



Artificial Intelligence

Opportunities, Risks
and Implications



Global Mindset, Asian Insights

Contents

Introduction	1
How AI is used today?	2
The implications and risks of AI use	4
Responses to AI: to regulate or not?	8
Conclusion	15





Introduction

Artificial Intelligence, or AI, is making huge strides today, around the world and in Singapore.

The rapid developments in AI have engendered high expectations of the benefits it can deliver for business and for society at large. Yet there are concerns about AI replacing human input and thus displacing jobs.

In 2017, Elon Musk, the CEO of Tesla, sparked a debate when he said that AI poses a “fundamental risk to the existence of civilisation,” and called for regulators to be proactive rather than reactive to developments in AI.¹

On the other hand, others like Oren Etzioni, CEO of the Allen Institute for Artificial Intelligence, warned Americans that if they took steps to slow down progress on AI, other countries like China would overtake them.²

In this publication, we examine the opportunities, risks and implications of AI use, particularly in the context of the accounting and finance industry, and how we can equip ourselves to deal with such future scenarios.

What exactly is AI?

AI is a rather broad term that has been bandied about in the media. At some level, data analytics and AI have considerable overlaps. One of the goals of data

analytics is to interrogate data to obtain insights in order to make predictions, as in predictive analytics. Machine learning can be seen as an evolution of predictive analytics, albeit with vastly more advanced state of the art techniques and knowledge such as the use of neural networks.

Most of the AI application we hear of today is **machine learning**, where predictions are made by a computer system based on its exposure to data. In other words, the computer system automatically spots patterns from data sets; then, depending on the choice of algorithm, it will present a possible interpretation for a decision to be made. An algorithm is a set of rules used for the analysis.

Then there is **deep learning** which, as its name suggests, presents a deeper and higher level of machine learning. It imitates how the human brain works, in creating new patterns of data for decision making. It can “learn” without supervision, from unstructured data – data that is not already neatly organised according to pre-determined categories.

Robotic process automation (RPA) is a form of business process automation, which some say is not AI, strictly speaking. For example, RPA can be used by a financial institution to automate the processing of credit card application.

1 The Guardian, “Elon Musk: regulate AI to combat ‘existential threat’ before it’s too late,” 17 July 2017: <https://www.theguardian.com/technology/2017/jul/17/elon-musk-regulation-ai-combat-existential-threat-tesla-spacex-ceo>

2 Oren Etzioni, “How to Regulate Artificial Intelligence,” The New York Times, 1 September 2017: <https://www.nytimes.com/2017/09/01/opinion/artificial-intelligence-regulations-rules.html>

One way to explain the nuanced distinction is that RPA mimics human behavior and actions, while AI mimics human thinking.

For today and the foreseeable future, as David Leow, Managing Director of Thaler Global, explains, the AI we have is termed “weak AI” and would fail the Turing test for strong AI, as their capabilities are easy to distinguish from those of humans. For example, current versions of Siri, Alexa, and Google Assistant have impressive voice recognition and are quite adept at scheduling appointments, but a real human assistant can think ahead to resolve logistical issues such as booking transport, and make and distribute concise meeting notes.

Nevertheless, advances are being made fast enough that it may not be long before a stack of receipts and

bank statements could be machine translated into a set of financial statements and tax returns ready for filing. As the goal of strong form AI is brought closer to reality, we will be able to take increasingly raw inputs to get increasingly more useful outputs.

The Turing test

A test, devised by the English computer scientist Alan Turing, which involves blind judging as to whether the responses coming from a human and a computer can be distinguished. If it cannot, it then means the computer has attained a standard of “strong AI”.³

1

How AI is used today

Scenarios of robots taking over the world may seem to be the stuff of science fiction in the 1980s and 90s. Today, humans already co-exist with robots in factories and offices, in hotels and restaurants, in schools and many other places. In this section, we look at some current uses of AI, particularly in the finance industry, to see how AI is already impacting businesses.

