms: A Global AI in Financial Services Survey," World Economic Forum, January 2020.

Function: Research	Use case: Interpret alternative data sources	Problem statement: Is it possible to use different data sources to generate new investment insights?

Solution:
Over the last decade, buy-side leaders have begun to access new alternative data sources (most commonly web-scraped data, geo-location data, social media data, and credit card data)[4] and use AI or ML advanced analytics techniques to generate new insights from these complex, big data sources. This promise of improved investment performance is fuelling rapid expansion in the use of alternative data: the global alternative data market size is expected to expand at a compound annual growth rate (CAGR) of 58.5% from 2021 to 2028.[5]

Function: Dealing	Use case: Trading intelligence	Problem statement: How can we improve traditional approaches for forecasting transaction cost and liquidity?

Solution:
Leading companies in this area have publicly announced that they are generating trading intelligence from trading-related data points and combining these internal datasets with external datasets with the goal of 'lower transaction codes or being rewarded for providing market liquidity.'

SALES, DISTRIBUTION AND CLIENT EXPERIENCE

An improved client experience is imperative for investment managers, and in particular, augmenting their current human touch with advanced analytics. Leading firms are beginning to use AI to help with this key distribution challenge, for example: by enabling a richer and more productive interaction with clients; by helping tailor solutions for clients to win and grow business; and by identifying clients most likely to churn and the actions to retain them.

Function: Sales and marketing	Use case: Augmented relationship manager (ARM)	Problem statement: How can a relationship manager (RM) provide a more personalised treatment and experience to clients?

Solution:
AI-enabled digital interface for RMs is the solution. Delivers insights and recommendations to the RM in real time, based on the content of the client conversation, helping the RM to provide better client service and identify and progress new business conversations. The enforced changes to working patterns because of COVID-19, such as the usage of Zoom, Teams, etc. have lowered the barrier to adoption for ARM for both of the key constituents – the RM and the client.

Function: Sales and marketing	Use case: Churn prevention	Problem statement: How can we reduce client churn?

Solution:
Reduce redemptions by better client segmentation to identify those most at risk of churn. Another way is through predictive modelling and targeted interventions that are delivered through marketing automation software enhanced with real-time AI capabilities.

[4] "Alternative Data Primer," BofA Securities, October 2020.
[5] Alternative Data Market Size, Share & Trends Analysis Report By Data Type (Credit & Debit Card Transactions, Social & Sentiment Data, Mobile Application Usage), By Industry, By End User, By Region, And Segment Forecasts, 2021 – 2028," Grand View Research, Aug 2021.

RISK, COMPLIANCE AND OPERATIONS

Whilst Assets Under Management (AUM) across the industry continues to grow, many firms face challenges delivering scalable, and profitable growth. Leaders are beginning to apply AI to augment their current rule-based solutions around risk, compliance and operations, with the goal of better outcomes and improved efficiencies.

Function:	Use case:	Problem statement:
Risk and compliance	Monitoring suspicious transactions	How do we reduce the number of false positives and make our transaction monitoring more efficient?

Solution:
Automated detection of money laundering activity is still mainly performed by rule-based monitoring solutions, predominantly designed for monitoring transactional accounts in banks. These solutions generate high volumes of false-positive alerts, particularly outside of the banking domain. Supplementing and replacing elements of the incumbent rules-based solution with AI approaches has the ability to improve the efficiency of monitoring (reduced false positives) and improve effectiveness (more true positives).

Function:	Use case:	Problem statement:
Operations	Document intelligence	How can we better extract information from the mass of unstructured data in documents?

Solution:
Using technology for intelligent reading, analysis and extraction of insights from voluminous and unstructured documents, e.g., investment management agreements, prospectus, Key Investor Information Document (KIIDs), legal agreements, compliance reports, and regulatory documents, to help employees be able to spend their time in more valuable tasks.

Function:	Use case:	Problem statement:
Operations	Communications mining	How can we better extract and understand information from the mass of unstructured communications in emails, chats and systems?

