



Reshaping Banking with Artificial Intelligence



HONG KONG MONETARY AUTHORITY
香港金融管理局





Table of Contents

Foreword	01
Executive Summary	03
1. Introduction to AI	05
1.1. What do we mean by AI?	07
1.2. A definition of AI	08
1.3. Data and computation algorithms	11
1.4. Recent AI developments	16
2. Technology foundations that enable AI deployment	23
2.1. Big Data infrastructure to store learning data	25
2.2. Efficient exchange of data with edge computing and 5G	26
2.3. Accessible cloud computing to enable AI algorithms	26
2.4. Improved economics and collaboration through open source technologies	27
3. Industry adoption of AI: drivers and barriers	29
3.1. The need to improve customer experience	31
3.2. The need to stay cost-effective	33
3.3. The need to better manage risk	34
3.4. Barriers to AI adoption	35
4. Examples of AI in banking today	41
4.1. Customer experience	45
4.2. Automation and quality control	48
4.3. Risk management	53
5. AI implementation	59
5.1. Key considerations in defining a high-level AI implementation journey	62
5.2. Key considerations in establishing a governance model	67
5.3. Measuring and assessing AI implementation	69
6. AI for Hong Kong banking: looking forward	71
6.1. Building an AI-ready workforce	73
6.2. Establishing an AI centre of excellence (CoE)	75
6.3. Fostering Hong Kong's banking AI community and infrastructure	76
6.4. Drafting guidelines on risk management in AI	77
7. Conclusion	79
Appendix A: Glossary of terms and definitions	81
Appendix B: Survey results	85
Appendix C: Summary of local FinTech-related university programmes	92
Acknowledgements	94



Foreword

Just a few years ago, AI was the stuff of science fiction. However, powerful new computer chips, cheap data storage and the emergence of neural networks – among other technologies – have pushed AI to centre stage as a readily available tool for business applications. In the last couple of years, a growing number of companies have started to adopt AI and experiment with how it might reduce their operating costs, increase efficiency, grow market share and improve customer experience. Potential applications of AI range from back-office operations to customer-facing chatbots, and investment advisory services. The potential benefits are enormous.

As the value proposition of a new technology gradually emerges, the risks and the legal and governance issues that it brings with it also need to be properly addressed. In this paper we aim to address these issues through an open-minded, in-depth examination of AI, from a technology-neutral stance.

Underpinning this paper is an industry-wide survey which looks at how far Hong Kong's banking sector has progressed in adopting this technology. We would like to thank The Chinese University of Hong Kong for its help in devising the questionnaire used in the survey.



This is the first in a series of papers looking into AI. It examines the technology in detail, shares a number of use cases being adopted in Hong Kong and explores some of the main findings from the survey. We hope that this and subsequent papers will offer Hong Kong's banking sector some useful reference points when considering the adoption of AI.

Lastly, we would like to express our gratitude to the banks and institutions that have participated in this paper by completing the questionnaire and spending time with us for interviews.

Edmond Lau
Senior Executive Director
Hong Kong Monetary Authority

Executive summary

Among the many new technologies to have had an impact on the corporate world in recent years, AI looks to be a critical disruptor for the banking industry. Apart from gaining traction in the corporate world, AI has been adopted as a key national strategy by countries such as China, the US, the UK, Canada and Japan¹. According to PwC's AI Impact Index², which looks at 300 AI use cases around the globe, the technology could contribute US\$15.7 trillion to the global economy by 2030, boosting the GDP of individual countries by up to 26%.

Financial services is one of the pillars of Hong Kong's economy, accounting for 7% of jobs and generating 18.9% of the city's GDP³. Seventy of the world's 100 largest banks operate in Hong Kong, securing its role as an international financial centre. Banking alone makes up 40% of all financial services employment and 63% of the value that financial services add to Hong Kong's GDP. Maintaining the competitive position of the banking sector is of fundamental importance to Hong Kong.

What does banking mean to HK?

 **70** of the 100 world's largest banks operate in Hong Kong

 **40%** of the 253,100 financial services jobs in Hong Kong are in banking

 **HK\$270.4B**
value added - amounting to 18% of Hong Kong's GDP

Why is AI important?

 **US\$15.7T**
potential contribution to the global economy by 2030

 **Up to 26%**
boost in GDP for local economies

¹ Dutton.T. (28 Jun 2018). An Overview of National AI Strategies. <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>

² PwC (July 2017). Sizing the prize: What's the real value of AI for your business and how can you capitalise? <https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf>

³ The Financial Services Sector in Hong Kong (April 2018). Hong Kong Monthly Digest of Statistics <https://www.statistics.gov.hk/pub/B71804FB2018XXXXB0100.pdf>



The Hong Kong Monetary Authority (HKMA) aims to foster discussion about AI and encourage its development across the broadest range of institutions, given Hong Kong's role as a major international financial centre. We have therefore commissioned PwC to publish this research, which has the following objectives:

- **To explain AI and summarise how it works** by laying out the basic concepts, describing the latest developments and defining the building blocks needed to adopt AI in a corporate environment.

- **To discuss potential AI applications in the banking sector** by sharing tried and tested use cases.
- **To look at the latest developments and implementation issues** as AI is adopted by the banking industry.
- **To promote and encourage adoption of AI** by raising awareness of the benefits and risks associated with the technology.

This research paper is targeted at:

- Operational and management-level staff of financial institutions, who may not necessarily have in-depth knowledge of AI.
- Stakeholders in banking digitalisation, such as:
 - AI and FinTech start-ups;
 - Technology incubators;
 - Regulators.

The HKMA-PwC research is based on:

- **Research** into academic and commercial publications on AI covering economics, public policy, regulation, technology and public use cases.
- **A Banking AI Survey** conducted in August and September 2019 on 168 HKMA-registered banks.
- **Deep Dive Interviews** with individuals from a start-up, two FinTech incubators, a research organisation and ten banks. The banks were selected to provide a balance of:
 - Multinational, Chinese and local banks;
 - Large, medium and small banks;
 - Established banks and banks that have been operating in Hong Kong for less than five years.

The research insights presented in this paper have been organised as follows:

- **Chapter 1** introduces concepts necessary to understand AI from both business and technical perspectives. It looks at the most important components and techniques in current AI development.
- **Chapter 2** analyses key technologies enabling the adoption of the current wave of AI.
- **Chapter 3** looks at the drivers for the adoption of AI in the banking industry, along with benefits and potential barriers to adoption.
- **Chapter 4** provides an overview of AI in Hong Kong's banking industry. It shares popular AI solutions being implemented by banks globally and defines where Hong Kong is on the spectrum of AI development.
- **Chapter 5** considers how to implement AI.
- **Chapter 6** looks to the future and recommends ways in which different stakeholders can help develop Hong Kong into an AI innovation hub.

1

Introduction to AI

Key Takeaways

Since the term 'AI' was first coined in 1956, its scope has evolved from the application of predefined logical rules to the performance of human-like cognitive functions. The goal of AI is to allow computers to mimic human intelligence so that they can learn, sense, think and act in order to achieve automation and gain analytic insights.

To mimic human intelligence, AI applications adopt two computation approaches: rule-based and non-rule based. Rule-based AI 'learns' using pre-defined rules and knowledge, and 'thinks' by inferring logical causes and effects according to 'if-then-else' rules. Non rule-based AI 'learns' with machine learning algorithms and 'thinks' using trained AI models.

Key enablers of AI development are the huge increase in the amount of data available and vastly increased computing power in both traditional computers and mobile devices, along with the ongoing development of machine-learning algorithms.

Some recent AI developments include natural language processing for better data pre-processing, deep learning and neural networks for learning complex rules, and collaborative learning models for model training.



1.1. What do we mean by AI?

Artificial Intelligence has progressed in fits and starts since the term was first coined in 1956 by the American computer scientist John McCarthy. In the 1980s and 1990s, AI was focussed on rule-based ‘expert systems’ that implemented predefined logical rules. Since the 2000s, data-driven AI has emerged as the main enabler of technologies. AI has now evolved to the point where it can interact and communicate with humans through analytics and automation. It can perform human-like cognitive functions (e.g. recognising patterns from data and understanding images, text or speech), as well as make predictions, recommendations and decisions.

In the banking sector, analytics and automation technology have been in place for several years. So how does AI differ from earlier waves of automation? To understand this from an industry perspective, we put the “what is AI?” question to several banks in Hong Kong. Some of their answers are given below:



AI is a new enabler for cost savings.

It is the technology that enables machines to perform business activities which humans can do today.

Compared to rule-based data analytics that we have been doing for a long time, AI is non rule-based business intelligence derived from analysing data.


It refers to using mathematical methods with predictive and decisive intelligence to substitute for humans in conducting repetitive tasks and seals the gap of various automation processes we have been doing over the years.

It is the future of banking, banks that fail to master AI will become obsolete.


It is clear from these responses that banks recognise the importance of AI but that individual banks, with their different strategic focuses and business functions, may not have a full picture of what AI is and what it does. To aid the in-depth discussion in the following sections, we first offer a definition of AI.

1.2. A definition of AI

In a nutshell, the goal of AI is to enable computers to mimic human intelligence so that they can learn, sense, think and act, in order to achieve automation and gain analytic insights.



While machines can ‘sense’ and ‘act’ in relation to humans and the environment through user interfaces, sensors and robotics, the core of AI is the ‘think’ component – e.g. making predictions, decisions or recommendations. To mimic human intelligence, AI applications adopt two computation approaches: rule-based and non-rule based.



Rule-based decision making, also known as ‘expert systems’, is one of the simplest and most common types of AI. It is commonly used by almost all financial institutions. It involves translating a fixed set of pre-defined rules and knowledge about decision making into a knowledge base and computer logic (a chain of rules in an ‘inference engine’⁴). Rule-based AI mimics thinking by inferring logical causes and effects according to a chain of ‘if-then-else’ rules in its knowledge base.

Rule-based AI systems became popular in the 1960s, and dominated the period from the 1970s to the 1990s, used especially for industrial robots performing repetitive tasks. Applications were limited however, as it can be very time-consuming and expensive to add new rules or knowledge to these systems so that they can react to the changing business environment. Also, in many scenarios it can be very challenging to explicitly define rules in a programmatic or declarative way. This led to a demand for, and the development of, non-rule-based AI.

Non-rule-based AI not only reacts and make decisions according to a pre-defined set of rules, it can also derive extra, meaningful insights by ‘learning’ automatically from its inputs.

Machine learning (ML) is one way to enable the ‘learn’ capability in AI. This involves using a set of learning algorithms driven by mathematical techniques which allow machines to learn from data, instead of being explicitly programmed to perform certain tasks. The training process uses the learning algorithm to derive relationships between data points from training data, which is commonly a subset of historical data. The outputs of the training are trained machine learning models, which can perform predictions or make decisions according to the data patterns observed from the input data, or from queries provided by users.

ML is able to identify subtle data patterns which cannot be easily described by humans, and extract insights from less structured data. As a result, ML has become the main technique driving the current wave of AI applications – from call centre voice analysis to autonomous vehicles.

However, ML in general is data and computing-heavy. In order to find data patterns or relationships and make accurate predictions, ML needs to process tremendous volumes of data. It needs to execute computing-intensive statistical techniques and mathematical algorithms for model training and for tuning and testing model selection.

According to a 2018 McKinsey report⁵, AI and machine learning have the potential to create an additional \$100 billion in the finance and IT sectors.

⁴ An inference engine is a tool used to apply logical rules to knowledge assets to deduce new information.

⁵ McKinsey (April 2018). Visualizing the uses and potential impact of AI and other analytics. <https://www.mckinsey.com/featured-insights/artificial-intelligence/visualizing-the-uses-and-potential-impact-of-ai-and-other-analytics>



1.3. Data and computation algorithms

The significant progress of both rule-based and non-rule-based AI systems has been driven by two key components:

1. The huge increase in the amount of data available
2. The availability of vastly increased computing power in both traditional computers and mobile devices, and the ongoing development of machine-learning algorithms.

1.3.1. Data

Data is vital to the ‘learning’ process of a non-rule-based AI system — from model training to testing. Even for rule-based AI, experts often perform data analysis and statistical hypotheses on sample sets of historical data using statistical techniques. Also known as ‘data mining’, this approach finds statistically significant data patterns when deriving the predefined rules to be programmed into the AI’s knowledge base.

Machine learning and data mining actually share many similar statistical models and algorithms. The former differs from the latter in that data mining is carried out by a knowledge worker on a particular data set with a specific goal in mind — namely to discover patterns and extract insights from the data set. By contrast, machine learning algorithms are carried out by computers automatically or semi-automatically.

Both the quality and volume of training data directly impact the predictive accuracy of most AI models. Proper data pre-processing, including collection, data profiling, cleansing, transformation, and labelling of training data, is critical to successful AI development.

Raw training data can come in various formats: including structured data (such as historical equity trading transaction records stored in a database), semi-structured data (such as social media feeds), and unstructured data — like audio, video and still images. ML algorithms require this training data to be cleansed and transformed into a defined format that the machine can easily read. Speech recognition technologies do this, for example, by converting voice recordings into machine readable texts that can be used to train a Chatbot. To derive insights from data relationships, data scientists may also need to assign meanings to training data based

on domain knowledge, i.e. data labelling. For example, if the task is to train an AI to classify scanned image documents, data scientists will need to label the type of document (i.e. labelled or output data) to each of the scanned images (i.e. input data).

Training ML models requires millions of pieces of labelled data. Pre-processing and labelling this data is labour-intensive and constitutes most of the work required in the ML model training process. As a result, the development of machine learning was slow in previous waves of AI. But with the recent advances in big data technology (as discussed in section 2.1.), large volumes of data can now be efficiently collected and processed.

According to a recent survey conducted by New Venture Partners (in which the majority of respondents were financial services firms), 21.1% of firms invested more than US\$500 million in Big Data or AI in 2019 (up from 12.7% in 2018)⁶. Investment in big data infrastructure over the past few years has led to the accumulation and analysis of granular, detailed data on a very large scale. This data foundation is providing the ‘fuel’ to power AI training.

How does AI derive insights intelligently from the training data? This requires advanced mathematical learning techniques enabled by efficient computing infrastructure and data.

1.3.2. Computation algorithms

Non-rule-based computation, particularly machine learning algorithms, has become the main way of powering the processing intelligence core of AI. This is due to its flexibility in enabling the machine to autonomously learn from sample data sets and then make predictions without explicit pre-programmed instructions. The power of machine learning is that models trained by machine learning algorithms can also iterate and improve their accuracy as they learn from new injections of training data⁷. This has become the foundation of different AI capabilities that mimic the ‘sense’ (e.g. computer vision) and ‘act’ (e.g. natural language generation) capabilities of humans.

There are four types of machine learning algorithms: supervised, unsupervised, semi-supervised and reinforcement learning.

⁶ NewVantage Partners (January 2019). Big Data and AI Executive Survey 2019. <https://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings-Updated-010219-1.pdf>

⁷ Rao, A. (25 July 2017). Demystifying machine learning part 1. <http://usblogs.pwc.com/emerging-technology/demystifying-machine-learning/>

Supervised learning is a type of machine learning algorithm in which humans guide the system by labelling the relationship between every variable of input data and every variable of output. For example, a supervised machine learning model can learn which emails are classified as spam by first having a real person go through thousands of emails to identify those which are spam and then label them. Techniques such as linear regression, logistic regression and decision tree classification are used to train supervised learning models.


An **unsupervised learning** algorithm, by contrast, identifies patterns in a set of unlabelled input data automatically. In other words, rather than learning the relationships between a given set of input and output data, unsupervised learning algorithms are designed only to learn whether data patterns or structures exist. One of the most common use cases is to identify different customer segments with different preferences and demographics.


The third type of machine learning algorithm is **semi-supervised learning**. It is used on data sets which are only partially labelled. As it can be costly and time-consuming to label large data sets, semi-supervised models are commonly applied to real-world machine learning problems. Semi-supervised learning combines supervised and unsupervised learning algorithms to either infer the pattern or structure of input variables, or make predictions about the unlabelled data and feed those predictions back to the model as training data.

The last type of machine learning algorithm is **reinforcement learning**. This determines optimal behaviour based on feedback from its environment. Reinforcement learning models incentivise machines to optimise behaviour by providing positive feedback, or 'rewards', which teach the machine to strive towards a preferred behaviour. In the absence of pre-existing training data, a reinforcement learning algorithm must learn to predict the consequences of its actions through trial and error. It therefore learns from its own prior experience rather than from a predetermined set of examples. Reinforcement learning is commonly used to automate the decision-making process in interactive scenarios.

Of all the machine learning types described above, 81% of the AI use cases that our survey and research studied involved supervised learning techniques. Each bank that was interviewed had applied a supervised learning solution to at least one use case, making this the most common type of AI adopted by our respondents.

While supervised learning is the most widely used learning algorithm, it is also common to see different algorithms used in conjunction to make more accurate predictions in complex scenarios.






Below is an example of how different algorithms work together in a complex scenario.

How do supervised learning, unsupervised learning and reinforcement learning work together in banking operations?


Most traditional anti-money laundering (AML) systems are rule-based and identify a substantial number of suspicious transactions. Significant human effort is then required to determine whether any of the transactions identified as suspicious are in fact false negatives (transactions identified as suspicious but which are actually not). Machine learning can be applied to identify false negative money laundering transactions and save human effort.



This problem can be solved with AI. With labelled historical bank transactions as training data, banks can use supervised learning to train a model to identify whether a transaction is a false negative or not. Given the nature of **supervised learning** (where patterns learnt by the model are based on historical labelled data), continuous monitoring and calibration of the model's performance is required to ensure that its results remain aligned with changes in regulations, policies, and bank procedures, and to enable the model to identify new patterns that are not available from historical data.


The continuous data tagging and monitoring and the calibration of the model's performance are performed by a 'robot trainer', as explained by some of the banks we interviewed. Such activities not only require knowledge of data and AI technologies, but also some specific financial domain expertise, such as trade finance. The recruitment and retention of experts in these areas can be a burden for banks. We recommend that banks starting the supervised learning journey should not only consider the technologies involved, but also the implications for change management, organisational structure, staff skill sets and ancillary costs.

⁸ Magnim (10 March 2019). Connection between Data Science, ML and AI. <https://becominghuman.ai/connection-between-data-science-ml-and-ai-d1c18d89b0bd>



Banks can also apply clustering and **unsupervised learning** methods without human labelling to identify new relevant transaction patterns or features for events such as new account opening, loan applications and various other banking transactions. With clustering algorithms, common transactions are grouped into clusters with similar behaviours. These are set as a baseline of normal transactions. Anomalies can then be identified as outliers that do not have a strong membership in any cluster group.

When an anomaly-based fraud detection model is trained, the system can recognise and automatically notify the user of any deviations from the normal pattern. Operational staff can then confirm whether the transaction is legitimate. Any false alerts can be set as feedback, enabling the machine to gain greater accuracy through **reinforcement learning**. Banks that have deployed machine learning-based systems report that the false positive rate for flagging money laundering risks has been reduced by 20% compared to previous systems⁹.



⁹ Faggella, D. (12 December 2018). Bank Reduces Money-Laundering Investigation Effort with AI. <https://emerj.com/ai-case-studies/bank-reduces-money-laundering-investigation-effort-with-ai/>

1.4. Recent AI developments

Having discussed the importance of data and computation algorithms for AI, we next look at how recent AI developments have brought the application of these two key components to a new level.

While traditional AI algorithms, particularly machine learning algorithms, are robust and effective, they require data scientists and domain experts to pre-analyse and break down problems into small chunks that can be fed into the learning process (with specific algorithms, data and features).

Fortunately, the recent emergence of big data and cloud computing (discussed in section 2) has increased AI's autonomy and removed some of the constraints associated with the need for human domain experts. As a result, many AI use-cases in banking have now become feasible.

We have summarised below three recent AI developments arising from the emergence of big data and cloud computing:

1. Natural language processing
2. Deep learning
3. Collaborative learning models

1.4.1. Better data pre-processing with Natural Language Processing

Machine Learning can only detect patterns in raw data effectively if the algorithm can read and interpret the data. Natural Language Processing (NLP) is a branch of AI that aids computers to understand natural human language, which is usually relatively unstructured and contains 'implicit' meanings. Differences in authors' writing styles or cultural backgrounds can also affect how the content of natural languages is interpreted. Data scientists could manually tag the meaning of each sentence as part of data pre-processing, but this would be too time consuming with the large volume of training data. NLP can be used in conjunction with voice-to-text technology (which transforms human voice to machine readable text) and computer vision technology (which extracts text from images or videos) to pre-process human language and tag the meaning automatically before applying a general machine learning algorithm to find patterns.

NLP uses different techniques to learn and understand language, enabling computers to understand not only single words and word combinations, but also grammar and hence written or spoken sentences. NLP categorises, analyses and comprehends the following elements using a combination of statistical analysis and rules pre-defined by linguistics experts:

- The language used, such as spelling, grammar, meaning and connotations of words
- The content of the text – namely the message the person wants to convey (i.e. logical analysis)
- The person behind the text, including their style and feelings (i.e. sentiment analysis)

In short, NLP has the ability to recognise linguistic patterns and interpret meaning in languages, which it achieves through the use of linguistic analysis tools such as part-of-speech tagging, automatic text summarisation, and entity and relationship extraction. NLP is widely used to process large volumes of unstructured text in big data to reveal patterns, to identify and tag relevant information and to summarise document contents – all with minimal human input.