ATMs

Chng Sok Hui, Chief Financial Officer of DBS Bank, says, “DBS has been an early user of AI, and the use of AI permeates the entire organisation.

“A little known fact is that our ATMs in Singapore are the highest velocity ATMs in the world. We do about 20,000 transactions a month.” In comparison, an average ATM in the US does 2,000 transactions a month.

This means the ATMs get emptied of cash quickly, and DBS was getting significant customer feedback on this issue. Back in 2010, some of DBS’s ATMs needed to be topped up multiple times a day.

To fix the problem, DBS used AI. It was one of DBS’s early experiments in using analytics and machine learning to predict ahead of time when their ATMs were going to run out of cash, and then plan their delivery strategy accordingly.

The outcome now is that the average ATM rarely runs out of cash. And by reducing the number of trips to top up cash at ATMs, DBS was able to reduce costs significantly.

Audit

“We used to audit our branches on a sample basis,” Chng says. “Then we asked: can we use analytics to identify the higher risk branches?”

DBS looked at leave patterns of staff, their resignations, as well as transaction data and behavioural patterns in each bank branch. They then applied a predictor of three different types of models.

“We got a significant lift in predicting which branches will have higher risk issues,” Chng says.

3 Ergun Ekici, “Passing the Turing Test: Redefining What It Means to ‘Think’,” Wired: <https://www.wired.com/insights/2014/08/passing-turing-test-redefining-means-think/>

Finance

DBS now has a daily balance sheet for the entire DBS Group by granular cuts, with the ability to drill down to transaction level. They have also implemented the daily net interest margin by key locations.

“We are now developing predictive tools for flash estimates of balance sheet, income and expenses as well as key financial metrics,” Chng says. “These tools are dynamic and provide powerful insights on steering our balance sheet and managing our businesses.”

Trading securities⁴

Turing Intelligence Technology (TurinTech), a UK-based AI company, provides a trading platform which automates the buying and selling of securities. The platform is based on conditions and logic rules, signal, trade, and event-driven trade execution, transparent trades, auto custom exchange trade, and trading that relies on complex machine learning models.

Besides up-to-date equity pricing, corporate facts and figures, and futures, TurinTech provides access to many other data streams, such as topics trending in social media, or the current popular search terms on Google. It also helps users extract and put to work keyword-associated text, from web pages, PDF files or other documents. All of this data is collected using the machine learning process.

What TurinTech does has opened advanced trading – traditionally the preserve of large investment banks – to even personal investors.

Securing alternative funding⁵

The Fundworks, a US-based financing provider, helps

small businesses secure alternative financing through merchant cash advances.

It uses Microsoft’s Azure Machine Learning solution to develop predictive models that identify high-performing brokers, merchants and deals. By harnessing powerful technology to track, analyse, and process thousands of transactions every day, The Fundworks is able to maximise efficiency and provide businesses with speedy access to capital.

Other AI case uses

In the legal services sector, Kira is a machine-learning software that is used to speed up the identification and analysis of clauses and other data. If a lawyer were to pull out ten data points on a contract, Kira can cross-reference thousands of contracts to these data points very quickly.

In the consumer healthcare sector, chatbots are already being used to sell cold and flu products. When customers visit the website of a consumer healthcare company, they are able to interact with a chatbot that recommends products based on the symptoms they share with the chatbot – arguably in a similar manner to the way in which they would consult a pharmacist at the drugstore.

To sell advertising space to clients online, bots are used by agencies to determine their clients’ target audiences, such as in terms of age or location. The bot would then suggest media where the agencies’ clients could buy space.

Robo-advisory services are used by financial institutions on their websites to advise their clients on the types of financial services and products that best suit their needs.



