

THE
ENTERPRISE
AI COMPASS:
NAVIGATING
TALENT,
GOVERNANCE
AND ETHICS





Introduction: AI's Arrival in the Enterprise

It is clear, from the levels of attention in the press, in service and product companies, and from researchers, analysts and consultants, that artificial intelligence has arrived in the enterprise.

However, despite these levels, it's not always clear exactly what we mean by 'artificial intelligence'. The term is often used as if it means no more than 'something we can do with computers today that we couldn't do yesterday.'

For the purposes of this paper, we suggest that today, the current generation of Artificial Intelligence (AI from here onwards) means three things. First, it means a set of mathematical and computational techniques which, coupled with large datasets and abundant computing resources, can carry out prediction and categorization with an accuracy and speed that rivals or beats humans. Second, it means the application of these techniques to emulate human capabilities such as visual perception, speech recognition, decision making and translation between languages. Third, it means the use of these capabilities in enterprises to enhance existing decision making technologies, but also to automate work which could not previously be automated, or to derive insight which could not previously be derived.

The use of AI within enterprises is dependent on the results of decades of foundational research conducted since the earliest days of digital computing, but has been made commercially viable by the explosive growth of data, and the availability of cheap and easily accessible computing resources from public Cloud

providers. This acceleration has been further boosted by the development of the ecosystem that now accompanies any rapidly maturing technology: the emergence of widely used Open Source frameworks, the launch of commercial products, the development of capabilities within service companies. However, these advances don't solve all practical problems: enterprises adopting AI still need to work hard to make sense of a rapidly developing market, to source and manage the datasets AI depends on, and to find people with the right skills and talent.

The increasing adoption of AI in enterprises has also, of course, sharpened focus on potential problems and risks, which, like the technology itself, have graduated from theory to practical reality with speed.

Government, regulators, researchers, the press and society at large have all raised legitimate questions about the potential impact of AI on privacy and employment, as well as the degree to which AI can be trusted not to reinforce existing bias or introduce new bias.

This all means that AI presents a series of opportunities and challenges to companies in the financial services sector. As enterprises which depend on insight and decision making, which have a continuous drive to improve operational efficiency, and which possess large and rich datasets, AI holds a lot of promise.

At the same time, as financial services companies which fulfill a trusted role in society, with strong obligations to be fair and transparent, and to protect customers and have a high level of regulatory oversight, must be particularly attentive to the risks of AI and their responsibilities.

This paper gives an overview of:



The current approach to adoption of AI within financial services, and the way this approach might develop.



The practical considerations for financial services companies adopting AI, including how to approach the AI market, and how to acquire, develop and partner for talent.



The policy considerations for financial services companies adopting AI, including how to govern AI, how to address ethical questions, and the emerging stance of regulators.



AI in Financial Services: The Early Path to Adoption

The use of technology for decision-making, insight and automation is well established in financial services: despite often being depicted as slow to innovate, large financial services companies are amongst the biggest spenders and early adopters of new technology.

As a result of decades of computerization, most financial services companies have architecture which can broadly be broken down into systems which interact with customers, systems which provide the transactional and operational core of the organization, and systems which take data from the other systems and use that data for operational insight and action. The introduction of AI technologies

has not immediately transformed this architecture, but has rather augmented it: customer interactions which have traditionally been forced into structured forms can now start to benefit from conversational and contextual capabilities; operational decision-making which was either dependent on deterministic rules or human judgement can now be augmented by recommendations by models; and the work of traditional analytics is steadily transforming into the work of data science and the training of models.

The benefits for financial services companies of developing their

architectures in this way are estimated to be high: UK-based research provider Autonomous Research projected cost savings through AI of \$1 trillion for the entire financial services sector by 2030. For banking alone, the benefits were projected to be \$447 billion, a 20 percent cost saving.

An Infosys survey of corporate leaders showed that 80 percent of global CEOs expect AI applications to drive at least 30 percent of their new revenue and cost-saving initiatives in the next 24 months.