Solution:
Using technology for data mining and NLP, establish a holistic view of the unstructured communications which dominate many processes, such as post trade, IT service help desk and complaints. This enables a view into the reality of processes that operations managers have never had before, giving them opportunities to identify underlying root-cause drivers of activity, so as to identify productivity improvements.

ESG

Changing socioeconomic factors have fuelled the growth in environmental, social and governance (ESG) demand, resulting in businesses wrestling with a renewed focus from clients on long-term value creation and sustainability. Many firms now aspire to achieve improved sustainability capabilities through the value chain from product offerings, through the investment process, to client, board and regulatory reporting. However, this has proven to be a challenge given that the ESG data needed is provided from many different and inconsistent sources.

This is where AI is being applied to ESG data: to help make sense of the mass of varied, unstructured and inconsistent data. Common AI technologies used in the ESG space include computer vision, NLP and generation, data quality detection and imputation, time series and forecasting, just to name a few.

5. THE AI LIFE CYCLE

Firms of all sizes that want to remain competitive in the market need to innovate to grow and protect their businesses. Part of that innovation is introducing the latest advances in technology. However, despite a rise in investment, successful application of AI and ML has been a slow process in most enterprises. Only a small fraction of companies have successfully launched internal AI and ML initiatives to transform their operations. This section explores the typical AI life cycle and how organisations can adapt and pivot to newer ways of working to make the most of the latest opportunities presented by AI.

The AI life cycle can be broadly divided into three stages – proof of value, enablement, and ramp-up or go live.

The AI lifecycle					
Proof of value Experimentation and prototyping		**Enablement** Stakeholder validation and garnering resources		**Ramp-up and go-live** Development and production	
Problem statement and rapid prototyping	Data management	Build vs buy (Exploring ecosystem)	Risk, control and governance	Model development lifecycle	Monitoring and maintenance

The proof of value stage solely aims at validating whether an idea is worth the additional investment and resources. A targeted problem statement and rapid prototyping with an end goal in mind is the key to succeeding at this stage. In parallel, technologists must also ideate with a data management strategy in mind. They must ask questions such as: What sorts of data are readily accessible? How is the quality and completeness of this data? Where do we obtain supplemental data from? A strong understanding of data is crucial to any ML development.

The focus of the enablement stage is to quickly turnaround management support and resources to support technology development. For organisations that do not possess internal technical capabilities or face legacy system challenges, there are always alternatives to outsource and hire.

An organisation is ready to ramp-up or go live when a full product development plan and scope is in place, including the necessary technology assets and personnel. Throughout the entire process, it is important to keep in mind the principles of responsible and ethical AI to ensure that outputs are always fair and representative.

PROOF OF VALUE

The first step in the AI life cycle is the assessment of value each use case will deliver to the business and whether the organisation has the desired datasets to meet, to deliver the value.

A proof of value asks questions such as: 'How valuable would it be to predict the market movement, and can we do it accurately with existing datasets?' and 'If we could predict could the market behaviour, what can we actually do anything about it?'

The benefit needs to be assessed against the cost of implementation vs. the opportunity cost of using resources elsewhere. It is designed to assess, explore and take the risk out of projects. Through multiple proofs of value, you can identify a portfolio of the most promising use cases.

The proof of value:

- Helps understand the various use cases that can be implemented with the existing datasets and the overall business value that each use case project adds.
- Prevents significant investments in projects that will not deliver business value.

- Identifies gaps in the current datasets for future investments.
- Helps businesses prioritise investments in AI or ML projects by clearly evidencing business value that will help secure wider organisational and stakeholder buy-in.
- Helps prevent over-engineered solutions (e.g., there is no need to consider a deep neural network if a simple statistical model delivers similar results).
- Takes some of the risk out of projects.

DATA MANAGEMENT

To fully harness the power of AI, certain preparatory steps and frameworks need to be implemented to ensure that the accuracy and reliability of the model meet expectations. To do so, however, organisations must first understand the important, yet delicate relationship between data and AI.