4 Microsoft, “Investment AI firm uses machine learning and Microsoft startup support to drive big business at low cost,” 29 March 2018: <https://customers.microsoft.com/en-us/story/turintech-banking-and-capital-markets-azure>

5 Microsoft, “Cash cloud: financing company uses cloud-based portal to offer fast, smart capital to support small businesses,” 26 January 2018: <https://customers.microsoft.com/en-us/story/the-fundworks-banking-azure>

2

The implications and risks of AI use

“Compared to human intelligence, AI is rather stupid [if one considers the machine learning process]. But stupid as AI is, look at what it has already achieved. We should be alarmed AI can do so much.”

– Enoch Ch'ng,
Former Associate Professor of Information Systems
(Education), Singapore Management University,
and Senior Advisor, Data Science Rex

While the promise of what AI can deliver is exhilarating, it is not without implications and risks.

“There is potential for AI systems to go rogue,” Richard Koh, Chief Technology Officer at Microsoft Singapore, says. “Its owners may not even know with certainty these systems’ algorithms’ decision-making or recommendation processes. The development, deployment and on-going nurture of AI thus comes with a lot of responsibility.”

Here, we present some practical scenarios of the implications and risks of the use of AI, with responses and insights from our interviewees on how they can be addressed. While some of these are not ethical issues, and neither are they unique to AI applications, it may be worth considering regulatory or mitigating measures to avoid these undesirable scenarios.

The problem of the reinforcing feedback loop

Given the nature of the machine learning process, the performance of an AI system is dependent on the data that is fed into it. This could lead to skewed results, because of the “follow the crowd” mentality that AI algorithms tend to engender.

An over-reliance on AI could lead to the danger of the self-reinforcing phenomenon. The next example on AI-powered stock exchanges illustrates these dangers.

AI-powered trading

Enoch Ch'ng, former Associate Professor of Information Systems (Education) at Singapore

Management University, and Senior Advisor at Data Science Rex, presents a scenario where all trading systems in the world were to be run on AI. If one AI system detects a signal to sell in a market or on an exchange, similar systems would detect the same signal and trigger a sell order as well. In turn, this would reinforce and amplify the impact, spiralling into a market crash.

One solution to pre-empt this problem is generally to intervene and calm the market. However, if the “sell” signal is triggered by self-learning AI systems, we may not know how or what to do to “calm” these machines, because we do not know what they are thinking or how they work.

The alternative may be to depend on the “circuit-breaker” as has already been instituted to halt high frequency, algorithmic trading systems when needed. But the use of such a “circuit-breaker” could have its downsides as well – it could equally cause another set of issues, resulting in operational chaos.

Definitions

High frequency, algorithmic trading: A high-speed, high volume computer system where formulae are used to automatically determine the parameters of instructions in the trading process, such as whether to initiate the order, of transactions and/or price.

“Circuit-breaker”: A mechanism to shut down or disable a system, such as in an emergency.

The 2010 “Flash Crash”

On 6 May 2010, stock indices in the United States, including the S&P 500, Dow Jones Industrial Average and Nasdaq Composite, collapsed and rebounded very rapidly within approximately 30 minutes. During that time, hundreds of billions of dollars were wiped off share prices, and then largely restored.

The report of the US Securities and Exchange Commission (SEC) on the “flash

crash” of 6 May 2010 highlighted the role of high-frequency trading firms (HFTs), which operate on automated algorithms and typically buy and re-sell shares in less than a second.

Since then, the SEC has introduced “circuit-breakers” for individual shares that stop trading across all markets, which could be modified to allow shares to continue trading within pre-set bands.⁶

China’s Social Credit System

China intends to launch a Social Credit System (社会信用体系) by 2020, which will assess and score the economic and social reputation of its citizens on a national scale.⁷

This would work very much like how rating agencies assess credit worthiness before an individual is allowed to take on a loan – except that the implications here are far deeper. Using data such as traffic offences, court records, and possibly the social media activities of citizens, an AI-powered system would determine an individual’s integrity. This would have the potential of disadvantaging a job candidate, jeopardising an individual’s chances of gaining entry to university, or even barring an individual from buying flight tickets to travel abroad.