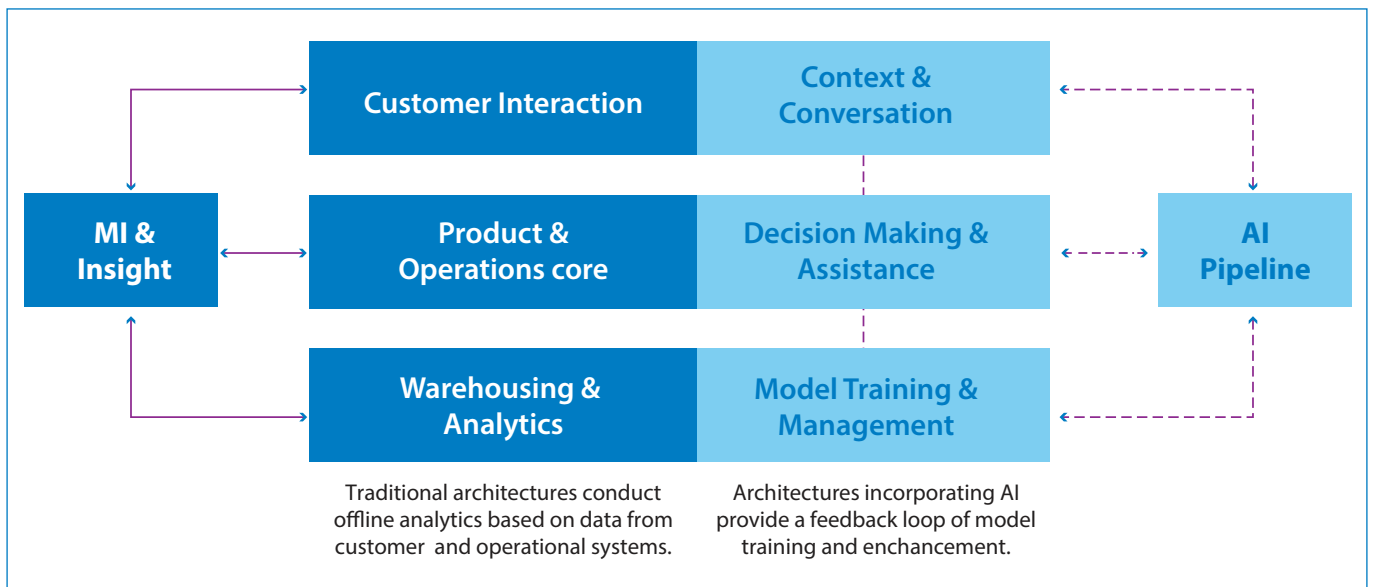



Figure 1: The Development of AI Architecture



HSBC's Journey

Over the last three years, the adoption of AI within HSBC has progressed from a few isolated experiments conducted by deep specialists, using limited datasets and traditional technologies, through the first wave of successful deployments in analytics-rich domains, such as financial crime detection, to increasingly widespread adoption through partnerships with Cloud providers such as Google, Amazon and Microsoft, with new companies such as Quantexa, and with innovation and development partners such as Infosys. Those initial early adopters are now part of a global community of practice. While adoption is still in its early days, HSBC is realizing benefits in customer

engagement, risk management and compliance, and is seeing the technical and talent barriers to AI adoption steadily fall. HSBC's perception, though, is that the market for AI products is, with some exceptions, focused on general purpose platforms and tools, with less emphasis on solutions for specific industries such as financial services. Consequently, HSBC continues to focus on building internal talent and developing the specific solutions needed by its customers and business, often through the use of Open Source frameworks. HSBC is also deeply aware of the growing attention being paid to the use of AI by customers, regulators, governments and society, and of the obligation of companies to make

responsible use of AI and of the data which powers it. The bank has defined principles for the ethical use of Big Data and AI, and established a panel to assess the application and development of these principles. The number of AI solutions in HSBC is projected to triple in the next three years.

The teams making the best progress are characterised by a combination of access to flexible infrastructure, the right datasets, internal talent and strong external partnerships.