Traditionally, NLP methods have mostly employed 'shallow' machine learning models, which use a limited set of input data features defined by human experts and based on their domain knowledge. This is because some linguistic information can be represented in many different ways and it can be challenging to derive accurate patterns from a multiplicity of features (the so-called 'curse of dimensionality'). With the growing use of deep learning, neural-based NLP models have achieved superior performance compared to traditional models – without the need for task-specific feature engineering.

NLP can satisfactorily process major languages (e.g. English and Mandarin). It can also collaboratively work with computer vision and voice-to-text technologies for applications and use cases such as real-time translation, regulatory requirements analysis, legal document reviews, compliance reviews, intent parsing for Chatbots, customer call record analysis, and sentiment analysis for social listening.


The mixed language environment in Hong Kong (where speakers often use English, Mandarin and Cantonese

interchangeably) and the relatively small population of Cantonese speakers has created barriers for the wider local adoption of NLP. Research and development by Hong Kong-based technology firms such as Fano Labs¹⁰ and organisations such as the Hong Kong Applied Science and Technology Research Institute (ASTRI) have addressed these challenges in recent years. This has helped put NLP centre-stage among AI applications in Hong Kong.

1.4.2. Complex rules and deep learning

Traditional Machine Learning usually requires data scientists to select key data features for prediction, as it is not efficient to run the learning algorithm with too many parameters. However, some subtle patterns may thus be missed and prediction accuracy may suffer as a result – especially with highly unstructured data, such as images or voice.

Deep learning, a subset of machine learning, represents a breakthrough in the field of AI. As the name suggests, deep learning ‘learns’ by modelling hundreds of thousands of deep neural network layers, inspired by the human brain, which can have billions of permutations. Deep learning is powerful enough to recognise and learn subtle structural patterns from vast amounts of data, enabling it to recommend decisions or make predictions. Deep learning requires a huge amount of computing power and data to train a model with sufficient predictive power. Its applications were therefore limited until the big data era of the last decade. Since then, deep learning has gained substantial traction.



¹⁰ Lam, A. (11 July 2019). Fano Labs: Multilingual ASR and NLP. <https://www.miotech.com/en-US/insights/article/5d26dd98d0728f0044b21639>


Neural networks (NNs) are the building blocks of all deep learning techniques, and are inspired by the structure of the brain. The basic unit of computation in a neural network is a neuron. A neuron calculates an output based on information from other neurons or from an external source. Through the hidden layer architecture formed by connecting neurons, deep learning extracts high level features and learns classification incrementally with each additional layer of neurons. Deep learning algorithms are difficult to explain because each hidden neuron represents a non-linear combination of all the previous neurons.

With the explosion of data in the big data era, huge amounts of structured and unstructured data — including images, voice data and digital footprints — have been captured in company databases or on the Internet. Learning algorithms trained using big data can discover connections between subtle or complex features of data that humans cannot easily find or describe logically.

The latest advances in unsupervised deep reinforcement learning from DeepMind's AlphaGo Zero research¹¹ show that in certain situations AI can be surprisingly powerful, even without initially inputting and labelling training data. AI algorithms have already demonstrated the ability to outstrip the performance of even the most experienced humans in speech recognition¹², facial recognition, playing games such as chess and Go¹³, and detecting loan or credit card fraud.

One important difference between deep learning and traditional machine learning is the performance as the volume of data increases¹⁴. Deep learning algorithms do not perform well when the data volume is small. With greater data volume, deep learning algorithms can outperform traditional machine learning algorithms.

Performance of machine learning technique vs amount of Data



¹¹ Silver, D., & Hassabis, D. (18 October 2017). AlphaGo Zero: Starting from scratch. <https://deepmind.com/blog/article/alphago-zero-starting-scratch>

¹² Li, A. (1 June 2017). Google's speech recognition is now almost as accurate as humans - 9to5Google. <https://9to5google.com/2017/06/01/googlespeech-recognition-humans/>

¹³ Silver, David, et al. (19 October 2017). "Mastering the Game of Go without Human Knowledge", Nature 550, 354–359. <https://www.nature.com/articles/nature24270>

¹⁴ Analytics Vidhya. (8 April 2019) Deep Learning vs. Machine Learning. <https://www.analyticsvidhya.com/blog/2017/04/comparison-between-deep-learning-machine-learning/>

Below are some of the common applications of narrow AI enhanced by big data and deep learning in banking:



Handwriting and voice recognition

Machine learning techniques are widely used to recognise handwritten words, images and voice recordings.

Traditionally, the detection of written characters has been conducted using rule-based OCR (optical character recognition) to find fixed character patterns. This approach can recognise printed characters well but not handwritten characters, which have many varied forms. Using deep learning to recognise structural patterns that may have many variations, computers can now read handwritten characters accurately and convert them into machine-encoded forms.

Similarly, the refined process of deep learning has now made voice recognition accurate enough to be applied outside a carefully controlled environment. According to our interviews, speech recognition has reached 95% accuracy and is gradually turning into a key method by which humans can interact with computers, for example through voice-instructed Chatbots or robot assistants.



Natural language generation

Natural Language Generation (NLG) is an automated AI process that generates custom content by transforming structured data into natural language. The aim of the algorithm is to convey a message by accurately predicting the next word and forming a sentence that is readable by humans.

Early attempts at NLG involved predicting the next word in a sentence based on the current word. However, there was a high degree of inaccuracy, as this method ignored the wider context and the structure of preceding sentences.

A new deep learning-based NLG approach uses the Long Short-Term Memory neural network (LSTM). This involves the neural network model building and remembering the correlations of multiple words after considering the sequences of all the key words in the previous sentences. This significantly increases the accuracy of sentences generated because the word predictions take context and structure into account. Some NLG approaches can further predict the customer's emotional profile based on previous conversations, or the content of marketing campaigns. They can then personalise the message and style of the language generated using a predefined database that includes words tagged to reflect different emotional responses.

Some banks have adopted deep learning based on NLG to generate personalised e-marketing materials, and have successfully doubled the click rates on digital advertisements for credit cards and mortgages loans¹⁵.

The above examples show that deep learning has had promising results in processing unstructured data by automating the process of extracting features from the datasets. Research is being actively conducted to improve the efficiency of deep learning algorithms and make them fit for

general purpose. More practical applications are expected to be enabled by deep learning once banks accumulate enough data to meet the heavy data requirements for training deep learning models.

¹⁵ Cheng, M. (7 August 2019). JPMorgan Chase has found software that writes better ads than humans can. <https://qz.com/work/1682579/jpmorgan-chase-chooses-ai-copywriter-persado-to-write-ads/>



1.4.3. Model training with collaborative learning models

Conventional machine learning and deep learning models are not designed to be trained collaboratively by multiple parties. They are also privacy-intrusive, as the training data set needs to be stored within a centralised database and be aggregated using servers within one company's data centre. The centralised nature of conventional model training has limited AI development, as no single organisation has access to all available data or computing power.

Federated Learning, recently adopted by tech pioneers such as Google¹⁶, is a new machine learning approach that allows organisations to collaborate on machine learning model training without giving away proprietary data. Conceptually, federated learning enables a shared machine learning model to be trained by multiple machines, based on distributed data sets. In practice, both the training and data reside in a local machine, but the results are encrypted and shared with collaborating machines.

Eliminating the need to store all the training data in a centralised database addresses growing concerns about data privacy, especially in light of new privacy legislation, such as the EU's General Data Protection Regulation (GDPR)¹⁷. Corporates such as WeBank have expressed interest in the development of federated learning, and WeBank has contributed its Federated AI Technology Enabler (FATE) framework for use by the open source community so as to assist other banks and corporations to onboard the technology¹⁸.

Federated learning can also be applied to consumer devices, which are potentially powerful machines for co-training a model using local data sets. According to statistics from Ericsson¹⁹, the total number of smartphones being used around the world had reached almost 5.6 billion as of June 2019. The machine learning process could become much faster and cheaper if all the unutilised computing power on smart phones was contributed via federated learning.

¹⁶ McMahan.B. & Ramage. D (6 April 2017). Federated Learning: Collaborative Machine Learning without Centralized Training Data. <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>

¹⁷ General Data Protection Regulation. <https://eugdpr.org/>

¹⁸ Federated AI Ecosystem. Collaborative Learning and Knowledge Transfer with Data Protection. <https://www.fedai.org/>

¹⁹ Ericsson (18 September 2019). Ericsson Mobility Visualizer: Interactive mobility data app. <https://www.ericsson.com/en/mobility-report/mobility-visualizer?f=1&ft=1&r=1&t=8&s=1&u=1&y=2019,2024&c=1>

²⁰ Patki, N., Wedge, R., & Veeramachaneni, K. (2018). The Synthetic data vault. Data to AI Group. <https://dai.lids.mit.edu/wp-content/uploads/2018/03/SDV.pdf>

²¹ Edge computing is a distributed, open IT architecture that features decentralised processing power, enabling mobile computing and Internet of Things (IoT) technologies. In edge computing, data is processed by the device itself or by a local computer or server, rather than being transmitted to a data centre. <https://www.hpe.com/hk/en/what-is/edge-computing.html>

Apart from federated learning, two other trends are important for the future development of AI. These are **Synthetic Data Creation** and **Transfer Learning**.

Synthetic Data refers to data that has been mimicked to include statistical properties that are similar to those contained in real data collected from actual business transactions. Since synthetic data is generated, there is no limit to the amount of data or scenarios that can be derived. This addresses a key challenge to AI development. Synthetic data is also not subject to data privacy concerns or modelling constraints, such as missing values and non-responses.


The ability to generate synthetic data is particularly important in use cases such as fraud detection, where there may not be enough real fraudulent data and related patterns available due to a bank's limited transaction volumes for a particular type of fraud or its location.

Recent research by MIT²⁰ shows that the prediction accuracy of machine learning models based on synthetic data can be just as good as that based on real production data.

Transfer Learning is another important machine learning technique that allows pre-trained models in a particular domain to be used by others as a starting point for developing

a new model in a similar domain. This accelerates the whole training process and performance of a new and related model by eliminating the need to retrain the same model from scratch. The insights which a data scientist has spent a tremendous amount of time, effort and proprietary data to develop can thus be passed on to other teams developing a model in a similar domain, resulting in an accelerated timeline and reduced effort. This machine learning method represents a new way for AI researchers and corporations to share knowledge accumulated in machines.

Efficient data exchange enabled by combining 'edge computing²¹ and high speed networks means that recent developments in collaborative learning systems hold a lot of promise for the banking industry. There have been a few successful pilot implementations of transfer learning and other collaborative learning in multiple scenarios. These include learning mobile user behaviour and training a common fraud detection system. Once transfer learning and federated learning mature, AI models may be able to learn more quickly in a collaborative manner, solve challenges for banks, and help the industry by building learning models from cross-domain data without compromising data security and privacy. This will be particularly useful for recently established banks, which will have little or no historical data.



2

Technology foundations that enable AI deployment


Key Takeaways

The four technologies that enable AI development are big data infrastructure, edge computing & 5G, cloud computing, and open source technology.

Big data infrastructure has delivered a huge boost to AI by efficiently storing and pre-processing learning data for AI training. 5G has cut down data transmission times and made data exchange much more efficient, while edge computing

allows machine learning to take place on devices like smartphones without the need for a cloud. Accessible cloud computing is enhancing the total computing power available for AI. Open source technologies make contributions from AI communities and major tech giants available to all, lowering costs and enabling software innovation.





The pace of AI deployment has been far from steady. While there has been a great deal of government-funded research, such as the 'AI Next Campaign' carried out by the Defense Advanced Research Projects Agency (DARPA)²², AI remained in the confines of academic research for many years. Insufficient computing power and a lack of freely available data were brakes on AI's progress.

In the last few years, however, AI has grown from a simple rule-based expert system that implements pre-defined logical rules to one that feeds data-driven, decision-making applications. This has not only been due to the advanced research that has been carried out on the underlying AI algorithms, but also to other related technologies.

This section mainly focuses on the key technology foundations that underlie the current and future waves of AI development. By keeping track of trends in these technologies, banks can better plan for AI deployment.

2.1. Big Data infrastructure to store learning data

As mentioned in Section 1.3.1, the quality and volume of available data are major determinants of AI development. The proliferation of computers, mobile phones, smart home appliances, Internet of Things (IoT) devices and social media means that data volume is increasing at an ever faster pace, while data structure has become more complex. According to the sixth edition of Domo's *Data Never Sleeps* report²³, 2.5 quintillion bytes of data are estimated to be created by the Internet every day.

Data scientists have always described big data as having five distinct dimensions: Velocity, Volume, Value, Variety and Veracity (The Five V's). Velocity refers to the speed at which data need to be analysed while it is being generated. Volume is the amount of data generated each second. Value is shorthand for the value of the data being collected and analysed. Variety means the number of data types, which are usually unstructured. Veracity is the quality and accuracy of the data. To derive meaningful insights, big data needs to be processed with advanced analytics – so-called 'big data analytics'.

²² DARPA. AI Next Campaign. <https://www.darpa.mil/work-with-us/ai-next-campaign>

²³ Marr, B. (21 May 2018). How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read. <https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#670d494360ba>

Compared to traditional data management tools, Big Data analytics infrastructure is typically designed from scratch to divide time-consuming data processing tasks for billions of records into manageable pieces. It utilises parallel computing processors and distributed data storage across hundreds of servers to shorten the data analytics process from days to minutes.

In the past decade, banks have made significant investment in big data analytics infrastructure for efficient data collection, processing and management. This has laid a foundation that is generating more 'fuel' for training machine learning models, resulting in better insights and more accurate results.

2.2. Efficient exchange of data with edge computing and 5G

Efficient exchange of data is another enabler of successful AI deployment, especially for real-time AI applications. Edge computing technology and 5G cellular networks can distribute data with significantly reduced communication times for data exchange.

Edge computing is done at or near the source of data, rather than relying on centralised data centres for processing, thus reducing communication turn-around times. With edge computing, most processes can be done locally on personal devices, such as cell phones or self-driving cars because the edge devices have sufficient computing power in their own right. Some are even equipped with chips specifically designed for running AI algorithms.

Edge computing has a number of benefits that are boosting the deployment of AI. It mitigates the problem of latency and Internet bandwidth (such as the delay in response time in cloud computing) by reducing reliance on centralised servers for data processing. It also enhances security and data privacy, as local devices can pre-process data and only share selected pieces with the cloud. This means that wearable devices or sensors have become more 'intelligent' in that they know what information to keep, what to transmit and how to use the information they have. For example, an iPhone with Face ID can perform facial recognition AI using the device's own processor without sending sensitive biometric information to cloud servers.

5G – the fifth generation of cellular network technology – is complementary to edge computing. The speed of 5G is expected to be 2Gbit /s, a hundred times faster than its predecessor 4G²⁴. There are also significant improvements in latency, from 40ms for 4G to 1ms for 5G. Fast internet speed and short latency allow multiple edge devices to exchange data much more efficiently, and thus to participate in machine learning model training collaboratively.

The deployment of AI has been accelerated thanks to the efficient exchange of data enabled by edge computing and 4G or 5G high-speed networks. Part of the machine learning process can be run directly on edge devices to power real-time AI applications, such as facial recognition. This allows smart devices to provide insights and predictive analysis in real time, which helps to create better customer interactions. High efficiency in data exchange also enables collaboration in machine learning model training, as it splits the training process across many individual local devices.

2.3. Accessible cloud computing to enable AI algorithms

Cloud computing was a precursor and facilitator of the emergence of big data. It refers to the provision of computer resources, such as computing power, networks and data storage, that are shared by users on an on-demand basis. The resources offered by cloud service providers are typically powered by massive server infrastructures across multiple data centres around the world, made available to a huge number of users at low cost due to their economies of scale.

Cloud computing technology has the potential to change the way companies do business. According to a PwC survey²⁵, companies that integrate the cloud into their business strategy benefit from improved time-to-market, flexibility to meet changing customer needs, and ease of integrating with existing infrastructure. The same benefits apply to AI model training. In almost all the interviews we conducted on this topic, banks also saw cloud computing as a key enabler for adopting AI.

Modern ML algorithms are data and computing-intensive. The cloud's scalability in terms of computing power helps shorten the training time for AI models by enabling them to

²⁴ Elkhodr, M., Hassan, Q.F. & Shahrestani, S. (10 October 2017). Networks of the Future: Architectures, Technologies and Implementations. <https://www.crcpress.com/Networks-of-the-Future-Architectures-Technologies-and-Implementations/Elkhodr-Hassan-Shahrestani/p/book/9781498783972>

²⁵ PwC (2011). Cloud Computing Navigating the Cloud. <https://www.pwc.com/cl/es/publicaciones/assets/navigating-the-cloud-final-3-24.pdf>

process a tremendous amount of data and to run algorithms ‘in the cloud’. Its on-demand flexibility also lowers the effort and lead time required to set up relevant infrastructure. These characteristics of cloud services suit financial institutions well, as many are still at the early, exploratory stages of their AI journey.

2.4. Improved economics and collaboration through open source technologies

Open source technology is another key enabler for organisations wishing to deploy AI. This decentralised software development model has been adopted by most of the technology giants. Microsoft, Google, Facebook and Amazon have all made remarkable progress in developing AI systems, and they have released much of their work to the public for free in recent years. China’s AI giant Tencent has also donated its federated AI framework, Acumos, to the open source initiative.

Unlike proprietary software products, which usually have a price tag for initial licensing as well as regular upgrades and maintenance, open source software products are often free to download and modify, thereby lowering the cost of ownership. While cost is one benefit, open source also provides a solid foundation for innovation. By bringing together large knowledge pools and significant expertise²⁶, open source communities collectively contribute to the continuous development and improvement of software.

In banking, open source development has provided data scientists and in-house technical teams with many handy off-the-shelf tools. These include Tensorflow²⁷ (from Google) and PyTorch²⁸ (Facebook), which can be used to quickly apply the latest machine learning technology to business problems.

Open source software also lowers costs and reduces risks for banks in their adoption of AI. The open source ecosystem widens the adoption of practical use cases, which can be a useful reference point for banks wishing to improve existing models that they have developed in-house. By riding on the core AI technologies (e.g. deep learning algorithms) contributed by the open source community, start-ups and vendors in the banking industry can reduce development costs and focus more on business applications. With cost savings and risk reduction, the adoption of AI applications in the banking industry can be accelerated.

²⁶ Carey, S., & Macaulay, T. (2 July 2019). What are the advantages of open source software? <https://www.computerworld.com/article/3412269/what-are-the-advantages-of-open-source-software-in-the-enterprise-.html>

²⁷ TensorFlow. Why TensorFlow. <https://www.tensorflow.org/about>

²⁸ PyTorch. From Research to Production. <https://pytorch.org/>



3

Industry adoption of AI: drivers and barriers

Key Takeaways

A number of drivers of and barriers to AI adoption were identified from our survey (see Appendix B).

The three drivers of AI adoption identified are the need to improve customer experience, the need to stay cost-effective, and the need to better manage risk. In Hong Kong, 82% of customers have experienced at least one problem or frustration with their banking experience. Banks can leverage AI to improve customer service by providing services like personalised recommendations and by delivering better customer support. At the same time, AI can also help banks to reduce costs through functions such as task automation. Finally, AI can simplify the risk management process and

develop more comprehensive insights into customer risk by taking data such as transaction history, market trends and customer credit history into consideration.

Three major barriers to AI adoption were also identified. The first is lack of explainability, which means that results or recommendations generated by AI are difficult to explain. A lack of resources, including AI talent and quality data, is another obstacle. Furthermore, as banking is a highly regulated industry, banks may face additional costs and risks in adopting AI due to the need to comply with AI-related regulations.



Our survey (see Appendix B) found that AI applications are on the increase across the Hong Kong banking sector – especially in retail banking. 89% percent of retail banks have adopted or plan to adopt AI applications; the total amount of capital investment by Hong Kong's retail banking sector is expected to increase by 70% in the coming five years.

Since 94% of banks that have adopted AI say that they will use it to shape their corporate strategy, it is important to understand the reasons driving this increasing adoption. In the following sections, these drivers are categorised and explained under three main headings:

- The need to improve customer experience
- The need to stay cost effective
- The need to better manage risk

3.1. The need to improve customer experience


As many business processes are digitalised and streamlined, customers are developing higher expectations of products and services. According to PwC's 2019 Survey *Virtual Banking: Customers Take Charge – Are You Ready?*²⁹, 70% of customers surveyed in Hong Kong are interested in more personalised and relevant information and tools. Below are two of the main expectations of these customers:



²⁹ PwC (November 2019). Virtual Banking: Customers Take Charge – Are You Ready? <https://www.pwch.com/en/financial-services/publications/2019-digital-banking-customer-survey.pdf>

These results suggest that consumer preferences are shifting to personalised digital services that are able to combine all their product and servicing needs, including touch points such as ATMs and branches, while also adjusting to the customer's different life stages. However, not all traditional banks are able to meet these preferences, leading to customer frustrations. Reported problems include long branch waiting times, poor online experience, and lack of any consolidated view of their finances. From the same survey, we found that the percentage of customers in Hong Kong experiencing at least one problem with their bank was 82%:

Customers experiencing at least one problem or frustration with their bank



These frustrations are among the reasons for the rapid development of FinTech companies, which are more capable of delivering a personalised experience to customers through digital solutions such as AI. For example, Alipay, the Chinese Fintech giant, has been using AI for credit scoring, smart customer support and fraud detection³⁰.