Aside from debates about whether such a mass surveillance programme would infringe on human rights, there are some very practical concerns too.

Imagine the scenario where the data collected for the purpose of assessing your “social score” is being manipulated as a result of a cyber-attack. The conclusions derived are highly dependent on the quality of data set. As an individual, you would not know what has been programmed into the system, what data is being used, and you would be in no position to lodge an appeal on the veracity of the personal data collected. Based on what the Social Credit System is meant to do as cited above, you could then be discriminated against, through no fault of yours.

“This is not a new dilemma,” says Chng Sok Hui, who cites what banks and personal credit bureaus have already been doing to determine the credit worthiness

and risk appetite of customers. “The uneasiness that people have with the pervasiveness of AI and the use of personal data might stem from the fact that tech firms are not being held to higher fiduciary standards, in the same way commercial banks are today.

“Individuals can protect their privacy and reputation against exploitation by increasing their own awareness and education on potential uses of data. But personal choices on the use of data needs to be better guaranteed via institutional safeguards, such as through mandatory opt-out mechanisms, or placing higher fiduciary standards on large tech firms, or through a more widespread adoption of a code of conduct that guides the use of personal data.”

Driverless car accidents

One scenario of AI use that has particularly worried the general public is that of driverless cars getting into or causing accidents.

Tragically, that problem became reality in the US in March 2018 with the first fatal accident involving an autonomous vehicle.⁸ Investigations concluded in part that the self-driving car’s sensors struggled to identify the victim, who was wheeling a bicycle across the road – it could not determine whether it was a person or a vehicle, because of the wheels of the bicycle. Moreover, the self-driving car’s built-in emergency braking system had been disabled, to prevent conflict with the self-driving system.



6 The Economist, “One big, bad trade,” 1 October 2010: <https://www.economist.com/newsbook/2010/10/01/one-big-bad-trade>

7 See for instance: Financial Times, “Inside China’s surveillance state,” 19 July 2018: <https://www.ft.com/content/2182eebe-8a17-11e8-bf9e-8771d5404543>

8 The Economist, “Why Uber’s self-driving car killed a pedestrian,” 29 May 2018: <https://www.economist.com/the-economist-explains/2018/05/29/why-ubers-self-driving-car-killed-a-pedestrian>

Who is at fault in such a scenario?

For Matt Pollins, Partner and Head of Technology, Media and Telecoms at CMS Singapore, this issue will depend on many factors, including the contractual allocation of responsibilities between the parties, road traffic regulations, whether a human operator is present in the vehicle and which party has assumed the relevant duty of care. In relation to a single incident, it is possible that more than one party could have a degree of responsibility – from the developers of the underlying AI software and manufacturers of the relevant components to the car manufacturer and the driver. Each case will require a fact-based assessment of matters such as causation and foreseeability, which will rarely be straightforward in the context of AI because so many parties are involved in its development.

This is demonstrated by the case of the March 2018 accident in the US, where a human safety operator was present in the vehicle, which was an experimental driverless car. The human safety operator was expected to perform an emergency brake, but did not. This demonstrates how many different factors can complicate the legal question of liability.

Even less straightforward, perhaps, would be the question of liability in the scenario of a driverless vehicle crashing into another driverless vehicle. Who, or what, is at fault?

When posed this question, Richard Koh said, “At Microsoft, one of the key foundation principle of designing AI to earn trust is Accountability. People who design and deploy AI systems must be accountable for how their system operate.”

Of greater critical importance perhaps is the cyber security risk for autonomous, connected cars. Here, a breach of a car’s network does not only compromise the privacy of a driver’s data, but it can be a matter of life and death – if the autonomous vehicle is remotely hijacked for a malevolent intent. A McKinsey report highlighted this as a threat to the car industry’s road map towards autonomous vehicles.⁹

Racially-biased sentencing

One of the key ethical risks associated with the use of AI is the potential for discrimination, such as on the basis of ethnicity.