Allianz's Journey

Predicting and pricing of risks is at the very core of the insurance business: it was therefore natural that the AI journey of Allianz started already in the 1980s in the area of insurance product design and pricing, with the use of data analytics methods. In this case, regression algorithms were used to calculate insurance tariffs.

Clear acceleration was triggered in the early 2010s, prompted by the mass emergence of new data sources, including web analytics, IoT and telematics. This resulted in the establishment of foundational capabilities: a core team of data scientists was formed and a global Data Science Platform was built to provide the infrastructure enabling acceleration of AI initiatives throughout Allianz. This acceleration brought concrete results, proving the potential of AI as a lever of strategic relevance throughout the value-chain, from lead management to claims handling automation through policy management and fraud prevention.

However, this broad-based, bottom-up approach also had limitations: AI developments tended to be focused on local / specific needs, not always fully integrated into final business process workflows, and rarely developed for scaling across the business. As a result, Allianz has now entered a new phase of AI maturity, focusing on scaling-up the business value of AI to the whole enterprise. AI is now a key dimension of Allianz's global transformation of processes, people and culture following a common, customer centric model: the Allianz Customer Model (ACM).

In a nutshell, it is clear that technology and leading edge data science skills are indispensable enablers for AI. However, harvesting the full potential of AI is a transformational project requiring a strategic and cultural shift. Allianz believes that it will eventually help redefining the business model of insurance from "detect and repair" to "predict and prevent".

This calls for addressing 4 inter-related challenges:

1



Ensuring data completeness, quality and accessibility across functions and geographies as a key lever for AI performance and scalability.

2



Leveraging state of the art technologies to develop AI products and to achieve full integration into operational value-chain.

3



Disciplined execution and steering of AI projects with strong integration into harmonized business processes.

4



Fostering a strong data-driven culture with no compromise on the principles of fair products, fair distribution and fair data usage.



Infosys' Journey

In order to help address these practical challenges, Infosys has created several assets such as the AI Workbench, the Data Governance Workbench, the iDevOps Workbench, and reference architectures for building a native AI / ML enterprise.

These assets aim to democratize and govern the AI system development process, provide the ability to consume and pre-process unstructured data, standardize the process for model development, enable seamless consumption of native AI components and enable control and machine learning. HSBC and Infosys have also jointly established the JuniperX Innovation Centre in London to accelerate the ability to place financial services data on the

public Cloud, improving access to Cloud-based AI tools and services.

Infosys has created the world's largest corporate learning infrastructure, both physical and digital, to support hundreds of thousands of employees on their learning journey.

From an AI perspective, Infosys realizes that amplifying value of human potential is needed to build, deploy and manage AI applications. Infosys is in the process of building an AI Centre to support the ever expanding needs of their clients in this space.

Through a unique ecosystem, supported by academic institutions, AI startups,

AI Cloud tech vendors and Infosys' own capability and investments in AI technologies, the firm believes it is possible to address the challenges in developing next-gen AI technologies that large enterprises need.

HSBC and Infosys are also jointly leveraging Infosys' WINGSPAN learning platform and global learning infrastructure to deliver digital capability training to its graduates on a yearly basis. The platform uses AI to map out the learning path for individuals from legacy technology to digital and new age technologies.