According to *The Future of FinTech in Hong Kong*³¹, a report issued by the Hong Kong Financial Services Development Council, FinTech companies are providing millions of people with access to financial services for the first time, largely due to the increased use of mobile phones and the decreased cost of servicing customers. Without the barriers of legacy systems to contend with, FinTech companies have the agility to define new customer digital journeys using emerging technologies, including AI. Traditional financial institutions are, therefore, being severely challenged in terms of providing customised services and innovative products.

Banks are recognising that they can raise their competitiveness against FinTech companies by better understanding their customers' behaviour, preferences and product needs. Some are looking to utilise their internal customer data and combine it with external data to form data partnerships that can generate detailed insights into customer behaviour, enabling them to provide personalised product recommendations to customers based on customer profile, transaction history and other data. This is one of the key drivers of their adoption of AI and big data analytics.

Open banking initiatives are further unleashing the possibilities of AI to enhance customer experience. These initiatives involve opening up internal bank data and helping banks connect and collaborate with third-party service providers, such as data or product providers, through an Application Programming Interface (API). APIs allow the secure sharing of customer-authorised financial data, and the distribution of partner-based products to bank customers. By also utilising big data analytics, banks can design unique, personalised products that enhance the customer experience.

Such initiatives are giving customers a more integrated banking experience across business applications, and are also enabling innovation and cross-industry business collaboration. Recently, a global bank announced a free personal financial management tool for the public in exchange for data from third-party service providers that would help it to enhance the tool's performance³². By aggregating data from separate sources, the bank expects to increase consumer engagement and put itself in a better position to proactively offer personalised products to suit customers' needs as and when required.

To remain competitive, there is a need for banks to provide a better experience to their customers in order to maintain and grow their market share. AI can help banks provide more flexible and comprehensive solutions to customers, thus creating a more personalised and unique customer journey.

³⁰ 光明网 (June 2017). 马云：支付宝从诞生起就尝试人工智能与人携手合作。 http://it.gmw.cn/2017-06/29/content_24934441.htm

³¹ FSDC (May 2017). *The Future of FinTech in Hong Kong*. https://www.fsdc.org.hk/sites/default/files/FSDC%20Paper_FinTech_E.pdf

³² Courbe, J. (14 May 2019). Building 'Open Banking' on a Platform of Trust. <https://bankingjournal.aba.com/2018/06/building-open-banking-on-a-platform-of-trust/>

3.2. The need to stay cost-effective

Given an ever-changing market and keen competition, it is vital for banks to find faster, more economical, and lower-risk approaches that will enhance their cost efficiency and overall profitability. This is another strong driver for banks to seek help from emerging AI solutions, which make large-scale automation of business processes possible.

After decades of work in digitalising the banking sector, use cases focusing on customer-facing and front-end operations, such as online banking, are fairly sophisticated. However, back-office functions such as risk management, compliance and technology have progressed more slowly. Most are still people- and paper-intensive. Paperwork has always been a big concern, as it is one of the main contributors to operating costs. Manual processing is not only slow, but also increases inconsistencies and human error.

Against this backdrop, the emergence of the concept of 'virtual banks' has introduced banks to a way of facilitating cost-efficiency in their daily operations. This concept was welcomed by the Hong Kong Monetary Authority (HKMA) in September 2017 when it launched a number of initiatives under the tagline *A New Era of Smart Banking*³³. Both financial firms (including existing banks in Hong Kong) and non-financial firms (including technology companies) may apply to own and operate a virtual bank in Hong Kong. The HKMA received over thirty applications, and eight licenses were granted in the first two quarters of 2019.

Virtual banks are acting as a catalyst for new developments across the banking sector, both digital and physical. They are focussing on hyper-personalisation and other new business models that can be realised through the adoption of new technologies, such as big data, and by leveraging an ecosystem of partners. With the increased adoption of new technologies, virtual banks are also able to enhance their own cost-efficiency. Without traditional physical branches and paper-based processes, virtual banks greatly rely on AI and data analytics to fully digitalise the customer experience while staying cost-effective.

The HKMA expects that virtual banks will enhance financial inclusion for retail customers and small and medium-sized enterprises (SMEs) in Hong Kong. For example, these banks should not impose any minimum account balance

requirement or low-balance fees on their customers according to the HKMA's *Guideline on Authorization of Virtual Banks*, they should also be able to offer more price-competitive products given their lower-cost income ratios. When it comes to welcoming SMEs as customers, traditional banks have been deterred by onboarding costs, but their automated procedures should make this segment more attractive for virtual banks.

In the face of increasing challenges from virtual banks, traditional banks are exploring automation opportunities to stay cost-effective. However, not every scenario can be improved by automation alone. For example, KYC and AML checks are time-consuming, as they require humans to read through all the documents. This situation cannot be improved by using automated solutions alone, as a computer cannot 'read' a document. For some more complicated processes, lack of IT integration leads to back-office staff entering data manually from one system to another to drive the workflow. To streamline business processes, alleviate operational costs and reduce the chance of errors, AI has become an important option for banks.

With the help of AI technologies such as natural language processing (NLP) and speech recognition, automated solutions are now able to handle cognitive processes such as reading through a document (e.g. with Intelligent Workforce) or listening to a recording (e.g. using voice and speech analysis). This is helping banks reduce costs. In addition to performing simple cognitive tasks, AI is also capable of making more accurate and precise predictions, in conjunction with the increasing amount of data sharing with business partners being enabled by open API. The HKMA introduced its Open API Framework for the Hong Kong Banking Sector in July 2018³⁴. This aims to facilitate the development and wider adoption of APIs within the sector. By the end of its Phase I implementation, twenty participating retail banks had made more than 500 Open APIs available, offering access to information on a wide range of banking products and services³⁵. With Open API, all information can be digitalised and understood by computers, and the systems involved can be more easily integrated to achieve end-to-end automation (including data collection, data integration and AI outputs).


³³ HKMA (September 2017). *A New Era of Smart Banking*. <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2017/09/20170929-3>

3.3. The need to better manage risk

Risk management is one of the major concerns of the banking sector: a careless mistake or ignorance of basic risk management can lead to disaster. However, integrated risk management can be a very complex and tedious process, as it requires high level accuracy and confidentiality. AI can simplify the whole risk management process by taking data such as transaction history, market trends and customer credit history into consideration, in the process developing more comprehensive insights into likely risk scenarios.

According to PwC's research on open banking³⁶, the use of alternative data – such as a borrower's employment history – enables banks to carry out a more comprehensive

credit assessment on each borrower. The availability of alternative data through partnerships with other types of institutions (non-banks), facilitated by API integration, is enabling non-traditional risk modelling, in which alternative data can be used for credit scoring and real-time dynamic pricing. Put together with AI models and big data analytics, this is enabling other aspects of customers and customer behaviour to be considered when making lending decisions. The loan market can thus be expanded while maintaining control over credit risk, as the bank's ability to consider the right variables means that more borrowers are able to obtain loans. The graph³⁶ below illustrates the difference when alternative data is considered as part of the credit risk assessment process:



AI can also help with other risks, such as money laundering and fraud events. For instance, when banks are going through the AML process, data points such as personal travel or historical transaction statistics can be used to help predict whether customers have any connection with potential overseas / multi-country money laundering activities. AI can also make more accurate predictions when identifying potential operational or legal risks.

Other types of risk management – such as market, regulatory or liquidity risk management – could also be improved with more data. AI enables the banking sector to focus on analytics and prevent losses in a proactive and real-time manner, rather than tackling them afterwards. Better risk management is clearly a strong motivator for banks to adopt AI.

³⁴ HKMA (17 July 2018). Open API Framework for the Banking Sector and the Launch of Open API on HKMA's Website. <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2018/07/20180718-5/>

³⁵ HKMA (31 July 2019). Open API Framework for the Banking Sector: One year on. <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2019/07/20190731-3/>

³⁶ PwC (June 2018). Opening the bank for a new era of growth. <https://www.pwc.com/il/he/bankim/assets/2018/Opening%20the%20bank%20for%20a%20new%20era%20of%20growth.pdf>

3.4. Barriers to AI adoption

While AI has been widely adopted in retail banking, total adoption across the whole banking sector still has room to grow. According to our survey, just 48% of the entire Hong Kong banking sector has adopted or plan to adopt AI applications. Local banks, in particular, are still hesitant about applying AI technologies.

Some of the major barriers to AI application are as follows:

- Explainability
 - Building trust and safety mechanisms and complying with regulations through transparent, explainable algorithms

- Resources
 - Identifying talent, collaborating and outsourcing to implement AI solutions
 - Developing data infrastructure to process AI solutions and store data
- Regulations
 - Establishing governance for ethical AI, data privacy and other security issues
 - The changing landscape of regulations required in a complex digital world

Customer-related barriers	Customers are concerned about data privacy or other security issues 2.602	Our business processes require face-to-face interaction with customers 2.072	Customers do not trust AI 1.659			
Environment-related barriers	Hard to find employees with AI expertise 2.467	How to design ethics of AI is still being debated 2.211	Costs in resolving customer doubt in AI solutions 2.120	Low government-provided incentives 1.754		
Organization-related barriers	Making needed operational changes is difficult 2.108	Changing corporate culture is difficult 1.868	Likelihood of success is low 1.737	Insufficient support from top management 1.425		
Policy-related barriers	Fear of legal consequences of adopting AI 2.240	Compliance challenges the replacement of human by AI 2.240	Obsolete laws prevent financial innovation based on AI 2.018			
AI readiness-related barriers	There is insufficient data 2.235	Have difficulties in analytical techniques 2.114	There is no business context to apply AI 1.766	The computing power is not sufficient 1.737		
AI application result-related barriers	Results of AI applications are difficult to explain 2.042	Results of AI applications are not robust 1.892	AI cannot ensure data security 1.721	Results of AI applications are biased 1.705	Directory Explainability Resources Regulations	
AI governance-related barriers	Our institution has not established a formal prioritization process for AI investments and projects 1.788	Our institution has not established formal processes to govern and manage AI projects 1.762	Our institution does not have a steering committee at executive or senior management level responsible for determining AI development 1.606	The CIO or a similar role cannot clearly articulate a vision for AI's role in our institution 1.482	CIO does not have a direct reporting line to the CEO 1.018	

In the following sections, these barriers will be discussed and their implications examined. It is important for banks to understand the potential execution challenges so that they

can better estimate the time and effort needed to implement AI solutions.


3.4.1. Lack of explainability

AI represents significant automation opportunities. However, as it becomes widely adopted in decision-making processes that do not offer clear explanations, it also represents a ‘black box’ threat, meaning that it may follow biases it has learnt from previous examples of poor human judgement. The lack of trust in such decision-making, and the lack of guarantees for stakeholders who may be affected by these decisions, are major barriers to AI adoption. This is true not only in the banking industry but in virtually the entire business world, according to PwC’s 22nd Annual Global CEO Survey³⁷. It is expected that concerns about AI’s trustworthiness will decline as it becomes more widely accepted and more records of proof are generated.

According to the *High-level Principles on Artificial Intelligence* circular published by the HKMA in November 2019³⁸, banks should ensure an appropriate level of explainability of AI models to all relevant parties. They should build their capabilities in understanding, validating and explaining how the AI arrives at its recommendations or decisions as soon as possible.

According to the survey, 49% of financial institutions in Hong Kong have faced major challenges in validating the fairness and accuracy of their AI models (e.g. biased or discriminatory decisions), and 32% have faced challenges in developing and launching AI applications due to a lack of understanding of AI models (e.g. black-box risk, adversarial attacks).

84% of CEOs agree that AI-based decisions need to be explainable in order to be trusted



Many banks, including those which have already started their AI journey, expressed concerns on the part of their executive team and board about the difficulty of assessing the implementation risks.

Banks should implement adequate measures during the design phase to ensure a level of explainability which is

appropriate and commensurate with the materiality of their AI applications. After banks define the level of their explainability requirements, they need to consider the choice of AI algorithms, data processing, additional software components and application design at the early stages of the design phase. This could be challenging for banks due to a lack of accepted standards and existing practice.

³⁷ PwC (2019). 22nd Annual Global CEO Survey. <https://www.pwc.com/mu/pwc-22nd-annual-global-ceo-survey-mu.pdf>

³⁸ HKMA (1 November 2019). High-level Principles on Artificial Intelligence. <https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2019/20191101e1.pdf>

3.4.2. Lack of resources for AI adoption

Shortage of talent

As an emerging technology, AI suffers from a shortage of expert talent. It is difficult for banks to compete with technology companies and FinTech companies to recruit the best on the market. Recruiting talent who are knowledgeable about both AI and banking is essential for AI development, but vacancies remain unfilled. Some banks are willing to take on inexperienced individuals and train them, but this takes time and resources. This talent shortage was ranked as the top challenge to their AI journey by some of our interviewees.

Some banks have chosen to outsource or develop other collaborative models to address the talent gap. According to our survey, 85% of financial institutions adopting AI technologies in Hong Kong have considered appointing external IT and technology companies to provide manpower resources for their projects. As partnerships are the most common approach to building AI capabilities, there is an increased demand for partnering with external consultants, innovation hubs, incubators and academic institutions.

How does your institution develop your current AI capabilities?	Present	Ranking	Future (coming 5 years)	Ranking
Partnering with technology firms	69%	2	83%	1
Internal research and development	75%	1	80%	3
Talent acquisition	68%	3	81%	2
Working with consultants	66%	4	77%	4
Participation in innovation hubs and incubators	44%	5	64%	5
Collaborating with academic institutions	32%	6	58%	6
Outsourcing R&D to technology firms	13%	7	33%	7

Although partnering with technology firms can help banks manage the talent shortage, it brings with it other risks and concerns.

AI is an emerging technology without established standards, meaning that a diverse range of solutions is available. However, the proprietary algorithms developed by AI companies need continuous maintenance and calibration. This is a challenge for banks in terms of managing partnerships and intellectual property rights with external vendors. These AI companies are usually start-ups, so traditional vendor assessment procedures may not apply.

Shortage of quality data

The quality of the data used for training and validating AI models has a great impact on model performance. It is critical for AI models to use production data for training and validation before deployment to ensure the algorithm is applicable to real life situations, otherwise there is a risk of unexpected AI model behaviour. However, banks face challenges in collecting machine-readable (internal and external) production data for AI model training in Hong Kong, for three main reasons:

- Insufficient amounts of open public data
- Concerns about the collection, sharing and use of customer data and data privacy
- The need for significant cleansing efforts for unstructured or poor quality data

As reflected in some of the interviews, there is no well-established database for public (or government) data in Hong Kong, bringing challenges for the development of AI models for some specific use cases. For instance, developing deep learning enhanced OCR recognition for identity cards needs a vast amount of data in the form of copies of real identity cards for training and accuracy testing. Banks may be unwilling to use copies of their clients' documents due to data privacy concerns. However, acquiring an official training dataset from the government can take more than a month, as both parties need to go through an internal compliance process. Given the iterative nature of machine learning and deep learning, in this case the model training period could be unexpectedly long for banks – possibly undoing the comparative advantage brought by AI solutions.

The shortage of cross-domain data is even more severe than that of public data. The lack of a common data exchange and sharing utility in Hong Kong makes it challenging to obtain non-bank datasets, such as customer lifestyle information. Even after customer consent has been obtained, a lack of comprehensive regulations defining how much personal information can be shared makes banks hesitate about using such data to avoid potential compliance or regulatory risks.

Flaws in internal data further aggravate the situation. While banks capture huge volumes of information every day, much of it is not in machine readable format that can be fed directly to the AI model. Converting information into AI-friendly formats requires significant cost and time in manual labour for data tagging and cleansing.

Moreover, banks are required to meet regulatory requirements on the handling of production data. Banks that are relatively mature in terms of AI already have policies in place. Many of the banks we interviewed said that they have leveraged the risk modelling guideline and risk management policies for AI model governance and control. However, banks which are new to AI need to accelerate the development of their data policies in order to mitigate model and compliance risks.

3.4.3. Legal and compliance challenges

Banks worry about regulations that may create future costs. Other hidden costs generated by the inherent uncertainties of new technology also make AI a difficult choice.

Data privacy and protection requirements

Like other data users, banks need to comply with the Personal Data (Privacy) Ordinance (PDPO). Banks are reminded to do so in the *Consumer Protection in respect of Use of Big Data Analytics and Artificial Intelligence (BDAI) by Authorised Institutions* circular released by the HKMA in November 2019³⁹. Data privacy legislation overseas, such as the EU's General Data Protection Regulation (GDPR) (effective from May 2018), the California Consumer Privacy Act (effective from January 2020) and others may also impact AI development. For example, Article 29 of the GDPR⁴⁰ requires transparency in profiling customers, where profiling refers to the evaluation of personal aspects using personal data, such as prediction of customer preference or

classification of customer segments. Potentially, customer segmentation and risk profiling by banks will need to be reviewed to ensure that proper safeguards are in place for the model to make objective recommendations. Activities including 'data collection', 'automated analysis to identify correlations' and 'applying the correlation to an individual to identify characteristics of present or future behaviour' are all subject to the regulation if their purpose is for automated decision making.

Banks may therefore need to consider the following potential risks when collecting and making use of customer data:

- Does the user agreement provide adequate consent to banks in relation to privacy rights for AI models?
- Does customer consent cover biometric customer data such as voice, facial patterns, and fingerprints, and other personal data used for AI development?
- Can banks be sure that their current data collection, storage, and analytic processes comply with privacy regulations and are sufficient to mitigate the risk of cyber-attack?

Uncertainty in ethical AI regulations

In addition to privacy, there are other concerns relating to ensuring that safeguards are in place to protect individuals' freedoms and rights in terms of choices.

In May 2019, France and Canada⁴⁰ announced a plan to co-create an International Panel on Artificial Intelligence (IPAI) to "support and guide the responsible development of artificial intelligence that is grounded in human rights, inclusion, diversity, innovation, and economic growth".

A number of bodies are drafting AI ethical guidelines which cover these areas. For instance, the Institute of Electrical and Electronics Engineers (IEEE) has published the first edition of an 'Ethically Aligned Design' paper⁴¹ that discusses ethical principles in implementing AI. The paper acknowledges the huge potential impact of AI on society, and the importance of addressing the non-technical aspects of AI, such as human rights, safety and privacy. It sets out eight general principles for achieving ethically aligned design, and contains discussions on topics such as the legal framework for AI accountability.

³⁹ HKMA (5 November 2019). Consumer Protection in respect of Use of Big Data Analytics and Artificial Intelligence (BDAI) by Authorized Institutions. <https://www.hkma.gov.hk/media/eng/doc/key-information/guidelines-and-circular/2019/20191105e1.pdf>

⁴⁰ Science and Economic Development Canada. (16 May 2019). Canada and France work with international community to support responsible use of artificial intelligence. <https://www.canada.ca/en/innovation-science-economic-development/news/2019/05/canada-and-france-work-with-international-community-to-support-responsible-use-of-artificial-intelligence.html>

⁴¹ Institute of Electrical and Electronics Engineers. (2016). Ethically Aligned Design. https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead_v1.pdf


The European Commission has commissioned a high-level expert group on AI to draft 'Ethics Guidelines for Trustworthy AI'⁴², which will serve as a starting point for discussions around European regulation and legislation and will set out a framework for achieving trustworthy AI. The guidelines note that AI should be lawful, ethical and robust. However, they focus on the ethical and robust issues, and offer no suggestions regarding lawfulness. The publication states that the guidelines are not intended to replace current or future policies and regulations. This to a certain extent indicates the relevance of having legal regulations in AI development, but also indicates the complexity involved in developing clear regulatory guidelines.

Uncertainty over technology assessment


A comprehensive technology assessment framework is recommended for assessing different AI applications or any complementary IT solutions to ensure their security and reliability. However, it is often a challenge for banks to assess

the technology risk associated with adopting an AI solution, owing to the complexity of the solution and inadequate in-house AI expertise. Also, the emerging nature of these technologies means that the industry has not had enough time to build up a well-recognised standard or practice for such a framework, making for increased compliance risks when adopting AI solutions. Technology risks aside (e.g. cybersecurity), regulatory compliance is a major concern for banks when adopting AI.

Nowadays, more and more popular AI libraries are open source and freely available for developers to leverage and modify when building their own AI solutions. However, should banks feel that managing the technology risks associated with open source solutions is a challenge, these resources may become an obstacle instead of a driver. It is very important for guidelines and industry-recognised certification processes to be established to assess and govern the use of new technologies.



⁴² High-Level Expert Group on AI. (8 April 2019). Ethics guidelines for trustworthy AI. <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>



4

Examples of AI in banking today

Key Takeaways

In the banking sector, AI has been applied in both customer-facing activities and back-office operations. The three most popular drivers for adoption of AI are to improve customer experience, automation and quality control, and risk management.

Improving customer experience is one of the major drivers for the adoption of AI. Banks have used AI to analyse customer data so as to offer personalised wealth management services, as well as for reading customer information and confirming customer identify so as to offer remote customer onboarding.

In automation and quality control, the most common AI use cases are cheque processing using optical character recognition; helping contact centres through speech-to-text and natural language processing; and assisting customers using Chatbots and robotic process automation.

Another common application of AI application is for risk management. AI can help banks to better assess risk and detect fraud by combining both supervised and unsupervised learning models to build detailed individual customer risk profiles. In bank operations, AI can help to monitor the bank's IT infrastructure, detecting and preventing potential attacks. Lastly, combining AI with alternative data can be used to assess credit risk for SMEs without the need for traditional credit data.



In the previous section we explored major drivers of investment in AI technologies, which included the need to improve quality and customer experience, reduce operational costs and strengthen risk management. Some barriers to adopting AI technologies were also explored.

The previously mentioned survey (see Appendix B) also examined how the banking sector is coping with the technology. In the survey, a total of 27 AI use cases were identified. These are either being adopted, being planned for, or are being explored.