Whilst it is evident that data is a big enabler to deliver business strategy, it is absolutely critical for the successful implementation of AI. 'Access to data' and 'quality of data' are two of the biggest challenges.[6]

There are two main ways in which AI needs data:

1. Supervised learning: Historical data is collected about past events and circumstances. Supervised ML algorithms use this historical data as an input called a 'training dataset' to train the model to predict new output values.

2. Unsupervised learning data patterns: AI can be used to find structure and patterns in data so that algorithms can learn and adapt on a continual basis. Relationships and knowledge linked to data sources are continually computed by the model and used for predicting new output values.

Organisations capture extensive amounts of data originating from a variety of different sources. This section of the report discusses the practical steps organisations can take to improve their data capabilities and how this will benefit their data science operations.

AI IMPLEMENTATION HURDLES IN INVESTMENT MANAGEMENT

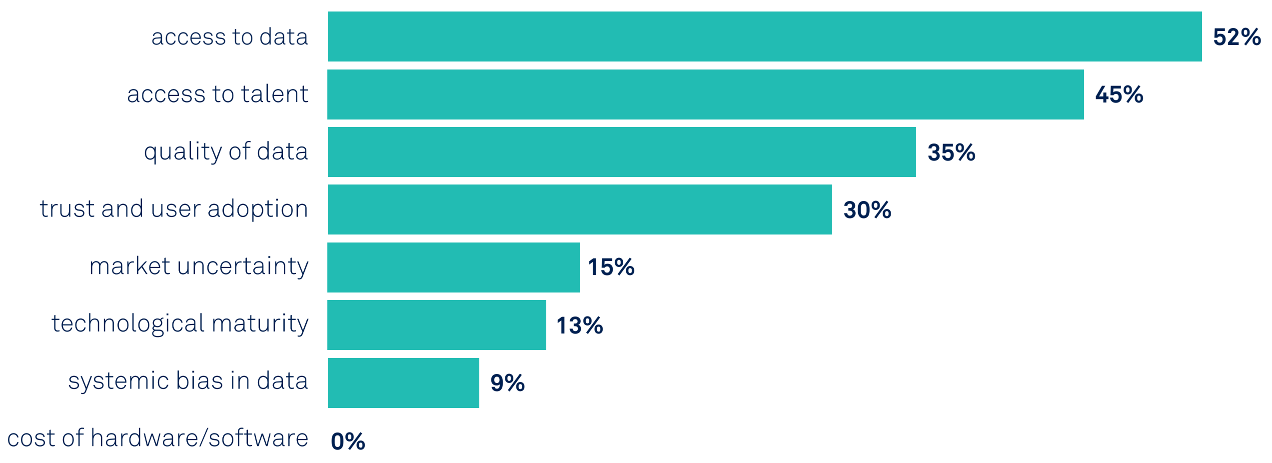

[6] "Transforming Paradigms: A Global AI in Financial Services Survey," World Economic Forum, January 2020.

Quality of data

Improving data quality is the most important and significant step that organisations can take to not only enhance their AI models but also improve other capabilities such as business intelligence. Good data quality ensures reliable outputs from the model, making data quality and trust synonymous with each other. So, what can firms do to start to improve their data quality?

A useful exercise to begin the journey towards better data quality is to conduct data profiling. Data profiling is the process of examining, analysing and creating useful summaries of data. The process helps create a standardised measure for assessing data quality issues, risks and overall trends.

Creating a set of standards that are designed to tackle flaws in the data will help streamline and accelerate the data cleanup required. Once identified, tools can be implemented within the process that can help automate the correction of the errors, those which are repetitive in nature. However, it is important that all key stakeholders agree with the definitions of anomalies, outliers and other aspects of the data that need decision-making.

Structuring and storing your data

The ability to handle large volumes of data via a common repository and standardised data access layer is critical to the success of AI. A standardised data layer allows for democratizing the data so that different teams can easily access it for their own functions. Creating a central repository ensures that data used across the business is consistent and that the right insights are generated.