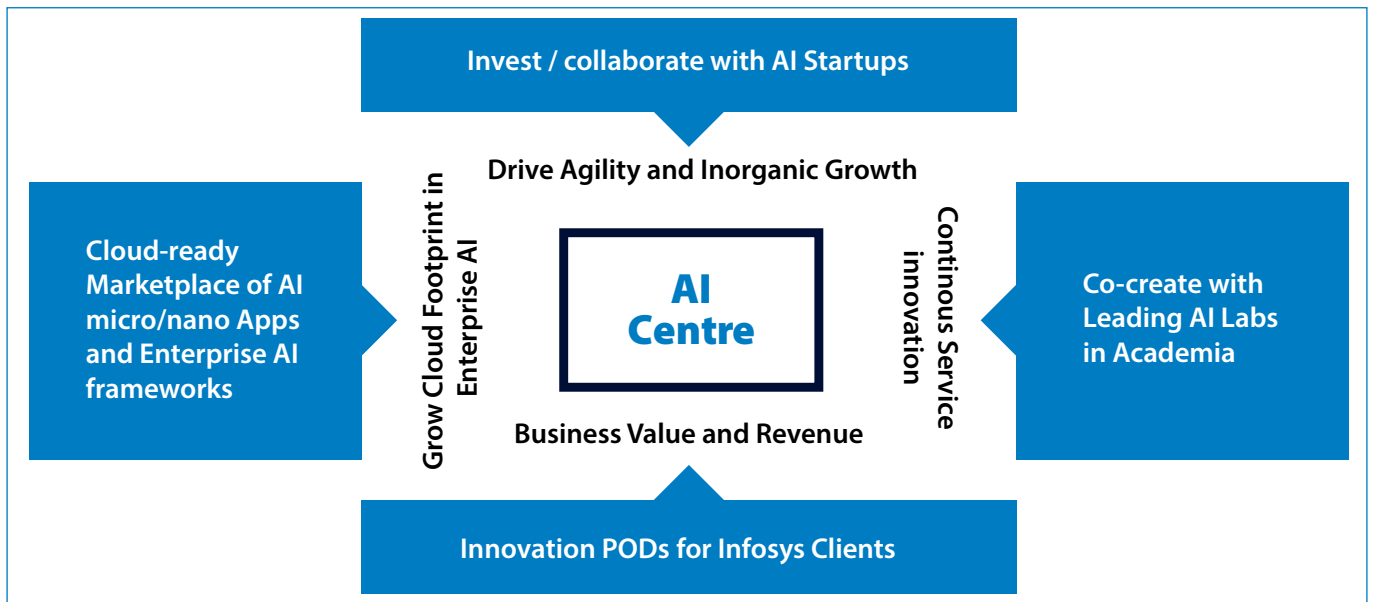


Figure 2: The AI Centre Construct



Practical Considerations for AI Adoption

Despite the rapid growth in the maturity of AI and its applicability in financial services, those companies planning to adopt AI and scale its use still face a number of practical challenges:

Talent, skills & understanding

The specialist skills required to develop AI solutions are still rare. Even companies who have strong analytics and data science teams, as well as strong software engineering teams, have few people with the right combination of these capabilities.

Indeed, many leaders within enterprises have a limited understanding of AI and its potential, meaning that opportunities for strategic or tactical value may be missed.

There has been much speculation on the impact of AI on employment, with some arguing that it will destroy jobs by automating work which could not be automated before, and others arguing that like many other advances in technology in the past, AI will result in a net increase in employment – although we may not know now what form that will take.

Investment in learning programmes:

The development of good AI solutions requires a combination of data science and software engineering skills which are rare to find together, and which are themselves in demand.

Leaders need to provide dedicated time and support to their teams, to assess and measure their overall level of AI competency, and to determine where training should be focused, including looking beyond traditional audiences for such training.

Work with universities and partners:

Despite its arrival in enterprises, leading AI skills are still strongly represented in academic institutions. At the same time, services companies are highly motivated to meet the emerging needs of their clients, and are developing extensive training capabilities.

Partnering with academic institutions and technology companies can provide access to existing knowledge – while offering those partners access to real-life problems.

Learning by doing:

As with most other technical skills, once they have an initial grounding, teams learn fastest by working on practical problems. Leaders should encourage their teams to apply what they've learned to real AI projects.

By giving teams the opportunity to identify and solve problems within the areas they support, or even within the operations of their own team, they will develop their design skills and their ability to spot with AI.



Rapidly developing market

The speed at which the market for AI products and services has developed, has made the current wave of AI adoption possible. However, the market is still rapidly developing, making it difficult for enterprises to choose who to work with. Furthermore, the AI technology market still lacks maturity: the ecosystem of tools to manage AI solutions throughout their lifecycle is at an early stage, and the range of industry vertical solutions is still small. Gartner predicts the business value created by AI will reach \$3.9 trillion in 2022. IDC predicts worldwide spending on cognitive and AI systems will reach \$77.6 billion in 2022.