The top five AI use cases being adopted by the banking sector relate to cybersecurity, client-facing Chatbots, remote

client on-boarding, biometric customer identification, and personalised advertising. As for planned use cases, the most attractive for banks include anti-money laundering, fraud detection and financial advisory. Overall, the current trend for AI adoption is towards the development of risk management and customer-facing applications.

To further illustrate the status of AI use cases in the sector, the 27 use cases have been mapped to the 25 banking functions shown in the table below. We have also quantified the average investment in each of these banking functions to create a comprehensive picture of current AI investment in the Hong Kong banking sector.

Significant investment made and mature AI capability in place		Core Banking	Customer	Channel	Corporate Functions	Propositions
		Account Management 67.00			Fraud Detection and Prevention 62.00	
			Customer Service 48.00	Mobile 56.00	Technology Management 54.50	
			Customer Acquisition 46.00	Internet 56.00	Clearance & Settlement 33.33	
			Segmentation & Insights 32.00		Compliance Documentation & Reporting 33.00	
		Product Innovation & Pricing 23.00		Call Centre 29.00	Talent Management 31.00	
		Collateral Management 16.00			Risk Assessment & Monitoring 27.00	Corporate Lending 26.00
						Consumer Financing 26.00
						Investment & Securities 20.33
						Wealth Management 19.50
					Finance & Reporting 13.00	
				Branch 7.00		
						Cash Management 7.00
						Cards 7.00
					Payment N/A	Investment Banking N/A

No AI applications and no AI capacity

This section summarises and discusses the three most popular areas for AI adoption: customer experience, automation and quality control, and risk management. For each of these areas, several common use cases are discussed.

Customer experience is an area that many banks in Hong Kong are looking to improve through AI. The most commonly adopted use cases are robo-advisors and remote client on-boarding. The adoption of AI has helped banks improve different banking functions, including customer acquisition and wealth management.

AI can also help banks in automation and quality control. Popular use cases are cheque processing, voice and speech analysis, Chatbots and intelligent process automation. With the help of AI in automation and quality control, banks can improve their core banking and corporate functions, as well as offering a better customer experience.

In risk management, AI can help banks better assess risk and detect fraud. Use cases identified are machine learning-aided transactional risk assessment and fraud detection, predictive maintenance in data centres, and credit risk assessment for SMEs with alternative data.

Given that these use cases entail more mature technologies, they should be applicable as both a starting point for banks in Hong Kong that are just beginning to adopt AI, and as a source of ideas for more experienced banks.

4.1. Customer experience

4.1.1. Financial advice with robo-advisors


Robo-advisor technologies have emerged that offer professional investment recommendations to the mass public at a reasonable cost. Traditionally, the general public have rarely had access to the sort of financial advice provided by private bankers, nor have they had the experience to conduct financial analysis on their own investments.

Robo-advisors act as low-cost alternatives by eliminating human effort and making investment recommendations available to anyone, anytime and anywhere. They utilise quantitative and machine learning models which are trained by inputting a huge amount of historical market data, with

the aim of building an optimised and diversified portfolio that will maximise returns according to the risk appetites of the customer⁴³. They are highly accessible as they are online 24/7, regardless of one's physical location.

Wealth management is one of the main services provided by banks in Hong Kong. However, adoption of AI technologies in this field is still at a relatively early stage. According to our survey, out of 234 AI applications that have been launched or are going to be launched by banks in Hong Kong, just twelve are for investment and wealth management. Within these applications, only two are related to providing financial advice, although there is growing interest in applying AI for the automation of financial advice.

Robo-advisors automate the portfolio management process as follows⁴⁴:



⁴³ Tao, L. (20 July 2018). What's your risk appetite? Your robo-adviser has the answer. <https://www.scmp.com/tech/start-ups/article/2102405/whats-your-risk-appetite-your-robo-advisor-has-answer>


⁴⁴ FrankenField J (1 October 2019). What is a Robo advisor? <https://www.investopedia.com/terms/r/roboadvisor-roboadviser.asp>

A Hong Kong-based robo-advisor⁴⁵ can use machine learning algorithms to create a portfolio that contains multiple exchange-traded funds (ETFs) catering to Hong Kong-based investors, based on their risk appetite. Trained by supervised learning using historical market data, a robo-advisor selects stocks and other assets based on criteria that include customer risk appetite, industry preferences and up to 200 other factors. Among these are transaction volume and price, fundamentals, macroeconomic factors derived from market data, financial statements and analyst reports. The portfolio performance is monitored and attribution analysis is performed to identify the key factors contributing to profit and loss. This is updated to the machine learning model for further stock selection⁴⁶.

Despite the potential benefits of robo-advisors, the quality of their performance is yet to be proved. Comparison of the portfolio performances of different robo-advisors shows that they do not always yield positive returns⁴⁷.

New regulatory frameworks are being developed to impose stricter regulations on robo-advisors. In July 2019, the Hong Kong Securities and Futures Commission (SFC)⁴⁸ issued *Guidelines on Online Distribution and Advisory Platforms*, which includes a chapter on robo-advice. It covers topics such as availability of information to customers, client profiling, system design and development, supervision and testing of algorithms, staff resources and automatic rebalancing.

The Financial Services Regulatory Authority of Abu Dhabi Global Market (ADGM) has also issued a guideline on robo-advisory activities. This⁴⁹ explicitly includes an algorithm governance section which recommends that the robo-advisory model must produce explainable, traceable, and repeatable outcomes. However, some robo-advisors have adopted sophisticated deep learning models which are difficult to explain, while others have adopted explainable models for which the investment details have not been disclosed for proprietary reasons.



⁴⁵ AQUMON. <https://www.aqumon.com/>

⁴⁶ AQUMON (29 April 2019) A股投资装甲—AQUMON智能投资引擎上线！ https://www.aqumon.com/webview?content=news&id=161&lang=zh_cn

⁴⁷ Fund Selector Asia. (9 September 2019). Robo-advisor performance: August 2019. <https://fundselectorasia.com/robo-advisor-performance-august-2019/.Activities.pdf>

⁴⁸ Securities and Futures Commission. (July 2019). Guidelines on Online Distribution and Advisory Platforms. <https://www.sfc.hk/web/EN/assets/components/codes/files-current/web/guidelines/guidelines-on-online-distribution-and-advisory-platforms/guidelines-on-online-distribution-and-advisory-platforms.pdf>

⁴⁹ Abu Dhabi Global Market. Supplementary Guidance – Authorisation of Digital Investment Management - (“Robo-advisory”) Activities. http://adgm.complinet.com/net_file_store/new_rulebooks/s/u/Supplementary_Guidance_Authorisation_of_Digital_Investment_Management_Robo-advisory_Activities.pdf

⁵⁰ CITIC Bank. Connecting etnet apps to in Motion to perform FX transactions. <https://www.cncbinternational.com/personal/e-banking/inmotion/en/index.html>


4.1.2. Remote client on-boarding

The process of opening a bank account often defines a customer's future relationship with a bank. As simple as it may seem, it often involves paper-intensive manual processes. A frictionless process, conducted either through a human advisor or digitally, can leave the customer with a highly positive first impression of the bank's service quality.

Remote client on-boarding uses machine learning and AI for fraud detection, displacing traditional knowledge-based authentication (KBA) methods for customer due diligence (CDD). It has the potential to save significant processing time and costs by digitalising end-to-end processes. For example, one Hong Kong bank claims that the time needed to open a bank account remotely using a mobile app can be shortened to just fifteen minutes⁵⁰. Customers can thus spend more time on the actual banking services they require, while banks can channel human resources towards performing complex banking transactions and higher value-added activities.

According to our survey, customer acquisition is an area that Hong Kong banks want to improve with the help of AI. More than 27 AI applications have been launched or are planned in this area, while another 22 are being explored or studied. Over 63% of these applications are related to remote client on-boarding.

Unlike face-to-face account opening, remote client on-boarding involves a stringent Customer Due Diligence identification process. This new account opening process is shown below, during which various computer vision technologies (such as machine learning enhanced Optical Character Recognition (OCR), facial recognition and voice authentication) are applied in the scanning and checking process:



Extracting ID information with machine learning enhanced Optical Character Recognition (OCR)

When a user uploads pictures or takes a picture of his or her ID document with a mobile device, other information on the ID document is also captured for further recognition. Machine learning algorithms then help identify the expected position of words based on the type of document recognised. Optical character recognition (OCR) is applied to extract words from the ID document and validate these against the user's input. Machine learning helps enhance OCR technology by handling the unstructured, incidental images in the user's uploaded documents with greater accuracy. It typically makes use of neural network models that can deal with complex backgrounds, noise, lighting, fonts and other distractions.

Verifying customers' identity using facial recognition and biometric liveness detection


Facial recognition technology can match the customer's face with that of the photo extracted from their ID document. Multiple digital images are sometimes required at different angles to verify unique security features on the ID card, such as microprint, hologram, layout and font⁵¹. The user may also be asked to record a 2-3 second selfie video, and a face extraction model is used to verify whether the video was shot by the user in a live environment.

Verifying customers' identity using voice authentication technology

The customer's identity can also be verified by what is called speaker recognition, or voice authentication. This makes use of the unique acoustic features of the voice to identify a user. AI-assisted voice recognition authenticates individuals more precisely by training on a large amount of voice samples. Under the AI approach, voice authentication performs better in noisier situations, thereby guaranteeing more functional usability and robust security.

All this information and documentation can then be sent to the bank's back office for further AML or KYC checking, as appropriate. This may require checking and vetting with external regulatory agencies and law enforcement agencies. AI can help to facilitate AML and KYC checking processes, with NLP extracting information from scanned documents and intelligent process automation (IPA) cross-checking information from scanned documents with different data sources in real-time.

The core risk of remote client on-boarding is in validating the authenticity of biometric information in order to rule out synthesised or faked biometric identity. The availability of open-source human image and video synthesis programmes online has lowered the entry barriers for fraud, and is giving rise to higher risks of identity impersonation. According to the Voice Intelligence Report⁵² published by call centre software solution provider Pindrop, voice fraud surged more than four times from 2013 to 2017 and is continuing to increase.



⁵¹ Jumio. ID Verification Solutions. <https://www.jumio.com/trusted-identity/netverify/id-verification/>

⁵² Pindrop 2018 Voice Intelligence Report. (2018). <https://www.pindrop.com/2018-voice-intelligence-report/#>

Legislation such as the 'Defending Each and Every Person from False Appearances by Keeping Exploitation Subject to Accountability Act of 2019', also known as the Deepfake Accountability Act, has been proposed to tackle the problem of Deepfake videos in the US. Facebook is joining forces with Microsoft researchers and academics to run the Deepfake Detection Challenge (DFDC), with the intention of creating software to detect Deepfake videos⁵³.

Despite significant efforts by various parties to defend synthetic biometric identity, this could still be a key challenge for banks, as counterfeiting techniques are constantly evolving in the digital world. To mitigate risk, the banking sector need to closely monitor the development of AI technologies and establish a comprehensive technology risk assessment framework.

4.2. Automation and quality control


4.2.1. Cheque processing

Bank cheques have been and will continue to be widely used for financial transactions around the world. Enormous volumes of written and computer-printed bank cheques are processed and verified manually every day. They often need detailed information checking to prevent incorrect processing. Considerable operational resources are therefore devoted to visually verifying every field in each cheque, such as date, signature, and courtesy amounts.

To reduce costs, a digital cheque clearing system (i.e. a cheque truncation system that involves the digitalisation of physical cheques into electronic images for transmission between banks) has emerged as a global trend⁵⁴. With the help of this system, cheque images and magnetic ink character recognition (MICR)⁵⁵ data (including bank codes, bank account and cheque numbers) are captured at the collecting branch and transmitted electronically. This can be seen as a first generation application of optical character recognition (OCR) for cheques, where the system can only capture limited information and recognise pre-set computer-readable fonts⁵⁶.

Unfortunately, significant manual efforts are still required on top of this OCR-assisted process. To further reduce cost, time and effort, banks have started exploring a solution capable of automating the entire process of cheque recognition, verification and data entry, even for handwritten cheques. This is a cheque recognition solution based on OCR and supplemented by deep learning models.

A comprehensive cheque processing solution typically starts with input capturing and pre-processing from mobile. The bank then extracts the MICR signature as well as other items, and passes them on for identification and recognition of the information needed, such as bank information and legal amount. All the recognised contents are then sent to the database for rule-based verification. The last step is to export the verified results for cheque clearing. This high-level process map⁵⁷ is shown below:



⁵³ Schroepfer, M. (5 September 2019). Creating a data set and a challenge for deepfakes. <https://ai.facebook.com/blog/deepfake-detection-challenge/>

⁵⁴ Jayadevan, R., Kolhe, S.R., Patil, P.M., & Pal, U. (2011). Automatic processing of handwritten bank cheque images: a survey. International Journal on Document Analysis and Recognition (IJDAR), 15, 267-296.

⁵⁵ Kagan, J. (6 May 2019). Magnetic Ink Character Recognition (MICR) Line. <https://www.investopedia.com/terms/m/magnetic-ink-character-recognition-line-micr.asp>

⁵⁶ Devscope. MICR NUMBERS. <http://www.devscope.net/Content/ocrchecks.aspx>

⁵⁷ MPS (23 February 2018). Cheque Processing. <https://multipasol.com/cheque-processing/>

AI is used to analyse the cheque layout during image capturing and pre-processing to identify relevant information to be extracted. OCR technology is then used to extract text from identified fields in the electronic images. Pre-set fonts using MICR are recognised simultaneously, identifying the payee account and bank. Next, a deep learning neural network model is employed for handwritten character recognition, a model which is able to continuously learn different styles of handwriting during processing to constantly improve its recognition accuracy.

Unlike in the traditional cheque recognition system, by incorporating deep learning into the solution banks are able to minimise human involvement because the entire cheque, including both computer-printed and handwritten fields, can be recognised and transacted automatically. With appropriate infrastructure and computing power, this automated cheque processing solution is perfect for large scale processing, as it takes less than one second to read an entire cheque⁵⁸. Furthermore, the self-learning nature of deep learning models means that recognition accuracy improves over time as more data is captured. According to one of the interviewed banks, a drastic improvement in recognition accuracy from around 65% to 95% was recorded after it applied a deep learning model to enhance OCR cheque recognition.

The multilingual environment in Hong Kong further increases the variety of characters to be processed, as Chinese OCR is significantly more difficult than English. The large number of Chinese characters and high textual information density make the successful recognition of characters challenging. The deep learning model thus requires a huge amount of training images, and significant human effort to prepare data for supervised learning. This has incentivised banks to develop deep learning models or seek local experts for training and data preparation to ensure the recognition procedure is accurate and applicable to the local market.

4.2.2. Voice and speech analysis in contact centres

Call centres have acted as a major service and sales channel for banks since their inception in the 1980s⁵⁹, and still remain a significant channel by which customers interact with bank representatives. In the digital era, many call centres have been transformed into 'contact centres', with greater capabilities for handling customer enquiries through e-mails, instant messaging applications and social media. These interactions with customers generate massive amounts of new data every day, both as a result

of transactions (e.g. waiving of credit card annual fees) and activities (e.g. the conversation between customer and agent when negotiating a fee waiver). The data points containing valuable customer insights (i.e. data stored in a structured format, such as databases, or in an unstructured format, such as a recorded customer call) are important for the bank to identify potential improvements in its operations, and to identify new business opportunities.

Voice and speech analysis is used in call centres to help improve their customer service, segmentation and insights functions. Over the years, banks have mined the structured data they have captured, but have found it costly to utilise unstructured data. For example, manual efforts are required to listen to and summarise samples of customer call records for compliance, investigation and training purposes. It is unfeasible to tag and analyse every call, given the significant cost of hiring people to listen and summarise.

In this field, two AI applications have been launched in Hong Kong and four are at the planning stage. But there is growing interest among Hong Kong's banking sector, with a further eleven applications being studied.

Common scenarios that involve applying AI-enabled conversation analysis at a contact centre include:

- Voice recognition technology to identify and verify a customer's identity based on his or her voice signature.
- Speech-to-text technologies that convert customer dialogue into text in real time.
- Natural Language Processing to analyse and summarise customer dialogue without the need for human involvement.
- Sentiment analysis, emotion detection and conversation pattern analytics, which deduce customers' intentions and preferences.
- Real-time recommendations for contact centre agents on advice to give to the customer.
- Sales compliance analysis. This can also be conducted in a much more effective manner with the support of AI technologies. Reports on problematic calls can also be generated by tracking various customised metrics.

To enhance the productivity of contact centre agents, AI can be deployed to analyse their performance, supplementing traditional contact centre key performance indicators (KPIs) such as response time and number of complaints handled.

⁵⁸ Keary, T. (20 February 2019). AI and OCR: How optical character recognition is being revitalised. <https://www.information-age.com/optical-character-recognition-tools-ocr-ai-123479324/>

⁵⁹ Call Centre Helper (13 May 2019). The History of the Call Centre – Updated. <https://www.callcentrehelper.com/the-history-of-the-call-centre-15085.htm>

Specific agent performance metrics, including silence ratios, speaking speed and interruption rate, can also be analysed to generate personalised and real-time feedback for each agent. For instance, one bank has applied an AI-based solution to analyse the correlation between an agent's seniority, work time and productivity⁶⁰.

The data generated from contact centre activities and other operational functions of a bank can generate insights into ways in which the bank could enhance the cost-effectiveness and efficiency of the contact centre. For example, one bank recorded a 15% increase in sales quality scores and a 10% increase in net promoter scores as a result of its ability to accurately analyse every call. This is also useful for identifying sales opportunities.

One key challenge to developing voice and speech analysis in Hong Kong is the difficulty of processing Cantonese, and the use of mixed languages (e.g. English and Cantonese) by many locals. Some technology breakthroughs have been made by local FinTech start-ups that are improving speech-to-text and NLP accuracy. According to one of the technology start-ups we interviewed, their technology recognition accuracy now stands at around 80% compared to human accuracy.

4.2.3. Chatbots


Chatbots are online, virtual conversation agents that allow automated responses based on customers' input, and simulate human interaction without the presence of a contact centre agent. They have become an increasingly popular service tool in places where banks find it hard to justify investment in property and labour for physical contact centres. Ideally, Chatbots provide faster and more consistent responses than a human agent. Human effort is also eliminated, since Chatbots handle a lot of repetitive questions.

Chatbots are one of the top three most popular applications among the 27 listed in our survey. The banking sector has shown widespread interest in this technology, especially due to its ability to help automate account management in core banking, and customer services in the mobile and Internet channels. Seventeen banks have already launched or plan to launch their own Chatbot, while 25 are exploring this option.

One of our interviewee banks stated that its Chatbots handle more than four million customer enquiries from Mainland China on a daily basis. Most Hong Kong respondents are exploring more extensive use of Chatbots, in light of their benefits in terms of cost savings and their ability to capture customer insights that are valuable for marketing and product design.

By using NLP, a Chatbot can understand what people are saying in different languages. By observing human actions, the robot also accumulates common in-house processes and systems knowledge. With machine learning and rule-based logic from an in-house IT knowledge base, the Chatbot can learn to diagnose a problem, decide whether it is something it can fix and, if not, pass it on to a human IT support team.

A Chatbot can handle speech input if a speech-to-text synthesiser is used. It can then process the input to identify keywords and key terms, and predict the user's financial intent using natural language processing (NLP). It can automatically generate responses by leveraging responses that are rule-based or based on the user's predicted request, finding corresponding structured data to form a machine-readable response. Finally, it applies natural language generation (NLG) to respond.



⁶⁰ Sestek. ING Bank Call centre Speech Analytics Case Study. <https://www.sestek.com/case-studies/ing-bank-speech-analytics-case-study/>

If needed, the Chatbot can then integrate with the bank's internal workflow engine to trigger different types of transaction requests. More advanced implementation of AI can reinforce machine learning by continuously feeding new data to the model, improving its performance.

Banks have also been piloting the deployment of cognitive robots for internal customer service helpdesks. Another virtual robot can handle simple problems that amount to nearly half of all support requests (e.g. email issues, password resets and Wi-Fi or network issues). This Chatbot is still in the early stages of development in that the services it can provide are limited and significant support from humans is still very much needed. But by continuously learning from, developing, and interacting with customers, it becomes able to solve more complicated problems and provide better responses. Following training, the robot can already understand 87% of queries. Staff from IT support are certain that virtual robots and humans can work together, enabling humans to focus on providing more complex support services⁶¹.

In an ideal situation, a Chatbot should be able to accurately understand a customer's intentions and provide prompt and consistent responses. Incorrect predictions of intent can result in confusion and customer dissatisfaction. In reality, however, the Chatbots deployed by many of the banks we surveyed are still at the development stage, only offering basic banking services such as information about balances and account details. Chatbots are still augmenting human actions rather than replacing them. Their underlying limitations, such as a lack of decision-making and research skills, require Chatbots to route complex questions to human staff.

4.2.4. Intelligent robotic process automation

Banks have been automating repetitive operational processes by integrating systems since computers became commercially available. Alongside the traditional system integration approach with heavy reliance on IT development, a recent trend has been for system integration to utilise general purpose robotic process automation (RPA) software. RPA is extremely efficient at handling repetitive, structured data automatically. However, with the amount of unstructured data growing rapidly, banks are looking to ramp up their RPA systems with more advanced versions of process automation software, such as intelligent process automation (IPA). IPA combines machine learning with RPA to deliver powerful tools that mimic human interaction and make decisions based on robotic outputs. It produces

automation capabilities that heighten business value and competitive advantages. With less need for human intervention, IPA can free staff from mundane activities so they can focus on processing higher risk or higher value cases.