Here, the example is often cited of the use by a US court of COMPAS, an AI-powered risk assessment computer programme, in deciding the length of jail sentence for a convicted criminal on the basis of his likelihood of re-offending.¹⁰ The convict then appealed, on the basis that the software’s algorithm was a proprietary secret, which made him unable to inspect or challenge its sentencing recommendation.¹¹

Separately, a study of the COMPAS software found that it was far more likely to judge black defendants to be at a higher rate of re-offending than whites.¹²

While the AI programme may initially seem to be more objective than a human judge, the dangers of racial and other discrimination through widespread AI use is real.

“This case highlighted the fact that while AI can demonstrate certain features of intelligence, it can only reflect the data that goes into it, as well as any prejudices contained within that data and of the developers of the relevant application,” says Matt Pollins.

“It is therefore key to ensure that the data fed into such AI programmes is diverse, and not skewed towards or against any particular group, whether in terms of race, or demographics, or otherwise. Fairness, transparency and explainability are all important factors in the trusted use of AI.”

Job disruption

At a more general level, the concern people may have with the pervasive use of AI is that they may be displaced in their jobs.

In a situation where AI and other technologies were to take over as much as 50% of jobs, as a McKinsey study found to be already possible,¹³ this would result in high levels of unemployment and, consequently, social unrest.

9 McKinsey & Company, “Shifting gears in cyber security for connected cars,” February 2017: <https://www.mckinsey.com/~media/mckinsey/industries/automotive%20and%20assembly/our%20insights/shifting%20gears%20in%20cybersecurity%20for%20connected%20cars/shifting-gears-in-cyber-security-for-connected-cars.ashx>

10 Harvard Law Review, “State v. Loomis: Wisconsin Supreme Court Requires Warning Before Use of Algorithmic Risk Assessments in Sentencing,” 10 March 2017: <https://harvardlawreview.org/2017/03/state-v-loomis/>

11 Note though that the appeal failed, on the basis that a human judge would have considered the same criminal records of the convict as the COMPAS software.

12 ProPublica, “Machine Bias: There’s software used across the country to predict future criminals. And it’s biased against blacks,” 23 May 2016: <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

13 McKinsey Global Institute, “Jobs Lost, Jobs Gained: Workforce Transitions in a Time Of Automation,” December 2017: <https://www.mckinsey.com/~media/McKinsey/Featured%20Insights/Future%20of%20Organizations/What%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/MGI-Jobs-Lost-Jobs-Gained-Report-December-6-2017.ashx>



Enoch Ch'ng underscores that such a scenario of high unemployment would have significant social and economic consequence – for instance, resulting in a lack of inflows into the financial market. After all, financial institutions offering products such as loans would still require customers who are drawing an income to be able to make loan repayments, so that their whole business model would work. Otherwise, the result would be financial instability, or a financial disequilibrium on a global scale.

Going by this argument from a macroeconomic perspective, it would not be in the interest of financial institutions to see all the world's jobs automated.

This is not the first time that a new technology has been predicted to result in job losses. It remains to be seen if such alarming rates of job losses will materialise. In any case, Singapore is looking into how it can prepare for an AI future.

Yeong Zee Kin, Assistant Chief Executive (Data Innovation and Protection Group) of the Infocomm Media Development Authority of Singapore, and Deputy Commissioner, Personal Data Protection Commission (PDPC), says, "For Singapore to adopt AI, it requires a whole range of skill sets across various roles in an organisation. For instance, we have been working with our TechSkills Accelerator team, looking into the skills necessary for an AI economy."

Starting from leadership positions, it would be vital to bring about the awareness of AI to the C-suite leadership of companies. They need to be aware that such technology already exists, what the positive examples of these technologies are, and how such technology can affect their business.

For those in managerial functions, there is a need for greater awareness to see how AI could replace certain technology currently in place, and/or consider what functionalities in the current operations in one's business can be enhanced by AI.