Given this potential market size, it is unsurprising that the explosion in AI technologies, which has made


its adoption in large enterprises commercially viable continues, while making it increasingly difficult for those same companies to make sense of the market.

Should they work with their established technology partners, all of whom are investing in AI? Should they establish new relationships with digital giants, who have only started selling AI solutions to enterprises relatively recently? Should they work with startups, who need access to enterprises' data, as much as they need access to revenue, or should they build their own capability using widely adopted Open Source frameworks?

As with talent, this is not a new problem for enterprises, that face a classic choice between buying (or otherwise acquiring) and building.

Data access & usage

AI solutions are dependent on data. Most financial institutions have large amounts of data, but not typically in a form which can be used directly by AI solutions. Even when data has been extracted from operational and transactional siloes into analytics platforms, enterprises have rightly drawn strong boundaries between the data used to design and test solutions and the data used for production. The AI development lifecycle challenges these boundaries, and enterprises need to find new approaches which allow data to be used for training without compromising privacy or security.



The Buy vs Build Matrix

For some capabilities, this choice is straightforward. General AI capabilities such as image recognition, text scanning, speech-to-text conversion and language translation require large datasets which most enterprises don't have access to, and have already been comprehensively addressed.

'If the goal is an AI project that handles routine activities and delivers immediate value, it's almost never a good idea to build your own', says Thomas Davenport, professor in management and information technology at Babson College. 'Even using open source tools, build-it-yourself AI can cost millions of dollars, and it can take months to

train a machine learning algorithm to do what most vendors have already accomplished. However, the market for domain-specific solutions in financial services is relatively young. While companies are emerging to address well known, data intensive problems such as financial crime detection, credit decisions and customer needs analysis, most companies outside financial services do not have the datasets associated with these problems, meaning that they must work in partnership with their clients to develop solutions. At the same time, given the early stage of adoption, there are many opportunities for financial services companies to establish differentiating capabilities, applying AI

to fields ahead of their competitors. The AI market is unlikely to undergo consolidation or standardization in the near future, meaning that enterprises wishing to realise the full potential of AI must develop their own capabilities and solutions, establish a network of partners, and have the appetite and ability to consume general AI capabilities from established providers.

As is often the case with the buy vs build choice, especially in emerging fields, the answer depends on context, and it is likely that all companies will be using all these options for some time.