Most banks are attempting to leverage technologies to automate their operations, particularly for clearance, settlement and compliance documentation and reporting. According to our survey, ten AI applications are already in place among banks in Hong Kong to perform these corporate functions, showing that the adoption of AI technologies in these areas is relatively mature compared to other functional areas. According to BIS Research, the finance and banking sector is expected to occupy around 40% of the global cognitive robotic process automation market by 2026⁶². Some banks have already adopted IPA to support or automate document processing, trade surveillance, compliance activities and customer services.

Here are some common applications of IPA:

Trade finance documentation

Trade finance transactions usually involve large amounts of documentation from different parties, such as the financial records of the borrower, sales orders or invoices, trade documentation such as bills of lading and customs declarations, along with other compliance and safety-related documentation. Trade finance relies on the authenticity of the documentation, and hence extensive vetting and cross-checking by the bank is required. However, owing to a lack of international standards for global trade documentation, rule-based automation cannot be applied.

IPA utilising deep learning-based computer vision and other machine learning techniques can be used to extract information from scanned trade documents. Other machine learning algorithms can be used to classify the document type by clustering documents that share similar document elements or features (e.g. procurement orders or invoices).

Based on the document type identified, cognitive agents can make use of AI-enhanced OCR (optical character recognition), HCR (handwriting character recognition) and NLP to extract trade information from the text. Document information captured can then be processed, validated, and reconciled (e.g. by the matching of procurement orders and invoices), and corresponding workflows can be triggered automatically using pre-defined rules.

⁶¹ Credit Suisse (2 February 2018). Amelia – Artificial Intelligence in Action at Credit Suisse. <https://www.credit-suisse.com/about-us-news/en/articles/news-and-expertise/amelia-artificial-intelligence-in-action-at-credit-suisse-201802.html>

⁶² BIS Research (2017). Global Cognitive Robotic Process Automation Market, Analysis & Forecast 2017-2026 Focus on Type (Services and Platform) and Industry (Finance and banking, Telecom and IT Services, Insurance and Healthcare). <https://bisresearch.com/industry-report/global-cognitive-robotic-process-automation-market-2026.html>

Some banks^{63,64} have already developed Trade AI Engines using NLP and machine learning-enhanced OCR to classify and process millions of pages of unstructured trade documents, digitising them into machine-readable format in order to generate insights into the networks of related parties and customer trade activity.

Cybersecurity


Banks need to safeguard their systems 24/7 to prevent system outages or data losses caused by cyber-attacks, so they need to perform regular cybersecurity incident simulations and drills to discover cyber threats and investigate security risks across the whole IT infrastructure. Such comprehensive simulations and investigations are labour intensive.

IPA can be deployed to mimic users' behaviour, including how users interact with systems daily and how they react to possible cyber-attacks, in order to discover security loopholes in an application, IT infrastructure or operational process. IPA can supplement the work of humans on large-scale incident simulations which traditionally involve hundreds of operational staff. The same simulation practice assisted by IPA could be applied in advance before rolling out new releases of applications to help banks identify and prevent a zero-day attack (i.e. an attack on a vulnerability on the day of its discovery).

Regulatory change management

Given the complexity of many new or newly updated regulations and the importance of implementing them correctly, compliance teams in banks have traditionally used a manual change management process. This includes reading and analysing the regulations, identifying the implications, and updating corresponding internal policies and procedures to implement the regulatory changes.

Some banks are exploring the use of RPA to monitor regulatory changes published on regulators' websites, and of IPA to analyse the regulatory documents and identify changes. IPA can identify which banking lines and functions will be impacted by regulatory changes, and trigger a corresponding workflow to handle the impact assessment and change management processes. The HKMA is studying the use of machine-readable regulations under its Smart Banking initiatives to support banks' regulatory compliance processes⁶⁵.



⁶³ Finextra (30 April 2019). Citi to develop AI-based trade finance compliance platform. <https://www.finextra.com/newsarticle/33745/citi-to-develop-ai-based-trade-finance-compliance-platform>

⁶⁴ Standard Chartered (22 May 2019). Industry-leading Augmented Intelligence (AI) engine to transform the traditional Trade Documentary System. <https://www.sc.com/en/media/press-release/weve-pioneered-trade-ai-engine-with-ibm/>

⁶⁵ HKMA (27 September 2018). Regtech in the Smart Banking Era – A Supervisor's Perspective. <https://www.hkma.gov.hk/eng/news-and-media/speeches/2018/09/20180927-2/>



4.3. Risk management

4.3.1. Machine learning-aided transactional risk assessment and fraud detection

Traditionally, banks use rule-based systems with predefined transaction patterns to identify potential money laundering activities. In cases of complex transactional behaviour and patterns, the volume of false positive alerts identified by this approach increases, requiring significant time to investigate alerts and filter false alarms. One recent advance is to apply machine learning to the detection of money laundering⁶⁶.

Machine learning models, trained by customer and merchant data from various internal and external data sources, are used to understand business activities and complex transactions. The trained models can be deployed to scan through all the suspicious transactions identified by the AML solution⁶⁶ to differentiate false positive alerts from transactions that really require investigation. In order to

continuously improve the model, new data patterns and the results of investigations by humans on identified suspicious transactions are looped back to the machine learning model for smarter future recommendations.

By combining both supervised and unsupervised learning models, detailed individual customer risk profiles can be built. Every transaction can then be scored based on the corresponding level of fraud or money laundering risk. Such risk prediction input has been widely adopted to enhance existing AML decision engine rules in use cases such as e-commerce transactions, mobile banking, loan applications and recurring banking payments. For instance, real-time AML assessment of customer fraud risk based on a customer's risk profile can be conducted during the customer on-boarding process. When higher risk is predicted or not enough information is available for decision-making, additional customer-specific follow-up questions can be triggered to protect the bank from the risk of fraud.

⁶⁶ Twomey, N. (19 December 2018). Corporate Compliance Insights. 5 Ways AI is Impacting AML and KYC Compliance. <https://www.corporatecomplianceinsights.com/5-ways-ai-is-impacting-aml-and-kyc-compliance/>



Banks have a long history of applying data analysis for fraud detection and prevention, making this use case popular among banks. In our survey, it ranked as the top area for which banks would like to apply AI. Currently, 36 applications have already been launched, 23 are being planned, and 94 are at the study stage. The survey also shows that most of these applications focus on Anti-Money Laundering (AML) related issues.

One area where AI technology is already being applied is in the real-time monitoring of fraudulent activities relating to credit card payments. Machine learning models for fraud detection are trained using large historical credit card transaction data sets. Using actual fraud cases that have been reported and verified, banks can tag historical data with relevant fraud and non-fraud labels to train classification algorithms. The trained model can then identify high risk suspicious transactions with pre-determined features indicating fraud. These range from transaction times and locations and other spending behaviour statistics to the historical transaction patterns of card holders. Geo-

location data is one pre-determined feature, as fraudulent transactions commonly occur far away from where card holders are located at the time of purchase.

Banks also apply clustering and other unsupervised learning methods to identify new relevant transaction patterns or features without human labelling⁶⁷. Clustering algorithms group common transactions into clusters with similar behaviours, and these are set as the baseline for normal transactions. Anomalies are then identified as outliers, lacking strong membership of any cluster group, for additional review. Human operations staff can then confirm if these transactions are legitimate, and at the same time the results can be fed back to the learning model to improve accuracy through reinforcement learning. One bank we interviewed claimed that its false positive rate for flagging laundering risk was reduced by 20% after a machine learning-based system was deployed to incorporate newly-found patterns into the original AML solutions.


⁶⁷ Westerlund, H.S. (25 October 2018). The big Challenges Faced in KYC Implementation Today. <https://blog.instantor.com/the-big-challenges-faced-in-kyc-implementation-today>

4.3.2. Predictive maintenance in data centres and operations risk

In the digital era, keeping data centres up and running is among the most critical of a bank's missions. A host of applications, from online banking to trade settlement and from accounting to risk management, all rely on data centres providing uninterrupted around-the-clock service. However, banks are finding it increasingly challenging not only to maintain, but also to expand their data centres, given the massive scale and complexity of banking systems. It is also challenging for their IT departments to keep up with the rapidly changing business environment. As customers

expect ever faster turnaround of products and services, bank IT departments have been busy scaling up their infrastructure, enhancing cybersecurity and ensuring that systems remain stable and robust.

AIOps (AI for IT Operations) greatly improves the availability and the consistency of performance of bank production systems. Hidden server abnormalities can be discovered by AI in advance through operation and maintenance log analysis, and the relevant IT team can then intervene and proactively respond to potential production or security issues. With AI-driven proactive alerts, the data centre operation model for faults can be transformed from passive firefighting to active fault prediction.



According to our survey, technology management is one of the top banking functions that banks are choosing to invest in. Eighteen applications are now live and 16 are planned. Within technology management, cybersecurity has the most applications, with 13 launched, seven planned and 14 under study. An intelligent identification system can help enhance and improve technology management accordingly.


AIOps has emerged among large corporations and financial institutions. It does not rely on human-specified rules, but uses machine learning algorithms to automatically process huge amounts of data generated by servers and applications. These include server logs, application error event logs, and sensor data within data centre premises. In a typical bank IT environment, with hundreds of application servers, there could be tens of thousands of server event trends or patterns, which are difficult for humans to identify.

Some banks have noted that adopting AIOps has helped to reduce the time needed to locate IT infrastructure issues and analyse problems and root causes, greatly improving IT service levels for emergencies and critical events⁶⁸. In addition to proactively predicting and resolving server issues, AIOps can help to identify performance bottlenecks and optimise server utilisation by routing workload intelligently so as to reduce downtime in operations, data delivery delays⁶⁹ and power consumption⁷⁰.

Model-based simulations are also performed by AI to optimise the data centre's server and network equipment arrangements. Using simulations, the extent to which equipment changes may disrupt operations can be predicted, thus assisting the IT operation team to plan ahead and determine an optimum maintenance schedule. Some data centres even use AI-enabled predictive analytics to perform disaster recovery (DR) drill simulations to supplement their actual regular DR drills, which are costly, time-consuming and require significant levels of participation from IT and operations staff⁷¹.

AIOps could also be used to detect and predict threats and other potentially malicious activities. It is very hard for conventional systems to keep up with all the new malware being created. However, banks can train AI with patterns of existing malware in order to detect future potential malware. AI can even be trained to detect the smallest signs of virus or malware attacks before they enter banks' systems, so that they can be isolated in advance.

In short, AIOps uses predictive analytics powered by AI machine learning to forecast and proactively prevent server issues, simulate different data centre scenarios, recommend changes and automatically optimise server workload and utilisation. All of these outcomes are helping increase the performance of IT infrastructure operations while reducing human effort.



⁶⁸ 郑仕辉 (18 February 2019). 交通银行的智能运维（AIOps）实践. <https://www.cebnet.com.cn/20190218/102551555.html>

⁶⁹ Instor (2019). How can Artificial Intelligence improve data centre efficiency? <https://instor.com/blog/artificial-intelligence-improve-efficiency/>

⁷⁰ Vi Kasinathan, D. (6 February 2019). Machine Learning and Predictive Analytics: Making Power Management Smarter. <http://www.virtualpowersystems.com/machine-learning-and-predictive-analytics-making-power-mana>

⁷¹ Taylor, A. (6 June 2018). Network world. Using AI and Predictive Analytics Data centre. <https://www.networkworld.com/article/3279274/using-ai-and-predictive-analytics-to-improve-the-data-centre.html>

4.3.3. Credit risk assessment for SMEs with alternative data


Established banks have adopted the credit history assessment principle when estimating a potential customer's default risk. Statistically, the longer the positive history of lending and repayment, the less risk arises for a particular credit transaction. These statistical models for credit risk assessment and the associated credit control framework have helped banks manage credit risk effectively in the past.

To address the challenge of banks not having enough in-house data relating to credit history to produce statistically robust predictions for particular market segments, such as SMEs, the Commercial Credit Reference Agency (CCRA) was set up in Hong Kong in 2004. Its role is to collect SME credit information (including credit limits, indebtedness and credit history) and share this information (with the SMEs' consent) with members of the Hong Kong Association of Banks (HKAB) and the Hong Kong Association of Restricted Licence Banks and Deposit-taking Companies

(The DTC Association). As of 2017, the CCRA had collected over 176,000 credit records from 119,000 SMEs, sole proprietorships or partnerships with turnover of up to HKD50 million. The database was further expanded from December 2017 to cover non-listed limited companies with an annual turnover of up to HKD100 million⁷².

While the CCRA provides a trusted source for building credit score models, a large proportion of SMEs still remain uncovered, given that there are 340,000 SMEs operating in Hong Kong⁷³. These are uncovered either because of their brief operating history or because the owners have supplied incomplete or overly complex financial information. This shortfall of data has therefore excluded them from accessing loans under the current credit risk assessment framework.

However, there are statistical ways to identify the correlation between a customer's non-financial transactions and behaviour and their credit default risk. This can include factors such as demographics, expenses and payment transactions and lifestyle data. This is termed 'alternative data'.



⁷² Hong Kong Monetary Authority (24 August 2017). Coverage of Commercial Credit Database to Expand. <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2017/08/20170825-3/>

⁷³ Trade and Industry Department, the Government of the Hong Kong Special Administrative Region. (2019). Support to Small and Medium Enterprises. https://www.tid.gov.hk/english/smes_industry/smes/smes_content.html

Financial institutions in mainland China have successfully adopted machine learning to build a credit assessment model using alternative data. The People's Bank of China has established a Credit Reference Center⁷⁴ with a centralised database of credit data for financial institutions to perform risk assessments. Financial institutions are also utilising other transaction data and demographic information to build a machine learning model with the potential of partnering with third-party payment platforms.

Following the rapid expansion of e-wallets and other P2P platforms, a significant number of retail and business payments are now conducted on third-party online payment platforms such as Alipay and WeChat. FinTech is rapidly growing as an add-on service to these platforms where the transaction records can be used as alternative data for credit assessment. According to iResearch Consulting Group⁷⁵, FinTech now accounts for nearly 45% of third-party online payment transactions, reaching CNY 29.1 trillion in 2018. One of the interviewed banks stated that it had built a machine learning model with multiple layers made up of more than 1,000 variables derived from around 200 attributes directly obtained from customers. The bank recorded a 20% improvement in the performance of its SME lending and its default rate dropped by 50% after the AI-enhanced model was applied.

According to our survey, most banks have yet to apply AI to credit scoring as a way of supplementing traditional credit risk assessment. Only two applications have been launched, with three planned. However, banks are showing active interest in this area. Fourteen banks indicated that some applications are being studied, with reference to a number of successful implementation cases of AI-enhanced credit scoring outside Hong Kong.

One obvious way to increase the performance of credit risk models based on alternative data is to increase the volume and types of variables feeding into the model. Federated learning, a new trend in AI, was adopted by one bank that we interviewed to securely learn from the encrypted data captured on electronic invoices, also known as 'e-Fapiao'. These electronic invoices are shared by organisations

centrally in China to improve KYC assessment and credit risk measurement of SMEs. Raw data from invoice centres is encrypted and aggregated before being shared with banks, and banks can then use that data as training variables to improve their models, while the data itself remains encrypted. This ensures each partner remains the owner and retains control of their data without disclosing unnecessary data, while banks can still train the model¹⁸.

Besides introducing alternative data into the core credit risk assessment process, AI is also being adopted for wider credit risk management practices. For instance, ING has worked with Google and PwC⁷⁶ to create an Early Warning System (EWS) for credit risk using AI technologies. This is a good example of how AI can complement the already mature credit risk analytical framework. The EWS is designed to collect large amounts of structured and non-structured data (e.g. real time news feeds) into the big data infrastructure, then use Google's tools, such as NLP and machine learning models, to process and prioritise data that may impact on a customer's credit risk, which credit analysts can then further analyse. This enhances the performance of credit risk management analysts by helping them to handle large amounts of data in different structures and languages as well generating early warnings about particular customers or loans, prompting decisions and actions for managing the risk exposure.

The discussion above indicates that, although banks' credit risk management is mature and makes extensive use of data analytics, AI still has the potential to add value. While federated learning has proven promising for improving the performance of AI based credit assessment, it faces computing challenges because the underlying data encryption increases training times by a factor of hundreds or even thousands. Some technology start-ups are working on projects to overcome part of this problem with specialised hardware and technology architecture⁷⁷. Federating learning is expected to be widely adopted once the technology becomes more mature, and as data collaboration alliances or partnerships are formed among banks and owners of alternative data sources.

⁷⁴ Overview for the Credit Reference Center, the People's Bank of China. http://www.pbccrc.org.cn/crc/zxgk/index_list_list.shtml

⁷⁵ iResearch. (13 May 2019). China's Third-Party Online Payment Transactions Approached 30 Tn Yuan in 2018. http://www.iresearchchina.com/content/details7_54532.html

⁷⁶ PwC (December 2018). Rapid shift to a digital credit risk alert system. <https://www.pwc.nl/en/topics/about/articles-december-2018/rapid-shift-to-a-digital-credit-risk-alert-system.html>

⁷⁷ Kim, D., Yu, T., Liu, H. H., Zhu, Y., Padhye, J., Raindel, S., Seshan, S. (1 January 1970). FreeFlow: Software-based Virtual {RDMA} Networking for Containerized Clouds. https://www.microsoft.com/en-us/research/uploads/prod/2019/02/FreeFlow-nsdi19_prepub.pdf

5

AI implementation

Key Takeaways

To help organisations better implement AI solutions, PwC has developed a suggested high-level AI implementation journey of six major stages: Strategy, Planning, Ecosystem, Development, Deployment, and Operate & Monitor.

The AI implementation journey starts with developing an AI strategy. This involves building a corporate-level AI strategy, aligning the strategy with regulations and industry standards, and refining internal policies and practices accordingly. At the planning stage, factors such as the impact of AI on a bank's revenue and reputation need to be considered. At the ecosystem stage, banks will need to assess their available AI talent and purchase external solutions where necessary. When developing and deploying AI solutions, banks need to follow an agile but well-governed development and deployment lifecycle. Lastly, the operation of the AI solution

will deviate from traditional IT applications because of the need to monitor and calibrate it, and the fact that the model is governed by a series of explainable, technology risk management principles.

The foundation of AI implementation is governance. To establish proper oversight and governance, banks should ensure that controls are in place to track performance, risks and impacts.

To help organisations with their AI implementation and to mitigate risks, PwC has introduced different tools designed to support the measurement and assessment of AI implementation from the perspectives of interpretability and explainability, robustness and security, and fairness.



PwC's research estimates that AI could contribute \$15.7 trillion to the global economy by 2030, as a result of productivity gains and increased consumer demand driven by AI-enhanced products and services. AI solutions are spreading across industries and impacting everything from customer service and sales to back office automation. According to our survey, 95% of banks in Hong Kong are also considering making AI part of their corporate strategy.

However, as discussed in section 3.4, the lack of explainability in the AI model, the lack of resources for its adoption, and legal and compliance challenges are holding some banks back from embracing AI.


Banks need to overcome these barriers and manage risks when implementing AI-related technologies in order to remain competitive and cost efficient by comparison with the new challenger banks that are entering the market globally.

To help organisations identify and address challenges when designing and deploying AI applications, PwC has published *A practical guide to Responsible Artificial Intelligence (AI)*⁷⁸, which outlines a suggested high level end-to-end AI implementation journey.

Responsible AI implementation involves a clearly defined journey of six major steps: Strategy, Planning, Ecosystem, Development, Deployment, and Operate & Monitor. Such a framework to guide banks through the journey is important, and should address questions such as who is accountable, how and when will AI affect the business operation, what decisions need to be made and when, what controls need to be in place and when, and whether its overall change management is aligned to the bank's business goals.

Governance is a major concern for banks when implementing AI solutions. In this section, PwC will also propose eight major considerations for governance, and offer an illustrative governance standard as a reference for banks.

To address AI implementation challenges it is important to have tools and accelerators in place before AI development and deployment begin. Tools are needed to assess how well challenges are being monitored, to what extent gaps in the plan are being identified and resolved, and whether the AI solution is providing the benefits expected. The final section introduces a Toolkit and explains how banks can benefit from it to address these issues.



⁷⁸ PwC (2019). *A practical guide to Responsible Artificial Intelligence (AI)*. <https://www.pwc.com/gx/en/issues/data-and-analytics/artificial-intelligence/what-is-responsible-ai/responsible-ai-practical-guide.pdf>

5.1. Key considerations in defining a high-level AI implementation journey

This section describes a high-level end-to-end AI implementation journey from strategy to operation. Banks are invited to leverage this to plan how to develop and implement their AI strategy, how to build an ecosystem or collaborate, and how to develop, deploy, operate and monitor their AI solutions.

5.1.1. Strategy

The AI implementation journey starts with developing an AI strategy. This involves building a corporate-level AI strategy, aligning the strategy with regulations and industry standards, and refining internal policies and practices.

Corporate-Level AI Strategy

- Organisation
- System
- Process
- Talent

Regulations & Industry Standards

- GDPR
- PDPO
- Fairness
- Transparency
- Accountability
- Ethics

Internal Policies & Practices

- Data
- Talent
- Algorithm
- Infrastructure



1. Building a corporate-level AI strategy

In building a corporate-level AI strategy, banks should first inform stakeholders at different levels about how AI will be used to drive enterprise values and goals within the organisation. This can be done by conducting initial market research and strategic assessments on common AI use cases and their applicability to the banks, and performing an organisational assessment to identify major gaps in processes, people, systems and data. Banks can then quantify the enterprise value of AI by conducting financial analysis to determine the ROI for AI, and impact analysis on factors like risks, controls and talent. A strategic agenda

can then be outlined for both short-term and long-term activities, and priorities can be defined as to what needs to be done to enable the AI initiatives.