Data science capabilities work hand in hand with a business' data infrastructure. Establishing a data-driven culture, and developing and maintaining data processing pipelines and data repositories is the backbone of a successful AI model.

ENABLEMENT

A key step in the AI life cycle is adopting AI at scale. There is a large ecosystem of fintech companies that offers solutions that range from solving specific problems (such as document automation and anti-money laundering (AML) fraud detection) to more generic data and ML development platforms. Firms must consider across the 4 'Cs' — cost, capabilities, customisation and connectivity when deciding on build vs. buy.

CAPABILITIES VS. CUSTOMISATION

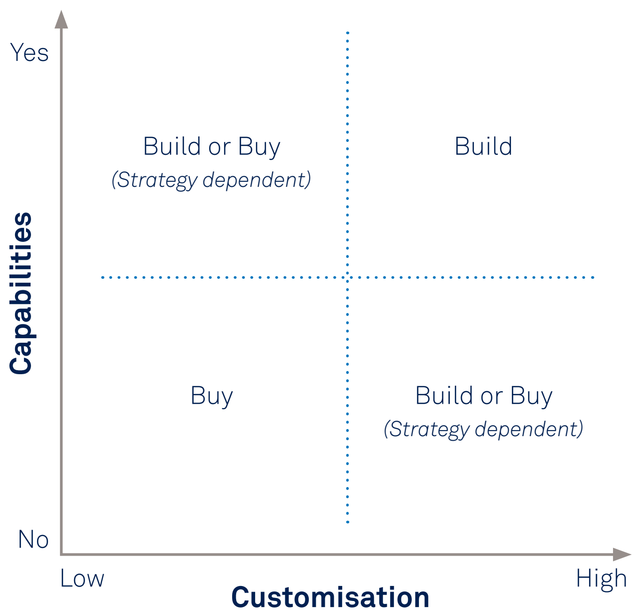

- **Capability**
- **Customisation**
- **Connectivity**
- **Cost**

The capabilities vs. customisation graph helps broadly in terms of the build-vs.-buy decision. The best option might be to leverage the fintech ecosystem in collaboration with internal capabilities.

RAMP-UP OR GO LIVE

Model development is an iterative process that involves several disciplines as shown below. The various circles indicate the different stages involved in the process.

1. Feasibility review and project scoping:

As mentioned in the first section, the first step is to evaluate whether ML is the right approach to the problem. Project scoping is about laying out goals and objectives, constraints, and evaluation criteria. The topics covered in this step includes overarching business goals, technical objectives, constraints and evaluation.

2. Data management:

The scoping of the project in the previous step would ideally indicate the scale of the project, and consequentially, the scale of the data requirement. Refer to the chapter titled, 'data management' within 'the AI life cycle' on considerations around governance, data quality and storage.

3. ML model development:

This is the stage that requires the most technical ML knowledge. Depending on the project scope, this model can be supervised or unsupervised, use simple tree-based architecture or deep neural networks. This is a highly iterative process where the data scientists experiment with different features from the input data and engineer those features into new ones whilst sampling different methods. On the basis of the project scoping, certain technical specialisation may be needed. For example, anomaly detection, NLP, or time series analysis. In recent years, availability of platforms that help with ML modelling process has significantly reduced the time and shortened the development cycle.

In ML modelling, one of the biggest considerations is the 'bias-variance trade-off. In essence, it means that model prediction can breakdown due to:

1. In-sample error
2. Model instability

If the model poorly fits training data, it is likely to poorly fit on unseen data and therefore will carry its own bias. The bias can be reduced by increasing the complexity of the model (i.e., adding more features or tuning parameters more aggressively) or choosing more sophisticated model architecture (e.g., moving from Bayesian models to deep neural networks), which will help fit the training data better. However, increasing the complexity of the model can lead to 'overfitting' due to the model learning the 'noise' from the training data that do not exist in unseen data. This instability is also referred to as 'model variance' and can lead to higher prediction errors. It is important to choose a model that finds an optimal balance between model bias and variance. In other words, this means choosing a model that fits the training data well enough to generate meaningful predictions but not too well that it leads to overfitting. governance, data quality and storage.