For the other employees within each organisation, there needs to be an understanding of AI on how AI affects the way jobs are carried out. Just as how search engines have changed the way documents are retrieved a generation ago, AI will similarly require employees to pick up new skill sets, and undergo training on supervising AI processes that are in place. Recently, IMDA and AI Singapore announced the AI for Industry programme to help working professionals better understand and use AI appropriately, with hands-on development in basic AI and data applications.

All in all, Yeong highlights the need for different training at different levels within an organisation.

3 Responses to AI: to regulate or not?



“CEOs and their Boards of Directors tend to be growth oriented and more focused on opportunities, threats and execution than ethical issues at this time,” says David Leow. “The most foresighted ones are also crafting policies around the use of AI and customer data that are consistent with regulations, community standards and their own corporate culture.”

When considering our responses to the potential and risks of AI, the first step is to demystify AI and the process of how it works. For Enoch Ch’ng, much of the risks associated with AI stems from stakeholders’ lack of understanding of how an AI system is designed, constructed, tested, used and managed. This includes the algorithm, the model used and the data used. If the understanding of stakeholders is limited, it could be due to a lack of transparency or access to information, or simply a lack of awareness.

Another way of looking at it is that transparency of what goes on within AI systems can help us achieve a higher understanding, develop mitigation measures and take appropriate actions. This will go a long way in enabling the use of AI systems to be more accountable to the businesses and the societies they are meant to serve.

In this section, we consider the responses of government/regulators, the private sector – financial institutions as well as tech firms, the legal fraternity,

and from thought leaders on the ethics of AI use. In doing so, this publication aims to catalyse dialogue between these sectors, as part of a longer-term mission of building a framework for the governance of AI.

Governments and regulators

“Technology is neutral,” says Yeong Zee Kin, who underscores that the emergence of new technologies does not automatically render the need for regulation. Rather, regulation comes into play when there are certain objectives to be achieved, such as the need for safety or technical standards.

In Singapore, there is no intention to rush into the regulation of AI, which can already be viewed as an extension of the general data protection regime. Rather, Yeong highlights that it is more vital to start the conversation about the governance framework surrounding AI.

Towards this end, PDPC has been involved in establishing a research programme on ethical, legal, regulatory and policy issues arising from commercial deployment of AI. This programme aims to provide a forum and thinking to build up a body of knowledge and a pool of experts on these issues. Without such a forum of knowledge in place, Yeong highlighted that Singapore would still be in a premature state to talk about regulating AI.

It was announced on 5 June 2018 that the Singapore Management University (SMU) School of Law was awarded a grant of S\$4.5 million from the National Research Foundation (NRF) to helm a five-year Research Programme on the Governance of AI and Data Use.

Advisory Council on the Ethical Use of AI and Data

On 5 June 2018, it was announced an Advisory Council on the Ethical Use of AI and Data will work with IMDA on the responsible development and deployment of AI.¹⁴ The advisory council would assist IMDA in engaging relevant stakeholders such as the ethics boards of commercial enterprises, as well as the private capital community, to raise awareness of the need to incorporate ethics considerations in their investment decisions into businesses which are involved in AI.

With the development of such a governance framework, PDPC also hopes to encourage good practices in AI use among companies, enabling companies to give an account of what internal controls were in place to mitigate risks, how decisions were made, and what feedback channels consumers

have. This will result in an overall boost of trust in adopting AI solutions for economic benefit.

“At present, the law and accounting industries are ready to reap benefits from the adoption of AI,” Yeong Zee Kin says. “For accountants, the use of AI can reduce the time needed to reconcile statements and identify anomalies or detection of fraud, enhancing the efficiency and effectiveness of accountants. For lawyers, AI-assisted document review means the reduction in time needed to review voluminous documents.”

PDPC’s Discussion Paper: Fostering Responsible Development and Adoption of AI

The Personal Data Protection Commission (PDPC) has released a discussion paper on how a possible reference AI and data governance framework for industries could look like.¹⁵

The discussion paper is intended for organisations to use for their internal discussion and adoption. Furthermore, PDPC encourages trade associations and chambers and professional bodies to adapt the proposed framework for their respective sectors, for instance as a voluntary code of practice.