One key strategic agenda item is to develop an enterprise talent and organisational strategy for AI. Banks are recommended to assess their internal AI capabilities and identify skills gaps regularly as part of their overall AI strategy. To mitigate the shortage of AI talent, banks might formulate a talent strategy for continually investing in their digital talent pipeline, as well as building in-house skills through enterprise-wide training programmes. At the same time, they are recommended to define a new organisational structure that allows for efficient internal collaboration for AI development, since banking domain knowledge is needed to build up relevant AI models for specific use cases.

To facilitate communication and collaboration between AI specialists and the bank's business functions, banks can consider implementing a new cross-functional implementation team structure – DevOps. DevOps blends business analysts, developers and system operators in one integrated team focussing on one specific process. This reduces communication overheads and enables faster solution deployment, enabling banks to develop and deliver AI applications more quickly.

2. Aligning strategy with regulations and industry standards

When developing a corporate-level AI strategy, banks need to take regulations and industry standards into consideration. Banking is a highly regulated industry, and regulations relating to fairness, accountability, transparency, and ethics may affect the use of AI (this is further discussed in section 5.2). To ensure regulatory compliance, banks should survey existing and imminent regulations relating to AI, interpret them and establish a consensus on their likely regulatory impact, and validate their AI strategy relative to the regulations with key stakeholders. Banks should also survey and consider following existing international and local industry standards. It is recommended that the AI strategy is validated against the survey results with key stakeholders.


3. Refining internal policies and practices

After completing steps 1 and 2 above, banks are recommended to refine their internal policies and practices on data, infrastructure, algorithms, and talent to provide clear direction and guidance on the adoption of AI for different business units. Banks can create an inventory of existing policies and identify those which will impact on AI, and then come to a consensus on refining these

policies based on their AI strategy, along with regulations and industry standards. Taking data policy as an example, a bank may store its customer data and transaction data across different business functions and in different silo systems. However, AI requires large amounts of customer data for model training to obtain better customer insights. There is thus a need for clear guidelines on a data policy, and specifically on the sharing of customer data across different business functions.

5.1.2. Planning

At the planning stage, banks should first identify the highest value AI opportunities that align with their business strategy, and decide the extent to which the AI solution should be deployed. Given AI's heterogeneous nature, there are no standardised formulae for evaluating business use cases. However, creating a comprehensive user story can help banks determine the level of performance needed from an AI solution. To clarify whether AI is appropriate for particular use cases, PwC proposes the following six factors for consideration in AI implementation planning:



This assessment of business use case feasibility needs to be conducted collaboratively by technical and business teams to ensure all aspects are evaluated and gaps for implementation are identified.

5.1.3. Ecosystem

After the use cases are evaluated and confirmed, banks need to assess their current AI talent and decide on an implementation model. Developing AI solutions always requires specialists, such as data scientists or algorithm programmers, who may not always be readily available for banks. To ensure the AI model is developed and deployed smoothly, a well-defined technology roadmap is essential, one which ensures internal and external collaborations are properly considered and managed.

Banks may decide to adopt a mature external AI solution as part of their overall AI architecture (e.g. a Cantonese NLP engine) in cases where the solution is not data sensitive. In cases that require customer data for model training, such as robo-advisors, it is preferable for banks to create their own AI team to develop the solution.


During AI implementation, a bank may need to collaborate with other banks, or even non-bank parties, to source data

to fuel the AI engine. However, it is recommended that banks carefully assess the financial viability and the data access controls of the vendor before incorporating such AI solutions into their infrastructure, given the fast-moving vendor landscape.

It is also important for banks planning to deploy AI to start with a cohesive change management plan. Unlike traditional IT system implementation, the deployment of AI solutions could create uncertainties in the workforce. A robust change management framework that caters to the workforce is key to the success of any AI planning.

5.1.4. Development and deployment

After identifying the project objectives and carrying out thorough planning, banks need to follow an agile but well-governed development and deployment lifecycle. This includes the acquisition and selection of data, model training, model evaluation and testing, interpretation, and deployment.



The first step of the cycle is to define, collect and prepare all the relevant data for use in model training. With appropriate selection of the data to feed the AI algorithms, the output AI model will be more accurate and fairer in its predictions. The selected data also needs to be divided into training data and testing data for model training and testing respectively.

The next step is to select an AI algorithm and use the selected training data to train the AI model. The AI development toolsets in model training are different from traditional models. The latter usually consist of a single implementation of each algorithm in one commercial package. With multiple open source libraries and tools available for any given AI or machine learning algorithm, it is common for a modern analytics development process to build models relying on several different tools and packages which yield different results and performance. If there is no integrated model development environment, or at least some form of version control, there is a greater potential for inconsistency and errors.

Once the training is complete, it is time to check that the model produces accurate and fair results. This is where the data identified as testing data is used. The evaluation allows the bank to test the model against data that has not been used for training. The result enables the development team to see how the model might perform on data that it has not yet seen, and thus how it might perform in the real world.

The final step is to deploy the trained model. Due to the nature of machine learning, new AI models are frequently updated and deployed. Banks need a robust handover process between the development and operation teams to ensure the deployment process does not disrupt existing processes. This is typically done by conducting automated continuous testing before deployment.

Once deployed, production data will start to arrive. With more real data, user feedback and model adjustment, the model can be continuously improved.

5.1.5. Operate and monitor

The operation of the AI solution will deviate from traditional IT applications because of the need to monitor and calibrate it, and the fact that the model is governed by a series of explainable, technology risk management principles. The model review cycle should be designed to comprehensively address the entire model ecosystem.

Periodic reviews and ongoing monitoring are recommended so that banks can ensure robust execution and manage cybersecurity threats. This would include continuous review of models, re-validation of data and resources, and review of contingency measures such as human intervention or the use of conventional processes. The introduction of safeguard methods, such as scenario planning, is also recommended to prevent unintentional interactions. At the operating stage, as the model result will be affected by accumulated production data, it is also vital to maintain the quality and fairness of input data.

Below is a case study of how a bank in Hong Kong (code-named Omega Bank) used the above AI implementation approach to develop a robo-advisor.



Omega Bank launches a Robo-Advisor solution to provide existing customers with professional investment recommendations

Omega Bank wants to offer professional investment recommendations to ordinary customers without incurring a significant increase in operating costs. To achieve this, Omega Bank decided to develop a robo-advisor.

Its robo-advisor implementation journey started with strategy planning. Omega Bank began by establishing guiding principles on ethical and explainable AI with reference to different frameworks, and implemented an overall governance process. To validate whether the use cases aligned with its strategic objectives and guiding principles, Omega Bank engaged both its business and IT teams early in use case discussions. It took multiple factors into consideration when evaluating the benefits, feasibility and limitations of the solution.

As banking is a highly regulated industry, regulatory compliance is critical for emerging technologies like AI. Omega Bank therefore held several meetings to discuss controls to ensure compliance, and developed governance and risk management guidelines following the HKMA's guidelines on Online Distribution and Advisory Platforms. A data privacy framework was also developed to ensure customers' data would be safely stored and processed in the implementation lifecycle, in compliance with the PDPO.

Talent and partnerships are also critical factors for AI development. The bank considered its internal AI capabilities and evaluated which AI tools needed to be bought and which could be developed in-house. Considering the domain knowledge and data available in-house and its ability to maintain these, the bank decided to develop an in-house

machine learning model for investment recommendations, while buying an NLP solution to analyse English, Cantonese and Mandarin from a local start-up.

To develop the AI model, the bank first identified what data was required, including customer profile (e.g. age, annual income, gender, and industry), the customer's investment preference (risk preference, market preference, and product preference), and the customer's targeted return on investment, and other useful data. With both data and algorithm ready, the bank performed data labelling, model training with supervised machine learning, and model evaluation. A model with a satisfactory performance that outperformed the market index was developed after several rounds of testing and enhancement. To better explain the recommendations from the robo-advisor, a surrogate model was built using a decision tree to explain the key decisions of the robo-advisor, based on some key input variables.

Once the robo-advisor went live, it outperformed the market and began to achieve customer traction. However, after six months of deployment, the bank performed a regular review and found that the robo-advisor had underperformed in one month due to an outlier event. The in-house data scientists documented the findings and re-calibrated the model by changing the selection of input variables. They then evaluated the updated model through thorough back-testing before redeploying the model. With continuous improvement and re-calibration, it is expected that the robo-advisor will be able to perform to the standard of a professional investment advisor.

5.2. Key considerations in establishing a governance model

The foundation stone of Responsible AI is end-to-end governance. Because decisions made by a Chatbot or an ML-driven prediction system are made by machines rather than humans, for many banks it can be unclear how such decisions should be governed. The board or senior management of banks should appreciate that they remain accountable for all AI-driven decisions. To ensure regulatory compliance, it is critical for banks to establish proper governance and risk management frameworks.

To do this, banks first need to answer such questions as: How does AI align with our core business strategies? What controls should be in place to track performance? What are the risks and impacts that AI will bring to consumers, businesses, societies and nations?

To help banks better build or update their AI governance model, PwC has summarised eight aspects that banks need to consider:

Consideration	Risks	Details
Compliance	<ul style="list-style-type: none"> • Compliance risk • Data privacy risk • Transparency risk 	A compliance review process should be defined to keep track of regulatory changes and ensure the bank's policies and processes are compliant.
Ethical	<ul style="list-style-type: none"> • Value alignment risk • Global alignment risk 	Enterprise-wide definition of AI use cases, data collection and model training should be defined based on ethical principles
Interpretability & Explainability	<ul style="list-style-type: none"> • Black box risk • Compliance risk • Lack of value risk 	The decision of the AI model should be interpretable and explainable to solidify the decision and avoid a 'black-box' threat
Robustness & Security	<ul style="list-style-type: none"> • Adversarial attack • API model theft • Open source software risk 	Rigorous validation, continuous monitoring and maintenance, verification and adversarial modelling should be planned to improve the security and robustness of AI. Banks should cover all AI systems, data, and communications to ensure security.
Fairness	<ul style="list-style-type: none"> • Fairness risk • Reputation risk • Lack of value risk 	AI decisions should be fair and avoid discrimination on the basis of age, gender, religion, race, etc.
Economic	<ul style="list-style-type: none"> • Job displacement risk • Uncertain or unassessable ROI • Loss of institutional knowledge 	Key performance indicators should be established to ensure the realisation of benefits, and also to mitigate unwanted outcomes that are sometimes overlooked when deploying technology changes.
Societal	<ul style="list-style-type: none"> • Reputational risk • Intelligence divide risk 	Consumer data-privacy settings should be examined to prevent possible social backlash. Personally identifiable information (PII) data should be securely stored to prevent downstream individual implications.
Control	<ul style="list-style-type: none"> • Lack of AI talent to detect or control rogue actions 	Effective AI operating models and processes should be designed to improve accountability and the quality of results.

Among these considerations, compliance is one of the most important when setting up the governance model. As the banking sector is highly regulated, banks must monitor the regulatory environment in which they operate and understand how AI will shape their future business practices. As mentioned in section 1.3.1, banks need huge amounts of customer data to train the AI model, so data protection is a critical consideration. When developing AI applications, banks in Hong Kong need to ensure compliance with the Personal Data (Privacy) Ordinance (PDPO), any relevant codes of practice issued or approved by the Privacy Commissioner for Personal Data (PCPD) giving practical guidance on compliance with the PDPO, and any other applicable local and overseas statutory or regulatory requirements.

Banks are also recommended to strive to develop, implement, and use AI solutions that are morally responsible and ethically defensible. To make ethical principles more actionable, banks could contextualise ethical principles in the form of specific guidelines for front-line staff. By doing this, banks may be able to identify the ethical implications of their AI solutions and the relevant principles that should be taken into account when designing and implementing AI models, better enabling robust mitigation of ethical risks.

To ensure AI implementation is compliant with major regulations and to better visualise the above considerations, banks are recommended to establish a governance standard across their business units. In this way banks will have greater control and confidence throughout the six stages of AI implementation. Below is an example of what a governance standard across different business units and departments could look like:

Illustrative

Consideration	Business Units of a Bank							
	Branch	CASA	Leading	Credit	Treasury	Operations	Risk	Corporate
Compliance	10	10	10	10	10	10	10	10
Ethical	10	10	10	10	9	10	9	10
Interpretability & Explainability	10	10	8	8	9	8	9	10
Robustness & Security	10	10	10	10	10	10	9	10
Fairness	10	9	10	10	10	9	9	10
Economic	8	7	9	9	9	9	6	8
Societal	9	9	8	9	10	10	10	10
Control	8	8	9	8	8	9	8	9

Score from 1 to 10. The higher the score, the more the BU should pay attention to the consideration.

5.3. Measuring and assessing AI implementation

It is important to accurately measure and assess governance issues when implementing an AI solution. A Toolkit is introduced to help assess features of Interpretability and Explainability, Robustness and Security, and Fairness during AI implementation.

5.3.1. Interpretability and Explainability

As discussed in section 3.4.1, lack of explainability is a major barrier to banks adopting AI. Before deploying AI applications, banks need to ensure that decisions made by AI are explainable to all relevant stakeholders, including customers, regulators and business sponsors. Banks should implement adequate measures during the design phase to ensure a level of explainability that is appropriate and commensurate with the materiality of their AI applications.

Banks should consider the following when building explainable AI:

1. Confirm the level of explainability: the verification and validation process for different AI use cases may require different elements of explainability. For example, an AI system for targeted advertising requires a relatively low level of explainability, as the consequences of it going wrong are negligible. On the other hand, the level of explainability for an AI-based recommendation system (e.g. a robo-advisor) must be significantly higher, because any errors could not only lead to investment losses, but also damage the reputation of similar systems and even of the bank itself.
2. Start at the design phase: in most engineering processes, developers must consider the capabilities of systems at the early stages of the design phase. Explainability is no different – it needs to be considered up front and embedded into the design of the AI application. This will affect the choice of AI algorithm, the means of data processing, and the type of data to be collected.
3. Choose the right algorithm: different algorithms have different levels of explainability. Decision trees, for example, are highly amenable to explanation. It is possible to build commercially useful models where the entire decision process can be diagrammatically illustrated. Neural networks or deep learning, on the other hand,

while amenable to graphical analysis, contain many more connections and have more subtle properties with respect to node interactions that are inherently difficult to explain.

There are two types of interpretability: global and local interpretability. Global interpretability means that the developer can comprehend the entire model at once. This level of interpretability is about understanding how the model makes decisions, based on a holistic view of its features and each of the data inputs, such as data type and data weight. Local interpretability, on the other hand, means the conditional interaction between dependent and independent variables for a specific observation. The tools described below can help banks in measuring and assessing both global and local interpretability:

1. Global interpretability - surrogate modelling tool: this approach involves building another model alongside the trained model using a more explainable algorithm, e.g. a decision tree. It can mimic the output of the trained black box model and help data scientists to understand approximately how the black box model works. This model reveals which variables are most important, and how they affect the predictions.
2. Global interpretability - partial dependence plot: this approach evaluates how variables affect the prediction by simulating all the combinations of input variables. It provides a visualised diagram of the relationship between variables and prediction.
3. Local interpretability measurement tool: this tool quantifies the importance of a selective input variable by testing the sensitivity of the model's result to the value change of that input variable. This tool measures and sorts the importance of input variables, and thus provides a better understanding of how the model works.

5.3.2. Robustness and security

Robustness and security are two other vital factors in successful AI implementation. AI solutions need to be robust and secure to be effective and reliable. Banks are encouraged to implement effective cybersecurity measures to avoid or control potential cybersecurity threats such as data poisoning and adversarial attacks.

The following tools can help banks measure and assess robustness and security:

1. Model Stability Assessment Tool: this tool assists users in assessing the prediction stability and sensitivity of trained models. Banks can use this tool to check if a trained model would generate the same outputs for identical inputs and generate similar outputs for statistically-similar inputs after model updates.
2. API Theft Assessment Tool: this tool attempts to perform model extraction attacks on a ML model to see whether it is possible to replicate the model using an open API. As ML models are considered important IP, banks can use this tool to test if their models could be subject to cyber theft, and enhance their security accordingly.
3. Synthetic Data Generator: this tool learns from existing data and generates statistically similar data while maintaining data complexity. If banks cannot obtain enough training or testing data, this tool can help provide data for model testing or training.

5.3.3. Fairness

In all stages of the AI implementation life cycle, bias can be introduced. This could lead to unfair predictions based on race, gender, religion or other factors. Bias can cause decisions made by the model to lead to unfair or inappropriate risk levels or product offers, or even to contravention of anti-discrimination laws. Banks should ensure that AI-driven decisions do not discriminate or unintentionally show bias against any group of consumers. In addition, banks should ensure that AI models produce objective, consistent, ethical and fair outcomes to customers.

To ensure fairness in the AI model, banks need to formalise enterprise-wide governance and consistent quality control policies and processes for evaluating data sources, understanding how the data is organised, and testing it to ensure that it is representative and that it achieves fair outcomes for customers. There need to be procedures to detect whether the training data is sufficient and the data quality is good enough to build a robust and fair AI model. Other than evaluating the data, banks are also recommended to mitigate bias by leveraging tools and techniques that can improve the transparency and intelligibility of models.

While there is no universal definition of fairness (there are more than twenty definitions of fairness, according to different regulations and organisations⁷⁹), it is recommended that banks establish company policies

and select approaches that will support the fairness assessment according to their own definitions and requirements. The following bias detection tools can help banks to check whether a model discriminates against various groups or individuals given a test dataset, as well as to correct biased outcomes for various groups or individuals (e.g. age groups):

1. Bias Detection — Comparing Performance Metrics: this tool analyses performance of the model across different data groups, by calculating the rates of correct/incorrect predictions. Banks can use this tool to compute performance metrics for each possible group in the test dataset.
2. Fairness Detection Tool – Measuring Fairness: this tool compares the outputs of disadvantaged versus advantaged groups using fairness definitions, e.g. equal opportunity or Gini equality. Banks can use this tool to quantify discrimination against disadvantaged groups and solve unfairness issues.
3. Bias Intervention Tool – Thresholding: this tool computes the decision boundaries of fairness for discriminated groups by re-adjusting thresholds until disparities between the different groups are minimised. Banks can use this tool to refine the model to reduce bias.

⁷⁹Verma, S., & Rubin, J. (2018). Fairness Definitions Explained. <http://fairware.cs.umass.edu/papers/Verma.pdf>

6

AI for Hong Kong banking: looking forward

Key Takeaways

We here summarise four recommendations for stakeholders in the banking sector to consider in support of the industrial development of AI: building an AI-ready workforce; establishing an AI Centre of Excellence (CoE); fostering a Hong Kong banking AI community and infrastructure; and drafting risk guidelines.

To better prepare for the current and next wave of AI, banks should build up an AI-ready workforce. Options include reskilling existing employees, and encouraging employees to participate in professional training or AI-related co-university programmes. An AI CoE is defined

as a team within a company that provides leadership, best practice, research, support and/or training in a given area. Establishing an AI CoE to manage the bank's AI investments and project portfolio will help banks refine their operating models. Fostering Hong Kong's banking AI community and infrastructure can help support banks in adopting and applying AI technologies. Lastly, drafting guidelines on AI risk management in consultation with working groups of banks and regulators would also help to better manage regulatory compliance.



To facilitate AI development in Hong Kong's banking industry, having a comprehensive implementation and governance framework such as that provided in Chapter 5 is only the first step. The ultimate goal is to develop Hong Kong into an AI Innovation Hub which attracts talent and investment from around the globe so as to maintain the city's competitiveness in the digital era. In pursuit of this goal, we present four recommendations in support of the industrial development of AI for banks and other stakeholders (such as AI technology vendors and innovation incubators in the banking ecosystem).

These recommendations aim at tackling challenges and risks arising from AI adoption from the perspectives of internal preparation, external collaboration and regulatory guidelines. They include building an AI-ready workforce, establishing an AI Centre of Excellence, developing a Hong Kong banking AI community, and managing AI technology risk.

6.1. Building an AI-ready workforce

According to our survey, 75% of the respondent banks which have considered developing AI solutions say they would prefer to manage the research and development stage internally rather than outsourcing it. Some state that this is mainly because banking domain knowledge is

necessary to build a suitable model for specific business use cases. This indicates that having an internal workforce ready for AI adoption is a key factor in determining whether a bank can successfully adopt AI solutions to address its business problems. However, as mentioned in section 3.4.2, the shortage of appropriate talent is a major challenge faced by banks.

The following four recommendations are made with the aim of helping banks get their existing workforce AI-ready and increase their future talent supply.

Reskilling existing employees

Reskilling existing employees is a sustainable way for banks to tackle the problem of insufficient AI talent in Hong Kong, by transforming traditional banking professionals into AI-ready professionals. Such employees can provide a solid foundation for long-term AI development within the bank.

As an international financial centre, Hong Kong enjoys the unique advantage of having a large portfolio of banking talent in multiple disciplines and with diverse backgrounds and nationalities. It also has the breadth and depth in terms of banking services offerings to attract the required talent. One recommendation is therefore for Hong Kong to upgrade its existing banking workforce and to define a talent strategy to nurture and upskill non-AI professionals to become 'citizen AI users' and 'citizen AI developers'.