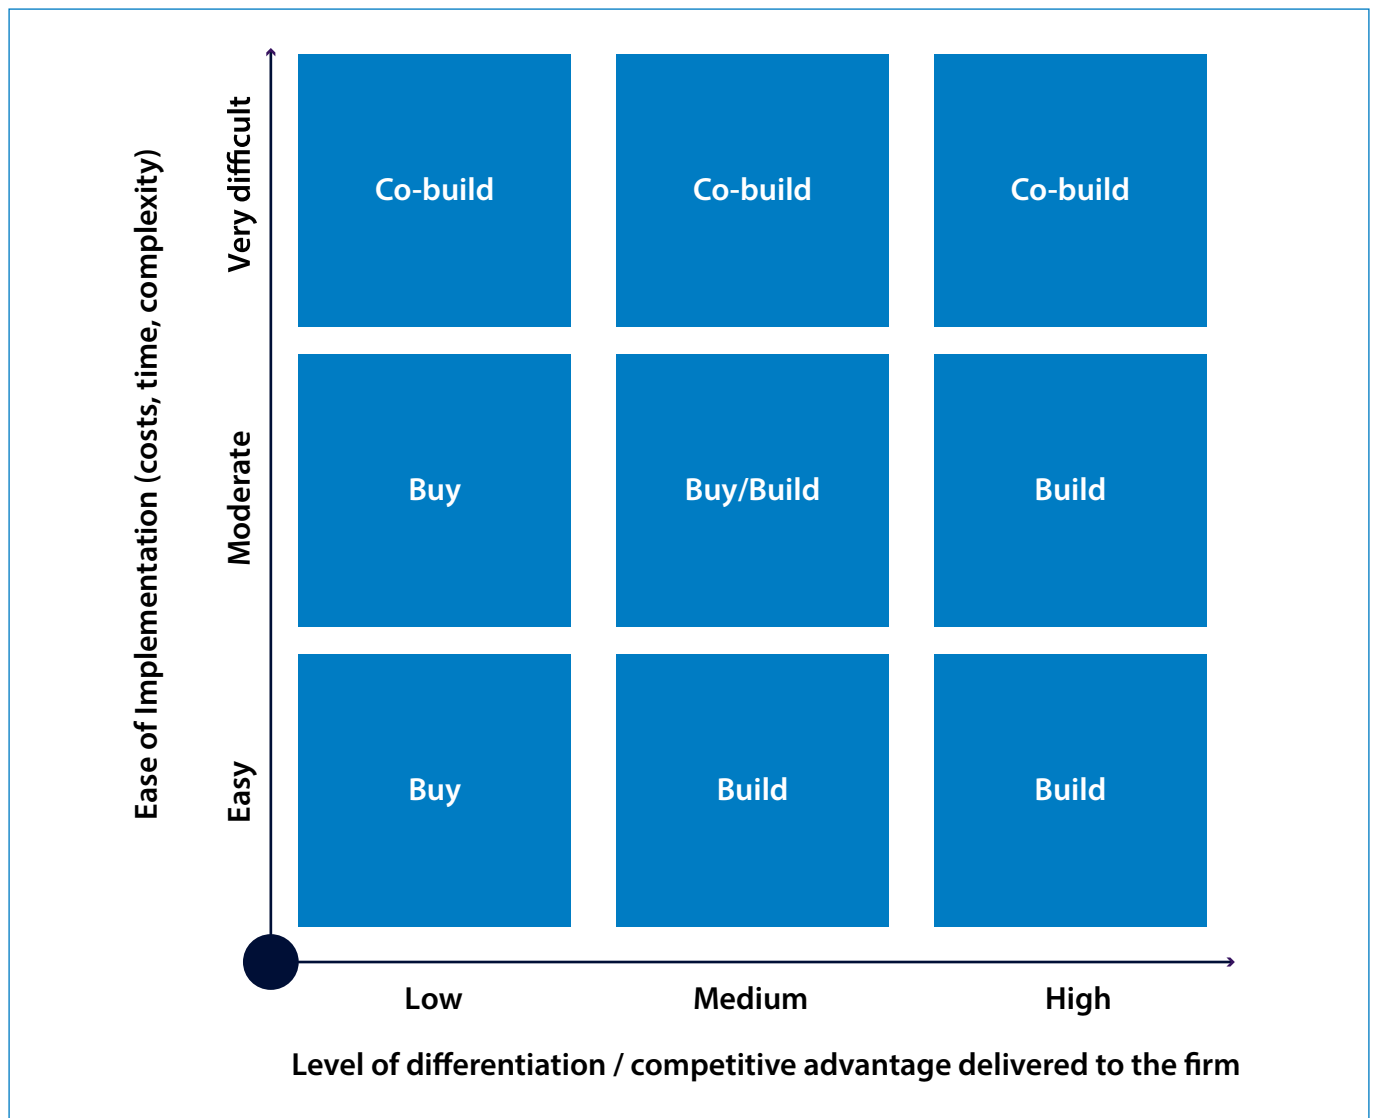


Figure 3: A Simple Decision Matrix for Buy vs Build

Policy Considerations for AI Adoption

While all applications of technology have potential policy implications, the potential of AI to replace or augment human decision making, as well as its need for large quantities of data, mean that it has drawn particular attention, even at this early stage of adoption, with legislators, policy makers, analysts, the media and the public all rightly assurance that AI will be used responsibly.

The full policy implications of AI are likely to take years to emerge, but financial services companies planning adoption should pay immediate attention to:

Governance

Financial services companies must consider how to adapt their existing governance mechanisms to address the specific features of AI, including not just governance of models and data, but of technology and operational resilience.