4. Deployment (and re-deployments):

Deployment is all about making the model accessible to users. Depending on the use cases, the deployment can be for real time (one prediction at a time) or batch inferencing (multiple predictions at a time). The deployment can also be small or large depending on the number of users and requests expected at any time, therefore, the right technology and infrastructure should be selected on the basis of this. This is an area where ML operations (Ops) specialists come in with their specialism in deployment and orchestration, as well as a good understanding of infrastructure and load-balancing, ensuring that the ML model, as a service, has high uptime and is speedy in responding (low latency).

At the point of 'go live,' the following aspects of the model should be well defined:

1) Concept and data drift: Concept drift is when the patterns that the model learned no longer hold i.e., the nature of the pattern or reality has changed. Data drift is when input data have changed significantly, compared with when the input data was used to train the model (i.e., changing distribution). It is imperative to have controls in place to identify concept and data drift (as applicable) on a regular basis and raise the process to follow in case of a breach.
2) Re-training: Models need to be re-trained if concept or data drift has been identified. Depending upon the type of business and the models used, re-training might be one off or can be scheduled at regular frequency.

5. Monitoring or maintaining and business analysis:

Once in production, models need to be monitored for performance decay and maintained to be adaptive to changing environments and changing requirements. This step can be an extension of the previous step and leverages ML Ops' expertise and personnel.

Tying everything together, model performance needs to be evaluated against business goals and analysed to generate business insights. These insights can include productivity of projects and potential for new projects.

6. REGULATORY CONSIDERATIONS

In recent years, we have seen existing law and regulation adapt to address the development of AI technologies. In parallel, we have seen the implementation of new standalone regulations impacting AI, such as the European General Data Protection Regulation (GDPR) (and its UK equivalent, UK General Data Protection Regulation (UK GDPR)) and the proposed EU regulation on AI, which would be the first ever legal framework on AI.

Regulators, governments and industry bodies around the world have also laid out their expectations by developing guidelines and principles for responsible AI. For example, the European Commission has produced the Ethics Guidelines for Trustworthy AI and the UK's data privacy regulator, the Information Commissioner's Office (ICO) has published detailed guidelines on Explainable AI. The picture is constantly evolving, but common global regulatory themes include the importance of fairness, human oversight, governance, transparency and explainable AI.

Financial regulations

- In the UK, the Prudential Regulation Authority (**PRA**) and Financial Conduct Authority (**FCA**) have both shown that existing regulatory principles apply to the use of AI, including rules on Senior Management Arrangements, Resilience Systems and Controls (**SYSC**).

> *"Boardrooms are going to have to learn to tackle some major issues emerging from AI – notably questions of ethics, accountability, transparency and liability."*
>
> — Financial Conduct Authority (FCA), August 2019, AI in the boardroom

- The FCA has reported on the use of ML within UK financial services, finding that key issues include software validation; governance, resilience and security of ML applications; and potential ethical issues when using ML and novel data sources. Following this report, the FCA established a public-private working group on AI, to explore these issues with both the public and private sectors, and determine the actions required to support the safe adoption of AI in UK financial services.

- The FCA has also reported on the supervision of algorithmic trading in wholesale markets. As set out in section 5, algorithms are amongst the instructions used for ML systems. Firms should be aware that the FCA can require them to produce a description of their algorithmic trading strategies within just 14 days, and that it recommends that firms have a detailed 'algorithm inventory' setting out coding protocols, usages, responsibilities and risk controls.

- Breaches of FCA principles in relation to AI also give rise to further exposures for financial institutions' senior managers (under the Senior Managers and Certification Regime), and to additional potential civil liabilities under the Financial Services and Markets Act 2000, which grants private persons a right to sue the firm in respect of losses suffered as a result of FCA or Prudential Regulation Authority (PRA) rule breaches.