Citizen AI users	The majority of staff, including front line and operations, should be trained to be able to use the bank's AI-enhanced applications. They should also be able to support good data governance, and get expert help when needed.
Citizen AI developers	5%-10% of business unit professionals who are power users should be trained to identify use cases and quality data sets, and work closely with AI specialists in developing new AI applications.

Providing industrial training

Besides internal reskilling, the banking sector should also support industry-wide professional training.

Hong Kong has a long history of industry-wide organisations (such as The Hong Kong Association of Banks (HKAB) and other professional education providers) offering effective industry or subject-specific training to practitioners in the banking industry.

An AI curriculum could be added to existing programmes and these organisations could leverage their platforms to accelerate industry-wide upskilling. In our interviews with banks, many suggested that government funding

(e.g. extending the coverage of Hong Kong's Continuing Education Fund) could also be provided to subsidise training in order for Hong Kong's banking sector to rise above the AI baseline.

Merely upskilling the existing workforce will not overcome the AI talent shortage. Some interviewed banks noted that it is very difficult and costly to upgrade banking professionals to become advanced data scientists and data engineers. However, data scientists possessing advanced analytical skills, and data engineers who can create, deploy, and manage AI applications, are key to elevating Hong Kong's position in AI. Solving the root causes of this shortage through education and talent transfer programmes is the only effective route.

Co-University programmes

The statistics given below suggest that one of the reasons for the shortage of experts is the mismatch between the supply of students and demand from employers.

Currently, limited numbers of university programmes are heavily FinTech focused. According to the Hong Kong Joint University Programmes Admissions System (JUPAS) and local university websites, there were 29 FinTech related programmes in Hong Kong in 2019 (ten Bachelor's and nineteen Master's degree programmes) (see Appendix C).

However, a significant talent gap still exists within the industry because not enough graduates possess the relevant technical skills and capabilities. The 2019 survey and dialogue "Hiring in FinTech" by Michael Page⁸⁰ revealed that 64% of employers found the recruitment process for FinTech talent challenging, while 47% agreed that the key challenge was a shortage of proven skills for FinTech related roles. Moreover, 80.7% of respondents in our survey recognised 'Hard to find employees with AI expertise' as a significant barrier for AI adoption. The talent that is available lacks the exposure and academic training for both emerging technologies and banking operations.

Obviously, the supply and quality of talent in Hong Kong is inadequate to support the growing demand for technology specialists by established institutions and virtual banks. A lack of data scientists and computer scientists with domain-specific knowledge in banking is hindering the overall development of the sector.

Some of the banks we interviewed mentioned that culture and lack of career support could be reasons for this

shortage. While Hong Kong's universities are well known around the world, few overseas students are willing to pursue technology-related careers in Hong Kong.

To mitigate this problem and nurture the talent required to develop AI, the co-development of AI-related programmes by financial institutions and universities is recommended. In Hong Kong, the HKMA launched the FinTech Career Accelerator Scheme (FCAS) in 2016 in collaboration with the Hong Kong Applied Science and Technology Research Institute (ASTRI) to meet the growing need for FinTech talent in Hong Kong⁸¹. One of the programmes in FCAS is the Gap Year Replacement, which allows students from participating tertiary education institutions to work on FinTech-related projects full-time at banks or the HKMA for six months or one year. Banks could consider joining this programme or developing their own co-university programme (e.g. a one year internship) to expand the FinTech talent pool. With early exposure to real industry experience, students can build both banking and AI-specific knowledge, thus reducing the talent gap.

Talent/Capability transfer programmes

Although Hong Kong is well known for the quality of its academic institutions, it has started losing out in the competition for the global talent pool, according to a report issued by IMD World Talent Ranking in 2018⁸², falling from 12th to 18th.

To increase its competitiveness and relieve the pressure of this talent shortage, a wider market for talent should be considered. Programmes to import technology talent should be launched and accelerated. For example:

TechTAS	A three-year pilot scheme which provides a fast-track arrangement for eligible technology companies or institutes to admit overseas and Mainland technology talent to undertake research and development work.
AIR@InnoHK	The HK Innovation and Technology Commission (ITC) ⁸³ established an AIR@InnoHK ⁸⁴ programme in the second half of this year to support research into AI and robotics technologies and its application in areas such as financial services. The research may cover big data analytics, machine learning, cognitive systems, intelligent agents, etc.

Besides relying on external resources, one recommendation is to launch a programme to send locals overseas to learn from more developed economies. Hong Kong could thus leverage the successful experience of other markets. For example, Singapore launched a 'Capability Transfer

Programme' in 2017⁸⁵ which aims to transfer capabilities from foreign specialists to locals. The programme provides a 90% (up to SG\$300,000) allowance for sending eligible AI staff overseas for training, or for hiring overseas experts to train local staff.

⁸⁰ MichaelPage. Hong Kong: Hiring in Fintech Survey and Dialogue. https://www.michaelpage.com.hk/sites/michaelpage.com.hk/files/16785-hk_fintech_brochure_mp.v8.pdf

⁸¹ FCAS (2019). FinTech career accelerator scheme (Gap year placement programme). <http://www.fcas.hk/>

⁸² Leung, K. & Chung, K. (20 November 2018). Hong Kong falls six places to 18th in global talent ranking, trailing Singapore. <https://www.scmp.com/news/hong-kong/education/article/2174058/hong-kong-falls-six-places-18th-global-talent-ranking>

⁸³ Innovation and Technology Commission. Technology Talent Admission Scheme. <https://www.itc.gov.hk/en/techtas/faq.htm#a1>

⁸⁴ News.gov.hk (16 June 2019). Smart living attracts IT talent. https://www.news.gov.hk/eng/2019/06/20190613/20190613_132750_200.html?type=feature

⁸⁵ SGTech. Capability Transfer Programme. https://www.sgttech.org.sg/SGTECH/Web/Initiatives/Talent_and_Capabilities/Capability_Transfer_Programme.aspx

6.2. Establishing an AI centre of excellence (CoE)

Whichever stage banks are at on their AI journey, they will need to continuously discover ways to enhance their business using AI in order to justify their investment. Rather than just relying on the resources of their technology departments, banks should advocate AI adoption across the business to adopt the AI solution more quickly and with higher quality results. To achieve this, they should consider establishing an AI Centre of Excellence (CoE).

A CoE is defined as a team within a company that provides leadership, best practice, research, support and/or training in a given area. For AI development, a CoE could act as a centre which collects required resources for AI development, such as use cases and initiatives. It could also act as a centralised knowledge base which provides suggestions and recommendations on the AI technology to be chosen and diffused throughout workflows and processes⁸⁶.

For banks that are exploring the possibilities of AI, an AI CoE can help to accelerate the development and adoption of AI within the organisation. Establishing an AI CoE to manage the bank's AI investments and project portfolio will refine its operating model by allowing teams from the business and the IT department, as well as AI specialists, to collaborate in solving specific business problems. More re-usable models and AI components will be created as a result, and the ROI for investing in AI projects will become easier to calculate over time by referencing project successes and lessons learned.

At the same time, given the ethical issues that can arise from AI, an AI CoE should also emphasise the provision of ethical guidance. By establishing a centralised ethics-related position or review board, banks should be able to more easily align the AI technology with their broader ethical, governance and privacy regulations.

It is suggested that the AI CoE should be capable of facilitating the following:

Internal Facilitation	Facilitating AI development by increasing user acceptance: <ul style="list-style-type: none">AI Education Programme: To develop and execute enterprise-wide educational programmes on key AI tools, platforms, use cases and regulations, covering every level from board members to front line staff.
Model Development	Building up a solid foundation and infrastructure for model development: <ul style="list-style-type: none">Reusable Digital Assets: To create inventories and advocate AI digital assets be developed in different projects and initiatives.Implementation Support: To provide advanced AI specialist resources to support AI development. Ideally the AI CoE should create and manage a one-stop digital platform for collaboration, support and resource management, representing a system environment with pluggable tools, shared data sets, and documentation on reusable components, methodologies and lessons learned.
Initiative Management	Helping manage and monitor all the AI initiatives, providing a clear direction in strategy setting, and serving as a centralised knowledge centre for AI implementation: <ul style="list-style-type: none">Use Case Sharing: To develop and share internal use cases of AI, and to research and introduce AI use cases implemented by peer banks to the organisation.AI Project Portfolio Management: To maintain the AI project portfolio, including ROI tracking and to create AI project prioritisation metrics according to the bank's corporate strategy.
Ethical Management	Guiding the business and development team in developing ethical AI solutions: <ul style="list-style-type: none">Ethical guidelines: To build comprehensive ethical guidelines according to governance, regulations and industry standards for AI solution development.Ethics check: To help test if the AI solution is ethical based on fairness, impact on customers, algorithmic bias and other factors.

⁸⁶ Ramel, D. (17 January 2019). How to create an enterprise "AI Centre of Excellence". <https://pureai.com/articles/2019/01/17/ai-coe.aspx?m=1>

6.3. Fostering Hong Kong's banking AI community and infrastructure

Throughout this research, a recurring challenge cited by banks has been the lack of AI use cases available to help them understand how to start adopting AI. A failure to identify beneficial business use cases brings challenges when justifying investment in AI, thus hindering its development. Setting up a CoE is one solution, as discussed in the previous section, while another is collaboration between banks. By understanding the state of play across the industry and leveraging others' experience, slower market followers can catch up.

To promote and support industry-wide collaboration, a neutral forum should be established to encourage banks in Hong Kong to adopt and apply AI technologies. The objectives of the forum would include:

- To facilitate the discussion and identification of use cases, lessons learned, and challenges faced by banks when implementing AI.
- To evaluate open source and other commercial solutions in terms of their suitability for banks and their compliance with banking regulations.
- To solicit and coordinate an industry response to relevant regulatory and industry consultations on AI.
- To co-create a digital framework on advanced technologies (e.g. federated learning).

Apart from including all authorised institutions in Hong Kong, certain other stakeholders should be involved in the forum so as to foster a more comprehensive banking ecosystem. For example, involving research organisations (such as local universities and ASTRI) and technology companies would help accelerate discussions on how banks could utilise the newest AI technologies and develop innovative use cases.

To achieve its objectives and generate new insights, the forum should host regular meetings to discuss both technical and non-technical issues. The forum should also serve as a knowledge centre. A document library could be established containing documented use cases and technical specifications. At the same time, quarterly reports could be generated to update participants on AI and financial developments. Besides discussion and reference, the forum could develop a sandbox infrastructure where talent nominated by member banks could co-create digital assets or AI prototypes. The findings from these pilots could

then be shared with other forum participants.

At the same time, the AI technology providers we interviewed also found it challenging to showcase their solutions in a banking context without access to actual banking data and systems. The need for a general purpose AI infrastructure that would enable banks to start their AI experiments and journeys is clear. The sandbox-like infrastructure could provide a simulated banking environment for the technology vendors. This would also help banks better understand and evaluate their AI solutions. Programmes offered by the Hong Kong Science and Technology Park and other organisations may help, but the group needs to be able to formulate a plan to utilise public resources in order to advocate the AI agenda in banking.



6.4. Drafting guidelines on risk management in AI

Another challenge for banks when pushing forward their AI agenda is uncertainty over regulatory risk. This has led to banks taking a conservative approach to AI deployment due to worries over non-compliance after significant investment. The deployment of AI technology impacts a number of existing practices that are subject to regulation.

In light of this, the HKMA's *High-level Principles on Artificial Intelligence circular*³⁸ provides guidance to the banking sector on the use of AI applications. The twelve principles listed in the document cover governance, application design and development, and on-going monitoring and maintenance for banks to apply in proportion to the nature of their AI applications and the level of risk involved. In

addition, the HKMA's *Consumer Protection in respect of Use of Big Data Analytics and Artificial Intelligence (BDAL) by Authorized Institutions circular*³⁹ recommends attention in four major areas, namely (i) governance and accountability, (ii) fairness, (iii) transparency and disclosure, and (iv) data privacy and protection.

To further facilitate communication on risk management, the HKMA launched its FinTech Supervisory Chatroom in 2017⁸⁷ with the aim of providing feedback to authorised institutions on their technology initiatives at an earlier stage and with a shorter response time (i.e. within seven working days).

With greater guidance, banks will become more confident and readier to apply AI solutions. The following are some topics that may benefit from discussion in the chatroom:

Topic	Description
Use of Production Data	Guidance on the controls regarding using production data for AI development and the sharing of production data when collaborating with third party AI technology providers.
Human decision controls on AI models	Expectations on the degree of human controls in place, and particularly on the degree of straight-through processing by the AI model that is allowed.
Deploying cloud-based AI solutions/models	A cloud platform provides flexibility and scalability in terms of the storage and computation power required for AI development. Many third-party AI solutions are available on cloud platforms only. Banks would appreciate greater clarity on the permissible scope of cloud deployment.
TSP and outsourcing arrangements	As with concerns over Open Banking APIs, clarity on controls and safeguards for straight-through processing and outsourcing arrangements (e.g. tagging of data for model learning) will help banks to design appropriate ecosystems to deploy AI.

⁸⁷ HKMA (28 November 2017). Fintech Supervisory Chatroom. <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2017/11/20171128-4/>



7

Conclusion

Artificial Intelligence is a term that has been in use for over half a century. Despite its familiarity, for much of its lifetime this nascent technology has failed to live up to its billing. While computing power grew exponentially, advances in AI were modest and incremental.

However, as this paper has outlined, a number of critical factors came into alignment in recent years and transformed something long anticipated into a reality. The consequent pace of technological progress has been astonishing, as anyone who has used freely available online translation software in the last couple of years will have noticed.

What brought this about? First, there was the transition from rule-based to non-rule-based learning, then the development of complex algorithms and – critically – the ‘fuel’ of vast amounts of data to facilitate machine learning. As a result, the basic building blocks, components and techniques enabling AI to support banking and other industries are now in place.

However, the ongoing technical difficulties of harnessing AI are now combined with other challenges. These include user acceptability, finding and retaining expert talent, integrating newly-enabled products and services into a well-established business strategy, and so on. One of the most critical issues for financial institutions – and any other organisation built on trust – is the concept of ‘explainability’. Encouraging familiarity with, and eventually confidence in, ‘black box’ processes is a daunting challenge from both a

risk management and a customer experience perspective.

Despite this, the potential to dramatically improve customer experience, reduce costs and better manage risk, is something no bank can ignore. The drivers for AI adoption are just too compelling. Further, the research cited in this paper shows that there is also a high degree of ‘market pull’, as customers demand more personalised and relevant tools to manage their financial and other goals.

To fully unleash the potential of AI in the banking industry, best practice needs to be leveraged and learned from. Hong Kong can develop into a leading AI hub by drawing on some of its core strengths, such as its deep and long-established expertise in banking and professional services. However, other aspects of the city’s human capital will also need to be intensively nurtured.

Few players in Hong Kong’s banking industry would claim to have a fully-fledged AI strategy or a thorough, all-round understanding of how the regulatory landscape will change because of AI. For these and other reasons, knowledge sharing of tried and tested use cases across Hong Kong’s banking AI community is essential. This should encompass success stories, as well as stories of those that fell short, either technically or due to a lack of real demand. In this way Hong Kong’s financial institutions will be able to develop an industry-wide understanding of the rules of the road for this new technology.





Appendix A: Glossary of Terms and Definitions

Terms	Definition
4G	4G is the fourth generation of broadband cellular network technology. A 4G system provides capabilities as defined by the International Telecommunication Union (ITU) in its IMT Advanced Standard.
5G	5G is the fifth generation of cellular network technology. It operates on three spectrum bands that provide faster download and upload speeds.
AIOps	AIOps stands for artificial intelligence for IT operations. It refers to the multi-layered technology platforms that automate and enhance IT operations. It utilises big data, machine learning and other advanced analytics technologies to enhance IT operation functions with proactive and dynamic insights.
Alternative Data	Alternative data refers to some non-traditional financial transactions and behaviours used by banks to gain additional insights into a company or investment.
Application Programming Interface	An application programming interface (API) is an interface or communication protocol for programmers to build software applications. It has been described as a “contract” between applications, ensuring that one application will always get a response in a specific format from the other after initiating a defined action.
Artificial Intelligence	Artificial intelligence (AI) is the intelligence demonstrated by machines when a machine mimics the cognitive capabilities of the human mind, such as learning, reasoning and recognising.
Big Data	Big data means the expansion of data in terms of volume (size of the data), variety (type and nature of the data) and velocity (speed of data generation and processing).
Chatbot	A Chatbot is an artificial intelligence (AI) software that simulates conversation via text or speech. It is commonly used in dialogue systems for various practical purposes, including customer service or information acquisition. Some Chatbots use sophisticated natural language processing (NLP) systems to process conversations.
Classification	Algorithm to predict the which one of the finite number of categories the model target belongs to.
Cloud Computing	Cloud Computing is the use of a network of remote servers hosted on the internet to store, manage and process data.
Clustering	A clustering algorithm divides data into groups (a.k.a. clusters) with similar traits and features.
Computer Vision	Computer vision is a field of artificial intelligence that trains computers to interpret and understand the visual world. It trains machines to accurately identify objects and respond accordingly.

Terms	Definition
Decision Tree Classification	Decision tree learning trains the machine learning model to make a prediction or classification decision based on a sequence of tests. The decision tree starts with a test of the inputs. Each test can have multiple outcomes, and these outcomes can lead to new tests (represented as new branches on the decision tree) or to conclusions about the target variables (represented as leaves on the decision tree).
Deep Learning	Deep learning is a subset of machine learning that is capable of recognising and learning subtle structural patterns from vast amounts of unstructured data. It mainly relies on a neural network (NN) to 'learn' the patterns from billions of input parameters by modelling hundreds of thousands of deep neural network hidden layers.
DevOps	DevOps is a set of practices that combines software development and IT operations. It automates the processes between software development and IT teams to shorten the development life cycle and provide continuous delivery with high software quality.
Disaster Recovery	Disaster recovery (DR) is a process to recover an organisation's operations with the aim of minimising the impact of an event on the organisation's IT infrastructure. DR enables the organisation to maintain business continuity or quickly resume mission-critical functions following a disaster.
Edge Computing	Edge computing is a system architecture that moves part or all of a computing process running from a centralised server to a local device (i.e. edge), such as a user's computer or mobile device. It minimises the amount of long-distance communication that has to happen between a client and server, hence reducing total network bandwidth and process running time.
Emotion Detection	Emotion detection is the process of identifying human emotion, usually from facial and verbal expressions.
Explainable AI	Explainable AI (XAI) refers to the methods and techniques that help to make the results of a solution understandable to human experts, addressing the issue of black box decisions.
False Negative	A case in which a model mistakenly predicts a negative. For example, an AML prediction model infers that a transaction is not money laundering when it actually is.
False Positive	A case in which a model mistakenly predicts a positive. For example, an AML prediction model infers that a transaction is money laundering when it actually is not.
Federated Learning	Federated learning is a specific category of distributed machine learning approaches which trains machine learning models using decentralized data residing on end devices such as mobile phones.
Handwriting Character Recognition	Handwriting Character Recognition (HCR) receives and interprets intelligible handwritten inputs. It involves the automatic conversion of text in an image into letter codes which are usable within computer and text-processing applications.
Intelligent Process Automation	Intelligent process automation (IPA) is the collection of technologies that combines process automation with Robotic Process Automation (RPA) and Machine learning (ML).

Terms	Definition
Internet of Things	The Internet of Things (IoT) is a category of technologies that connect devices and enable them to exchange data over a network. The 'things' include mobile phones, smart electrical appliances, wearable devices, and sensors embedded into mechanical or electronic components of machines.
Interpretability	The degree to which the model is understandable by a human user. This is usually measured by the model's explainability, transparency, and provability.
Label	A label is the prediction target—the y variable in simple linear regression. The label could be a continuous variable (e.g. stock price) or discrete categories (e.g. fraud cases)
Linear Regression	Linear regression is a simple statistical method to model the direct proportional (i.e. linear) relationships between independent input variables and the dependent output variables. The model can be trained to predict the output variables' values by inferring the linear relationships
Logistic Regression	Logistic regression is a variant of linear regression for classification tasks, e.g. which output variable is categorical (e.g. sunny, rainy, cloudy) rather than continuous (e.g. degree in temperature)
Long Short-Term Memory Neural Network	Long Short-Term Memory neural network (LSTM) is a recurrent neural network (RNN) architecture that is capable of handling sequential relationships within unstructured data. In contrast to simple RNN, which usually predicts using short term memory, LSTM has a mechanism to define and refresh (add and forget) a longer term memory.
Machine Learning	Algorithm to help machines learn relationships from training data and infer rules for prediction.
Narrow AI	Narrow AI is a machine intelligence that allows a machine to apply a certain level of cognitive capabilities within a specific context or problem.
Natural Language Generation	Natural Language Generation (NLG) is an AI technique that uses computational linguistics to create human-like textual content from structured data.
Natural Language Processing	Natural language processing (NLP) is a group of AI techniques used to read, recognise, process, analyse and understand human languages so that computer algorithms can enable human-computer interactions using natural language, including language translation, sentiment analysis and interactive conversation
Neural Network	A neural network (NN) is a machine learning model with multiple hidden layers of interconnected prediction units inspired by the interconnected neurons in the human brain. Based on the number of hidden layers, NN can be further classified into a shallow neural network (SNN) or a deep neural network (DNN). DNN can build model using huge (e.g. millions) of input variables and represent the hierarchical patterns of inputs at various levels of abstraction and semantic concepts.
Optical Character Recognition	Optical character recognition (OCR) is a technique to convert handwritten or printed texts on image files (such as a scanned document or photo) into machine-readable text.
Outlier	An outlier is an observation point that is distant from other observations in the data having similar input variables.
Personally Identifiable Information	Personal Identifiable Information (PII) is any information that can be potentially used to distinguish one person from another and can be used for de-anonymizing anonymous data.