If AI is to augment or replace human decision making, it will require appropriate governance, and this governance must reflect the new and special features of the technology. At the same time, given the nature of their business, financial services

companies necessarily already have extensive governance mechanisms in place, and these mechanisms cover many of the areas to be addressed by AI. AI governance in financial services, therefore, is likely to be an extension and evolution of existing governance mechanisms.

Increasing technical dimensions to existing governance:

As technology has become increasingly pervasive in financial services companies, mechanisms to govern dimensions such as risk have had to pay increasing attention to technology and its implications. Adoption of AI will increase this need.

Further emphasis on operational resilience:

Business leaders and regulators within financial services have increased their focus on the ability of technology to provide the required level of operational resilience. As AI enables technology to perform an even wider range of functions, this focus will increase.

Increasing business dimensions to technology governance:

Technology governance within financial services companies is already attentive to the business implications of the services run by technologists. As increasingly sophisticated solutions such as AI are given increasingly important roles, the business dimensions of technology governance will become even more important, as will the role of technology leaders in ensuring that business leaders understand the technology on which they depend.

Introduction of ethical principles and panels:

In addition to existing governance structures, financial services companies must ensure that they are aware of the ethical implications of AI solutions, and the way in which data is used to build those solutions. Financial services companies should expect to develop ethical guidelines, and to implement governance mechanisms which can oversee both the implementation and evolution of those guidelines. This new aspect of governance is explored in more detail below.



Ethics & Regulation

Financial services companies must consider how they will ensure that AI is used in a way which is consistent with their values, and with the expectations of the people, companies, economies and societies they serve, particularly given the impact which financial services companies can have on all of them.

AI ethics is a young field which, until recently, had received more attention in academia than in industry. Just like AI technology, though, AI ethics and the questions it addresses are rapidly making their way out of universities and into the boardroom, as well as onto the agenda of legislators and regulators.

Today, the speed of AI adoption seems to be matched by the speed in concerns about its implications. Some of these concerns are premature, and are dependent on the achievement of pervasive human-level AI which we are unlikely to see for many years – if indeed this turns out to be possible at all. However, there are immediate, pressing, legitimate concerns about the use of the current generation of AI to take decisions which were previously taken by humans. While financial services companies do not face some of the immediate, difficult questions about safety of life which apply to industries such as transportation and health care, they do make decisions which affect the financial well-being of individuals, companies, economies and societies.

Emerging ethical questions about AI in financial services can be grouped into four big themes:

Bias

How can we be confident that a decision taken by an AI algorithm is fair and free of unwarranted bias? How do we detect bias within a model or within a dataset used to train a model?

Explainability

How do we meet our obligations to explain the reason for a decision when the decision was taken by a model, and when the model comprises weights on the connections between nodes in a neural network, rather than formulae or code we can examine and comprehend directly?

Privacy

How can we ensure that customers' privacy is protected when their data is being used to build and train models? To what extent is it legitimate to use a customers' data to build models that provide a general good, such as the detection of financial crime?

Control

How can we be confident that the model continues to behave and perform as expected, when it is continuously learning? How do we detect and respond to unusual behaviour when it is taking place at speeds beyond human capability?

Financial services companies should engage proactively with regulators as they develop guidelines for the use of AI, and be transparent about the ways they are using and planning to use AI.

It is unlikely that precise rules for the use of AI will emerge quickly, and they may not emerge at all. Instead, enterprises should consider the development of general principles which they can apply through their normal governance structures. They should create the time and attention to carefully consider those cases which don't obviously fit or contradict their principles, and remain open to revising those principles as they are challenged by tough cases.

It is also important for enterprises to recognize that existing, non-AI solutions are subject to the same ethical considerations. Human decision making can be subject to bias, can draw from a wide range of personal data, can be difficult for human decision makers to explain, and requires control through governance mechanisms, training and review. Similarly, traditional software can contain embedded bias, can operate across large datasets, can be difficult to understand, even when code or formulae are fully explicit, and often contains bugs or exhibits unexpected behavior in unexpected conditions.

Finally, as well as raising ethical concerns, the use of AI also has the positive potential to enhance and reinforce the application of ethical principles. For example, as we develop techniques to detect bias in data and decisions, AI can be used as an additional check on other forms of decision making, and to reveal unexpected prior bias.