- With regard to outsourcing, SYSC 8 requires certain FCA-regulated firms to take reasonable steps to avoid any undue operational risk and not to outsource important operational functions in a way that could materially impair the quality of its internal control. The European Banking Authority (**EBA**) Guidelines also state that outsourcing must not lead to a situation in which an institution becomes an empty shell that lacks the substance to remain authorised. These requirements will apply to the use of AI, particularly when provided by a third party.

Market abuse

- When using AI during the investment process, firms should be aware that the FCA has brought attention to the risk that AI may further financial crime. It has focused on procedures to counter this, including testing to assess the impact of systems on market integrity and post-trade monitoring.

- It is possible that AI used in the financial services industry may commit market manipulation under the Market Abuse Regulation (MAR), as a result of being exposed to certain markets and data and being programmed to achieve the pre-set objectives, even if such manipulation is not intended by the users or governance team. The FCA has previously commented that '*the FCA cannot prosecute a computer, but [it] can seek to prosecute the people who provided the governance over that computer*.'

- If trading or investing on the basis of big data analysis, financial institutions or investment managers also need to ensure that the datasets they use do not contain confidential information (whether from within the firm or elsewhere) that amounts to inside information. If using AI tools to make or determine orders, they should be sure that the tool will not behave in a manipulative manner, whether immediately or later through iterative 'learning.' Where AI is used in generating published or disseminated information (for example in generating published research), institutions need to ensure that such information is not misleading.

Competition

- The UK Competition and Markets Authority (**CMA**) uses its powers to restrain technology with an anti-competitive objective. This includes anti-competitive conduct by an AI system that has, through iterative 'learning,' identified that a desired outcome can be achieved through anti-competitive strategies.

- More recently, in January 2021, the CMA's Data, Technology and Analytics Unit revealed its plans to launch a substantial programme of analysing algorithms to identify potential harms to competition and consumers, particularly with respect to:

  (1) Personalised pricing or consumer journeys which result in 'artificial' changes to consumer behaviour.

  (2) Algorithmic discrimination.

  (3) Unfair ranking and design relating to certain products, services or suppliers.

  The CMA is also seeking evidence on how the use of algorithms could lead to collusion between competitors through explicit coordination, hub and spoke coordination, or autonomous tacit collusion. This is intended to equip the CMA with the knowledge to effectively regulate the use of algorithms, and requires firms to offer remedies.

- In April 2021, the CMA launched its Digital Markets Unit (**DMU**), to oversee a pro-competitive regulatory regime for digital activities. The DMU will introduce and enforce a code of conduct applicable to companies with substantial market power in digital activities. It remains to be seen to what extent financial institutions will be affected by the CMA's regulatory agenda.

Data protection

- The scope of data inputs used by financial institutions when using AI tools is broad and might include structured market data, unstructured big data (such as news reports and social media content), and customer personal data, the latter of which will be subject to the **UK GDPR** and UK Data Protection Act 2018 (**DPA 2018**). Data processing is likely to be part of each of the AI use cases set out in section 4.

- The UK GDPR promotes fair and transparent processing by requiring firms to provide individuals with meaningful information about the logic involved, and the consequences of the processing. In addition to the risk of enforcement action, financial institutions should be aware that UK GDPR (and section 168 of the DPA 2018) gives individuals the right to bring civil claims for compensation for personal data breaches, including for distress.

- In July 2020, the ICO published guidance on AI and data protection. It addresses how to apply data protection principles to AI systems that process personal data, as well as how to manage any trade-offs between different principles that may be needed, and auditing AI systems to ensure they are built and used in a UK GDPR-compliant manner. Although this is a key focus area for the ICO, as of late 2021, there does not appear to be significant enforcement actions on the horizon in this space.

- It is also worth noting in this context, however, that the UK Government's Department for Digital, Culture, Media & Sport (DCMS) has recently published proposals as part of its 'Data: a new direction' consultation on the future of data protection regulation in the UK, which intend to simplify rules around the use of personal data for AI, as well as current requirements around the need and thresholds for explainability. This may lead to significant reductions in compliance obligations for organisations, should the proposals be adopted in the form of amendments to the UK GDPR.