Terms	Definition
Power User	A power user refers to the user of computer hardware, software or a related device who is capable of using its advanced features.
Regression	Algorithm to predict continuous output based on a set of input variables
Reinforcement Learning	Reinforcement learning (RL) is a machine learning model that trains machine agents to reinforce actions taken under specific conditions so as to maximise cumulative rewards
Reward	In reinforcement learning, a Reward (or reward function) is a mathematical function defined by data scientists to score and feedback on the machine agents' behaviour. The preferred machine behaviours are predefined by humans in the reward function. It gives rewards to machine agents when their behaviours are aligned with the reward function, and penalises machine agents otherwise. The aim of the reward function is to reinforce the behaviour of the machine agent under specific conditions through continuous feedback.
Robo-Advisor	A robo-advisor is a financial advisor platform that provides automated, algorithm-driven financial planning.
Robotic Process Automation	Robotic Process Automation (RPA) is the use of software that acts as a metaphorical software robot or AI worker to handle high-volume, repetitive tasks.
Robustness	The ability of the model to withstand or overcome adverse conditions and maintain consistent performance and accuracy of predictions.
Semi-Supervised Learning	Semi-supervised learning is a class of machine learning that makes use of a training dataset combining a small amount of labelled data with a large amount of unlabelled data, through bootstrapping the labels. It is a combination of supervised learning and unsupervised learning.
Sentiment Analysis	Sentiment analysis refers to identifying, extracting, and qualifying affective states and subjective information using natural language processing, text analysis, computational linguistics and biometrics.
Supervised Learning	Supervised learning is a type of machine learning that learns the relationship of inputs and outputs and infer the rules to predict the output given a set of inputs. The model is trained using labelled training data consisting of a set of input-output pairs.
Synthetic Data	Synthetic data is artificially generated data which are not obtained directly by observing real-world events or actual measurement. Synthetic data is created algorithmically to closely mimic the characteristics of production data but is anonymised so that personally identifiable information are excluded.
Test Dataset	A test dataset refers to datasets for assessing the performance of a trained model. It is independent of the training dataset, but follows the same probability distribution.
Training Dataset	Training datasets are datasets of examples used for learning which fit the parameters of the model.
Transfer Learning	Transfer Learning is a class of collaborative machine learning that transfers a model trained for a specific task and corresponding parameters and uses them as the starting point for model training on a second related task, so that model training on the same set of data does not have to be repeated
Unsupervised Learning	Unsupervised Learning is a type of machine learning that learns from non-labelled and non-categorised training data. It aims to identify commonalities of features in the data.


Appendix B: Survey results





The purpose of the questionnaire was to collect high-level information from selected Authorized Institutions under the HKMA about the status of AI in the banking sector in Hong Kong.

The survey contained seven sections: ‘AI Use Cases’, ‘Resources investment’, ‘Incentives and benefits’, ‘Barriers to AI adoption’, ‘Risks and accountability’, ‘Ethical issues’, and ‘Regulations related to AI adoption’.

The team has analysed and summarised the data in the survey and generated the following insights for retail banking as input for this research paper.

Retail Banking Industry




No.	Questions	Retail Bank
3	How much capital will you invest in AI in the next 5 years?	Total capital investment in AI will increase by 70%
4	Are banks expanding their AI workforce in the next 5 years?	 <p>92% 8% ■ Yes ■ No</p>
5	Are banks developing AI capabilities through partnering with external technology firms for AI implementation?	 <p>95% 5% ■ Yes ■ No</p>
6	Are banks developing AI capabilities through managing the research and development stage internally?	 <p>82% 18% ■ Yes ■ No</p>
7	Are banks with AI adoption utilising AI to shape their corporate strategy?	 <p>95% 5% ■ Yes ■ No</p>

No.	Questions	Retail Bank
8	What are the reasons for banks to apply AI technology?	<ol style="list-style-type: none"> 1 Improving customer experience 2 Strengthening risk management 3 Reducing costs
		Lack of employees with AI expertise 70%
		Insufficient data 52%
9	What are the barriers affecting banks' adoption of AI?	<p>Design ethics of AI 48%</p> <p>Data privacy & security 44%</p> <p>Legal compliance challenges 44%</p>
		Lack of expertise or resources 75%
10	What are the major AI risks identified by banks?	<p>Biased decisions of the AI models 74%</p> <p>Lack of quality data 65%</p>

The team also created heatmaps for some of the questions to generate better insights from the results.

Barriers to AI adoption (In Section 3)




This heat map has been created based on banks' responses to the extent of influence of 27 barriers to AI adoption identified in the survey. Five extent options were given to the respondents, namely, 'very serious', 'serious', 'moderate', 'minor', and 'ignorable'.

The score of each barrier has been calculated by the weighted average of each option with their corresponding multipliers. The higher the score, the more the barrier is affecting AI adoption within the banking industry. The following is the formula used to calculate the score.

	Multiplier
Very serious	4
Serious	3
Moderate	2
Minor	1
Ignorable	0

$$\text{Score} = \frac{(\text{Count of 'very serious'} \times 4) + (\text{Count of 'serious'} \times 3) + (\text{Count of 'moderate'} \times 2) + (\text{Count of 'minor'} \times 1) + (\text{Count of 'ignorable'} \times 0)}{\text{Total number of counts}}$$

AI use cases adoption (In section 4)



This heat map has been created based on banks' responses to the four stages of adoption, namely 'launched', 'planned to be launched', 'being explored and studied', and 'no plan to explore or study' of 27 use cases identified in the survey. Each use case is mapped back to the 25 banking functions categorised by the team as follows.

Banking function	AI use cases
Core Banking – Account Management	
AI on the maintenance of bank account or to provide value added services (e.g. cash management) to customer	<ul style="list-style-type: none"> Client-facing Chatbots
Core Banking – Product Innovation & Pricing	
AI to identify new product opportunities; or to support in product pricing to maximize value	<ul style="list-style-type: none"> Financial product pricing
Core Banking – Collateral Management	
AI to improve effectiveness and efficiency in managing collection	<ul style="list-style-type: none"> Collection information management
Customer – Customer Services	
AI to provide direct or indirect (supporting AI) customer services	<ul style="list-style-type: none"> Client-facing Chatbots Sentiment analysis
Customer – Customer Acquisition	
AI to acquire new customers	<ul style="list-style-type: none"> Personalised advertisement and recommendations Remote onboarding
Customer – Segmentation & Insights	
AI to develop better insights on what customer needs and their segmentation	<ul style="list-style-type: none"> Customer churn prediction Personalised advertisement and recommendations Transaction Data Enrichment (TDE)
Channel – Mobile	
AI applications in mobile banking channel	<ul style="list-style-type: none"> Remote onboarding
Channel – Internet	
AI applications in internet banking channel	<ul style="list-style-type: none"> Remote onboarding
Channel – Call Center	
AI applications to support call center operation e.g., identification, routing and management of in/outbound call, customer services, customer sentiment analysis, etc.	<ul style="list-style-type: none"> Sentiment analysis
Channel – Branch	
AI applications to support branch, priority banking, teller and ATM operations	<ul style="list-style-type: none"> ATM cash management Anti-Money Laundering (AML)
Corporate Functions – Fraud Detection and Prevention	
AI in fraud prevention, detection and management	<ul style="list-style-type: none"> Fraud detection Know-Your-Customer (KYC) Transaction Data Enrichment (TDE)
Corporate Functions – Technology Management	
AI in technology management e.g., automatic coding, DevOps, power consumption management, etc.	<ul style="list-style-type: none"> Cyber security applications Intelligent identification system
Corporate Functions – Clearance and Settlement	
AI in clearance and settlement operations to mitigate operational risk and / or increase efficiency	<ul style="list-style-type: none"> Contract analyser Operational automation Settlement
Corporate Functions – Compliance Documentation & Reporting	
AI in compliance and regulatory reporting processes	<ul style="list-style-type: none"> Contract analyser Operational automation Report generation using Natural Language Generation (NLG)

Banking function	AI use cases
Corporate Functions – Talent Management	
AI in talent acquisition, identification, development and / or retention processes	<ul style="list-style-type: none"> Surveillance system for employee conduct
Corporate Functions – Risk Assessment and Monitoring	
AI in risk appetite setting, policy setting, risk assessment, monitoring and / or mitigation	<ul style="list-style-type: none"> Economic forecasting / nowcasting Risk management
Corporate Functions – Finance & Reporting	
AI in finance and reporting processes in FIs	<ul style="list-style-type: none"> Report generation using Natural Language Generation (NLG)
Corporate Functions – Payment	
AI in both performing payment transaction and innovation and collaboration with non-FIs in new business model	<ul style="list-style-type: none"> Nil
Propositions – Corporate Lending	
AI in lending assessment, monitoring, fault management, pricing, etc.	<ul style="list-style-type: none"> Credit scoring
Propositions – Consumer Financing	
AI in consuming financing products, operations and management	<ul style="list-style-type: none"> Credit scoring
Propositions – Investment & Securities	
AI in investment advisory, asset management and securities trading services	<ul style="list-style-type: none"> Algorithmic trading Automatic trade execution Economic forecasting / nowcasting Capital optimisation and portfolio management
Propositions – Wealth Management	
AI to in wealth management products, operations and management	<ul style="list-style-type: none"> Economic forecasting / nowcasting Financial advice Transaction Data Enrichment (TDE)
Propositions – Cash Management	
AI in cash management services provided to customers	<ul style="list-style-type: none"> ATM cash management
Propositions – Cards	
AI in managing cards business	<ul style="list-style-type: none"> Card management systems
Propositions – Investment Banking	
AI in investment banking e.g., opportunities identification, risk identification, and management, trading, talent management, internal operation, etc.	<ul style="list-style-type: none"> Nil

The AI adoption score for each banking function has been calculated using weights associated with each of these adoption stages. The higher the score, the more use cases have been identified for that function and the more mature the AI capability. The following is the formula used to calculate the score.

	Multiplier
Launched	3
Planned to be launched	2
Being explored and studied	1
No plan to explore or study	0

(Count of 'launched'×3)+(Count of 'planned to be launched'×2)+(Count of 'being explored and studied'×1)+(Count of 'no plan to explore or study'×0)

$$\text{Score} = \frac{\text{Total number of counts}}{\text{Total number of counts}}$$



Appendix C: Summary of local FinTech-related university programmes

The purpose of this study is to investigate the AI talent shortage in Hong Kong. There were 29 FinTech related programmes in Hong Kong in 2019 (ten Bachelor's and nineteen Master's degree programmes). These courses are classified as FinTech-related because the overall content is heavily technology and finance-focused. Graduates are expected to have extensive knowledge in areas such as data science, business analytics and information systems, and possess technical capabilities in various software technologies and programming languages.

Undergraduate/ Graduate	AI/Fintech Course	University	No. of intakes / graduates	Reference Link
Undergraduate	Computational Finance and Financial Technology	City University of Hong Kong	7	https://www.jupas.edu.hk/en/programme/cityu/JS1000/
Undergraduate	BEng in Artificial Intelligence: Systems and Technologies	The Chinese University of Hong Kong	30	https://www.jupas.edu.hk/en/programme/cuhk/JS4468/
Undergraduate	BSc in Statistics	The Chinese University of Hong Kong	N/A	https://www.sta.cuhk.edu.hk/Programmes/UndergraduateStudies/BScinStatistics.aspx
Undergraduate	Bachelor of Arts and Sciences in Applied Artificial Intelligence	The University of Hong Kong	15	https://www.jupas.edu.hk/en/programme/hku/JS6224/
Undergraduate	Bachelor of Science (Major in Statistics)	The University of Hong Kong	74	https://saasweb.hku.hk/programme/stat.php
Undergraduate	BBA Information Systems	The Hong Kong University of Science and Technology	18	https://www.jupas.edu.hk/en/programme/hkust/JS5314/
Undergraduate	BSc in Data Science and Technology	The Hong Kong University of Science and Technology	N/A	https://dsct.ust.hk/
Undergraduate	Bachelor of Science (Honours) in Data Science	Lingnan University	25	https://www.jupas.edu.hk/en/programme/lingnanu/JS7225/
Undergraduate	Information Systems and e-Business Management Concentration	Hong Kong Baptist University	N/A	https://bba.hkbu.edu.hk/eng/programmes/curriculum/isem.jsp
Undergraduate	BSc (Hons) Scheme in Computing	The Hong Kong Polytechnic University	70	https://www.jupas.edu.hk/en/programme/polyu/JS3868/
Graduate	Master of Science in Finance (FinTech and Financial Analytics)	Hong Kong Baptist University	50	https://gs.hkbu.edu.hk/programmes/master-of-science-msc-in-finance-fintech-and-financial-analytics

Undergraduate/ Graduate	AI/Fintech Course	University	No. of intakes / graduates	Reference Link
Graduate	Master of Science In Financial Technology	The University of Science and Technology	50-60	http://www.msclfintech.ust.hk/program/design
Graduate	Master of Science in Business Analytics		50	http://www.msba.ust.hk/program-courses/program-design
Graduate	MSc in Big Data Technology		N/A	https://seng.ust.hk/academics/taught-postgraduate/msc-bdt
Graduate	Master of Science in Financial Technology	The Chinese University of Hong Kong	50	http://fintech.erg.cuhk.edu.hk/programme-information/objectives
Graduate	MSc in Data Science and Business Statistics		35	https://www.sta.cuhk.edu.hk/Programmes/PostgraduateStudies/MScinDataScienceandBusinessStatistics.aspx#overview
Graduate	Master of Data Science	The University of Hong Kong	50	https://www.scifac.hku.hk/prospective/tpg/MDASC#About-the-programme
Graduate	Master of Statistics		70	https://www.scifac.hku.hk/prospective/tpg/MStat
Graduate	Doctor of FinTech	The Hong Kong Polytechnic University	20	https://fb.polyu.edu.hk/dfintech/programme/overview/
Graduate	MSc In Information Technology		50	https://www.comp.polyu.edu.hk/en-us/programmes/detail/7
Graduate	MSc In E-Commerce		50	https://www.comp.polyu.edu.hk/en-us/programmes/detail/5
Graduate	MSc In Software Technology		50	https://www.comp.polyu.edu.hk/en-us/programmes/detail/8
Graduate	Master of Science in Business Analytics		N/A	http://www51.polyu.edu.hk/eprospectus/tpg/2020/23090-maf-map
Graduate	MSc Business and Data Analytics		140	http://www.cb.cityu.edu.hk/postgrad/?category=Programme&admitProgCode=P84
Graduate	MSc Business Information Systems		60	http://www.cb.cityu.edu.hk/postgrad/?category=Programme&admitProgCode=P05B
Graduate	MSc Electronic Business and Knowledge Management	City University of Hong Kong	60	http://www.cb.cityu.edu.hk/postgrad/?category=Programme&admitProgCode=P16
Graduate	MSc Electronic Commerce		45	http://www.cb.cityu.edu.hk/postgrad/?category=Programme&admitProgCode=P17
Graduate	MSc Financial Mathematics and Statistics		40	https://www.cityu.edu.hk/pg/programme/p68
Graduate	Master of Science in Data Science		100	https://www.cityu.edu.hk/pg/programme/p70



Acknowledgements

Authors and editorial team

Hong Kong Monetary Authority commissioned PricewaterhouseCoopers to conduct this research and publish the results. The following authors and editorial team members have individually contributed significantly to this research paper.

No	Organisation	Name	Role
1	Hong Kong Monetary Authority	Edmond Lau	Project Sponsor
2	Hong Kong Monetary Authority	Colin Pou	Project Sponsor
3	Hong Kong Monetary Authority	Nelson Chow	Project Lead
4	Hong Kong Monetary Authority	Brian Lam	Research Lead
5	Hong Kong Monetary Authority	Frederick Cheung	Research Analyst
6	Hong Kong Monetary Authority	Hinson Law	Research Analyst
7	Hong Kong Monetary Authority	Natalie Chan	Research Analyst
8	Hong Kong Institute for Monetary and Financial Research	Giorgio Valente	Research Lead
9	Hong Kong Institute for Monetary and Financial Research	Chris Ip	Research Analyst
10	Hong Kong Institute for Monetary and Financial Research	Eric Lam	Research Analyst
11	PricewaterhouseCoopers	Anand Rao	Editorial Panel Lead
12	PricewaterhouseCoopers	Brian Yiu	Editorial Panel
13	PricewaterhouseCoopers	Jeff Hao	Editorial Panel
14	PricewaterhouseCoopers	Karen Wong	Editorial Panel
15	PricewaterhouseCoopers	Akihiko Katayama	Co-author
16	PricewaterhouseCoopers	Dick Fong	Co-author

No	Organisation	Name	Role
17	PricewaterhouseCoopers	Alan Chan	Research Lead
18	PricewaterhouseCoopers	Chester Luo	Research Analyst
19	PricewaterhouseCoopers	Circle Ku	Research Analyst
20	PricewaterhouseCoopers	Ivana Sun	Research Analyst
21	PricewaterhouseCoopers	Sabrina Yang	Research Analyst
22	PricewaterhouseCoopers	Sandra Huang	Research Analyst
23	PricewaterhouseCoopers	Monica Uttam	Editor
24	PricewaterhouseCoopers	Peter Craughwell	Editor
25	PricewaterhouseCoopers	Sanjukta Mukherjee	Editor
26	Cygnet Communications Limited	Dr. Simon Alderson	Editorial Advisor

AI Survey Respondents

This white paper research has benefitted significantly from the input provided by 168 Authorized Institutions which submitted their responses to the Banking AI Survey. The Chinese University of Hong Kong has contributed to the design of the survey questionnaire. The information collected is valuable in understanding the existing AI landscape in Hong Kong and the adoption of AI by banks operating in the city

Individual participants in feedback sessions

The content has also been shaped by the following individuals, representing a balanced portfolio of financial institutions, startups and FinTech facilitators. The insights shared by these individuals have proved a valuable sounding board for our understanding of existing AI use cases, banks' challenges in adopting AI technologies, and ways in which Hong Kong can improve its promotion of and innovation in AI.

No	Organisation	Name	Role
1	Bank of China (Hong Kong)	Eddie Wong	Senior Project Manager, Data Application Development Division, Information Technology Department
2	Bank of China (Hong Kong)	Frazer Lam	Deputy General Manager (IT Development), Information Technology Department
3	Bank of China (Hong Kong)	Joel Tse	Senior Project Manager, Risk Management & MIS Division, Information Technology Department
4	Bank of East Asia	Howard Lau	Head of Information Technology Department
5	Bank of East Asia	Sean Lee	Section Head & Application and Data Architect, Information Technology Department
6	BNP Paribas	Vivien Jong	Managing Director, Asia Chief Digital Officer, BNP Paribas Wealth Management
7	Citibank	Banny Wong	Director, Digital Channels and Experience
8	Citibank	LeAnne Wong	Head of Data and Analytics
9	Cyberport	Ivy Chan	Head of Ecosystem and Collaboration
10	DBS Bank (Hong Kong) Limited	Alexander Lau	Digital Innovation Head, Institutional Banking Group
11	DBS Bank (Hong Kong) Limited	Angela Wong	Group Portfolio Analytics, Risk Management Group
12	DBS Bank (Hong Kong) Limited	Gaurav Aggarwal	Head of Data and Analytics
13	DBS Bank (Hong Kong) Limited	Matthew Ng	Innovation & Ecosystems
14	DBS Bank (Hong Kong) Limited	Medhy Souidi	Head of Fintech & Startup Xchange, Innovation & Ecosystems
15	DBS Bank (Hong Kong) Limited	Vincy Shiu	Credit COO Office, Risk Management Group
16	Element AI	Hardeep Arora	Delivery Lead, APAC

No	Organisation	Name	Role
17	Element AI	Marie-Claude Ferland	Advisory Lead, APAC
18	Element AI	Victoria Yung	Commercial Lead, Financial Services APAC
19	Fano Labs	Miles Wen	Chief Executive Officer, Honorary Assistant Professor, University of Hong Kong
20	Hong Kong Applied Science and Technology Research Institute (ASTRI)	Hugh Chow	Chief Executive Officer
21	Hong Kong Applied Science and Technology Research Institute (ASTRI)	Terry C. K. Lee	Head of Public Affairs
22	Hong Kong Science and Technology Parks Corporation	Cathy Lo	Senior Manager, Strategic Planning
23	Hong Kong Science and Technology Parks Corporation	George Tee	Chief Technology Officer
24	Hong Kong Science and Technology Parks Corporation	Kent Chen	Senior Manager, ICT Cluster
25	JPMorgan Chase	Andrea De Zordo	Managing Director, Global Equities & Prime Services
26	JPMorgan Chase	Johnny Chow	Executive Director
27	JPMorgan Chase	Samik Chandarana	Head of Analytics and Data Science
28	OneConnect - Ping An OneConnect Bank	Jeffrey Ng	Head of FinTech Solutions Department
29	OneConnect - Ping An Tech	Eugene Huang	Deputy General Manager, CTO & COO
30	The Hongkong and Shanghai Banking Corporation	David Knott	Global Head of Architecture and Innovation
31	WeBank	Alan Ko	Head of FinTech Innovation (Int'l), Technology
32	WeBank	Kevin Lee	Senior FinTech Innovation Manager (Int'l)
33	WeBank	Toby Chen	Deputy General Manager, A.I. Group