Conclusion

AI has arrived in enterprises in a form which, while still in its early stages, has enormous promise, especially to an industry as rich in data, decision making and insights as financial services.

In order to realize this promise, companies must address practical considerations such as how they will acquire skills, make sense of the market and manage data, as well as policy considerations such as how they will govern AI, use it in a way which is responsible and consistent with their values, and participate in the development of regulation and industry guidelines. They must also recognize that it is impossible to do this alone: they will need to form strong partnerships with companies they have worked with for years, as well as companies and institutions they have never worked with. Realizing the promise of AI also means remaining positive and optimistic about this promise, despite legitimate concerns about risks. Used well, AI has the potential to enable financial services companies to serve their customers better, as well as to co-operate on long-standing industry problems, such as fraud, financial crime and financial exclusion.

The financial industry is currently in the first wave of AI adoption: navigating this wave and the waves to come will require a large degree of traditional human intelligence and imagination.



About the Authors



Andy Maguire

Group Managing Director and Group Chief Operating Officer, HSBC


Andy Maguire is Group Chief Operating Officer, HSBC, responsible for leading HSBC Operations, Services and Technology and the bank's Chief Operating Officers. He joined the bank in December 2014 and was appointed to the Group Management Board in August 2015. Andy began his career at Lloyd's Bank, where he worked in retail, corporate and private banking. He then worked as a consultant specialising in large transformation programmes in Europe, the Americas and Asia. He joined Boston Consulting Group in 1998. At Boston Consulting Group he was the Head of the Banking and Customer Services practices, Managing Partner of the UK and Ireland and a member of the firm's global executive committee. From 2010-14 he led Boston Consulting Group's relationship with HSBC.



David Knott

Chief Architect, HSBC

David is an experienced IT leader, architect and strategist with over twenty five years' experience of designing, delivering and running IT for companies in the banking, insurance, media, utilities and transportation sectors. As HSBC's Group Chief Architect he is responsible for the architecture and design of technology across the organisation. His focus is on making HSBC's architecture simpler, safer and easier to use, while getting maximum value from new technologies. David holds a PhD in Philosophy, specialising in Ethics, and is closely following the rapidly emerging field of AI Ethics.



Dr. Christof Mascher

COO, Allianz SE

Christof Mascher is a Member of the Board of Management and Chief Operating Officer of Allianz SE since September, 2009. As a Head of Allianz Operations he is responsible for driving Allianz' transformation into a global digital operating company by building, developing and deploying digital assets for the Allianz Group. Joining Allianz in 1989, he held various leadership positions in operations and IT within the Group. Christof focuses on the challenge of transforming altered/new customer demands into business opportunities. Christof Mascher holds an M.A. in Philosophy from the University of Vienna and a Doctorate degree in law from the University of Innsbruck.



Mohit Joshi

President, Infosys Ltd

Mohit Joshi is President and the Head of Banking, Financial Services & Insurance, Healthcare & Life Sciences at Infosys. He is also responsible for firm-wide sales operations and reporting processes, including large deal pursuits and top account growth. Mohit has over 20 years of professional experience working across the US, India, Mexico, and Europe. He is based in London.

References

1. http://www3.weforum.org/docs/WEF_New_Physics_of_Financial_Services.pdf
2. <https://www.loc.gov/law/help/artificial-intelligence/regulation-artificial-intelligence.pdf>
3. <https://enterpriseproject.com/article/2019/1/ai-skills-5-ways-build-talent-internally>
4. <https://www.forbes.com/sites/insights-intelai/2019/05/22/should-you-build-or-buy-your-ai/#6cd8ac7a441d>
5. <https://www.statista.com/statistics/607716/worldwide-artificial-intelligence-market-revenues/>
6. <https://www.forbes.com/sites/louiscolombus/2019/03/27/roundup-of-machine-learning-forecasts-and-market-estimates-2019/#c50cef97695a>



For more information, contact askus@infosys.com



© 2020 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.