Ethical AI and discrimination

- There is an inherent risk that AI which incorporates biased datasets will create biased outcomes, leading to unfair or discriminatory decision-making.

- This is already covered under existing anti-discrimination laws such as the Equality Act 2010 and, to some extent, financial regulatory principles. Specifically, financial institutions regulated by the FCA must act in a way that is consistent with the FCA Principles (PRIN 2.1.1R), some of which relate directly to firms' treatment of customers. Specifically, FCA Principle 6 ('A firm must pay due regard to the interests of its customers and treat them fairly') would extend to any use of AI by an FCA-regulated firm.

- Whilst outright discriminatory practices are illegal, there are no codified laws on the ethical use of AI in the UK. Although discrimination can be tricky to spot, particularly when hidden proxies give rise to discriminatory outcomes, ignorance is no defence and enforcement actions are on the rise.

- Ethical values differ from country to country. They are hugely influenced by cultural considerations and are continually evolving. The core themes are fairness, accountability, transparency and human oversight. In particular, if a firm is unable to explain how a discriminatory outcome was achieved, the burden of proof will be against that firm in any regulatory enforcement action or litigation in the UK.

Liability in contract or tort

- AI usage (whether by a financial institution's suppliers or by the institution itself with its customers) may give rise to unintended consequences and expose institutions to claims for breach of contract or in tort, and test the boundaries of existing exclusion clauses. Organisations need to assess whether their existing terms and conditions remain fit for purpose where AI is concerned.

Towards standalone AI regulation

- On 21 April 2021, the EU Commission published its legislative proposal for a new EU regulation on AI (the '**AI Act**'). It is expected to become law in late 2022, with compliance mandatory from 2024 onwards.

- The definition of an 'AI system' under the AI Act is highly contested and will be difficult to determine with legal certainty. However, as it currently stands, it is very broad and could potentially capture almost any software used by firms, even if it does not involve any recognisable form of AI.

- The AI Act will apply particularly onerous requirements to 'high-risk AI systems,' which it has defined to include the following: biometric identification, recruitment or selection processes, and credit assessments of natural persons. AI is commonly used for these purposes within financial institutions, and it remains to be seen whether the exact nature of the AI systems used will be categorised as 'high-risk' under the proposed AI Act.

- However, financial sector firms will likely be subject to the new obligations. These apply whether they are users, providers, importers or distributors of AI systems subject to the AI Act.

- Whilst firms may be able to rely on existing procedures, systems and controls to meet at least some of the obligations of the AI Act, they are also likely to need to put in place new ones. Under the AI Act, the more extensive requirements applicable to high-risk uses include: establishing prescribed systems for risk management, quality management and post-market monitoring; preparing technical instructions for use that include the capabilities and performance limitations of the system; and registering the software on a publicly available European Commission database (including the instructions for use, even if these are commercially confidential).

7. MITIGATING THE RISKS OF MISUSE OF AI

Although AI changes the risk profile of an organisation, various aspects of the conceptual framework applicable to risk management and governance for different risk categories also apply to AI. Organisations will need to decide whether to treat AI as a new standalone risk category or to include it within the existing risk profile of each business process that uses AI. Regardless of the approach taken, there are a number of essential mitigation measures:

- **Overview and exposure:** Financial institutions need to ensure they have a complete overview of where AI applications are being used, and an inventory of where they are exposed. Some uses are very visible (e.g., in an investment decision process), whereas others may be more hidden (e.g., NLP enablers in intelligent automation). The inventory should also contain adequate information on the intended use of the AI so that its actual performance can be assessed against this objective. This information will form the basis of an organisation's risk assessment for each use case, and enable it to identify the use cases that present the highest risk and the mitigation measures required.

- **Governance and culture:** Financial institutions should create a robust governance and internal reporting structure with a culture of transparent and ethical use of AI embedded in the organisational leadership, including the company's boards, general counsels, senior data, compliance, risk and policy teams overseeing AI risk management. They also need to determine and document management responsibility for the institution's use of AI, with a consistent approach taken across the organisation.

- **Systems and controls:** Given the PRA and FCA's willingness to apply existing regulatory principles to AI, financial institutions will need to be mindful of an overreliance on automation, insufficient oversight and ineffective systems (as for any existing processes). For example, where robo-advice (automated investment services) is being provided and involves the processing of personal data, the FCA has stated that regulated financial institutions will need to 'ensure clear oversight over the auto advice proposition, as well as clear allocation of responsibilities' and consider how any explanation might need to meet the expectations of both data and financial regulators.

- **Testing:** Ongoing testing and monitoring of AI solutions, far beyond the development stage of the solution, should involve a wide range of stakeholders and span multiple areas (e.g., HR, technology, operations and engineering). Depending on the use case and risk involved, controls may be intensive and ongoing, or less frequent (e.g., data set vetting and post-model calibration). The tests and controls used should adapt and evolve as the technology develops.

- **Due diligence and regular audit:** Financial institutions should make it a top priority to assess where and how AI is used in their organisation; AI functionality, including the source of data, transparency and explainability of the system; the limits and boundaries set for the use of AI; and the apportionment of contractual and tortious liability between programmers, suppliers, the company and its clients in standard terms of business. Further, they should ensure consistency between existing policies which contemplate the use of AI, e.g., GDPR compliance policies, human rights policies, competition policies, codes of conduct, and new product approval process guidelines.

- **Explainability and transparency:** Financial Institutions must be able to provide clear explanations of how decisions involving AI are made, at every stage of the process. These explanations must be set out in a transparent and accessible way, and made available to employees, customers, regulators and other relevant stakeholders. This applies whether their AI is bought or built in-house. Guidance from the ICO and Alan Turing Institute, published in May 2020, sets out how to approach AI explainability from the point of view of key legal requirements, and the technical options and measures available for producing and delivering explainability in context. The guidance identifies six different AI explanations, as well as advice on governance to ensure explainability at the organisational level. The nature of the explanation required varies according to the use case, and the impact of the AI system on the affected individuals. It is important to have explainable AI systems for a number of reasons, including:

- Performance: Organisations need to consider and assess the trade-offs between certain risks and values, including the trade-off between false positives and false negatives; between statistical accuracy and data minimisation; and between increased complexity and 'overfitting' (see Section 5 above).

- Fairness: AI is increasingly used to support decisions that impact the lives of human individuals, including hiring decisions and access to credit. This can lead to ethical questions and engage anti-discrimination laws (see Section 6 above). One of the ICO's six AI explanations is the 'fairness explanation,' which involves explaining the steps taken to ensure AI decisions are unbiased. Practical ways to provide a fairness explanation include ensuring that training data is representative, engaging domain experts to advise on the right data for the relevant objective, looking out for hidden proxies for discriminatory features and training personnel in how to detect bias.

- Interpretability: Models can often be very difficult to visualise, or can contain millions of parameters or activation functions. This can lead to 'black box' issues. Several techniques are now available to help with interpretability. The goal of these techniques is to give insights on 'feature importance,' which assigns scores to inputs based on how useful they are at predicting a target variable.

"Whilst discrimination is a broader problem that cannot realistically be 'fixed' through technology, various approaches exist which aim to mitigate AI-driven discrimination."

— ICO, 2020 Guidance on AI and Data Protection

"The fairness explanation is about helping people understand the steps you took (and continue to take) to ensure your AI decisions are generally unbiased."

— ICO and Turing Institute, 2020 Guidance on Explaining Decisions made with AI

- **Record keeping:** the measures and procedures described above should be recorded and documented to a high standard. These records may be the first items requested by regulators, litigants or politicians if any AI issues arise. They help to demonstrate that an institution properly understands where and how AI is used within the business, has clearly allocated responsibility for that AI and is actively managing AI risk.

With thanks to EY and Clifford Chance for sharing their insights and expertise.

THE INVESTMENT ASSOCIATION

The Investment Association
Camomile Court, 23 Camomile Street, London, EC3A 7LL

www.theia.org
@InvAssoc

November 2021