14 Infocomm Media Development Authority (IMDA), “Artificial Intelligence Governance And Ethics Initiatives,” Factsheet, 5 June 2018: https://www.gov.sg/~sgpcmedia/media_releases/imda/press_release/P-20180605-1/attachment/Artificial%20Intelligence%20Governance%20and%20Ethics%20Initiatives.pdf

15 For the full discussion paper, please see: <https://www.pdpc.gov.sg/Resources/Discussion-Paper-on-AI-and-Personal-Data>

The paper proposes a four-stage governance framework:

- A. **Identifying the objectives of an AI governance framework** - Organisations that use AI in their decision-making process should set objectives such as for them to be able to explain how their AI systems work, and to put in place good data accountability practices.
- B. **Selecting appropriate organisational governance measures** - This stage considers the good practices such as setting up internal governance policies and structures and establishing operations management controls that can help organisations be accountable to regulators, affected individuals and/or stakeholders.
- C. **Considering consumer relationship management processes** - To build consumer trust and confidence, the management of communications with affected individuals and providing measures for recourse are important.
- D. **Building a decision making and risk assessment framework** - Decision-making and risk assessment considerations are incorporated into the framework at this stage. The risk and severity of harm to the customer are some of the factors that affect which decision-making approach should be adopted, and how organisations calibrate their governance and consumer management processes.

For financial institutions: The Fairness, Ethics, Accountability and Transparency (FEAT) Committee

On 2 April 2018, the Monetary Authority of Singapore (MAS) announced that it would be working with stakeholders in industry to develop a guide to promote the responsible and ethical use of AI and data analytics by financial institutions.

Slated for completion by the end of 2018, the guide would cover all segments of the financial sector including fintech firms. The guide will set out key principles and best practices for the use of AI and data analytics, helping financial institutions to strengthen internal governance and reduce risks of data misuse.

To co-create the guide, MAS convened the Fairness, Ethics, Accountability and Transparency (FEAT) Committee, bringing together leaders and practitioners in data analytics in the financial sector.¹⁶

David Hardoon, Chief Data Officer of MAS and co-chair of the FEAT committee, said in November 2017 in an interview with The Business Times, "The nascent use of AI and data analytics in the financial sector gives us an opportunity to move ahead of the curve and set the right conditions for its ethical and proper use. We look forward to working with the industry and the public to understand the opportunities and challenges ahead."¹⁷

Responses from the private sector

Banks and financial institutions

Putting AI into perspective

For Chng Sok Hui, the benefits of AI use are both incremental and disruptive.

"To understand the real potential of AI, we must evaluate the impact of AI together with the whole emerging ecosystem - digitisation, robotic technology, automation, big data, machine learning, the Internet of Things (IoT), among many other emerging technologies," Chng says.

"There will be many incremental benefits from gaining efficiency across things we are doing today - reduce the process time, minimise costly errors, improve the return on marketing and sales investment, use data to create better product features, improve the current processes which are of a cookie-cutter nature to be more contextually aware and personalised."

But there are also opportunities from more disruptive initiatives, especially when decisions can be made and executed automatically throughout the digitised process chain in real time. "It will fundamentally change how we evaluate the competitiveness of the company and the worth of the individual," Chng adds.

"In the early stage, when we are experimenting large scale of deployment of 'narrow AI', there could be systematic risk or biases perpetuated from the design of the system. In the later stage, when artificial general intelligence (AGI) is being deployed, the risk will come from uncertainties of how its consciousness will evolve in real world.

16 Monetary Authority of Singapore, "MAS and financial industry to develop guidance on responsible use of data analytics," 2 April 2018: <http://www.mas.gov.sg/News-and-Publications/Media-Releases/2018/MAS-and-financial-industry-to-develop-guidance-on-responsible-use-of-data-analytics.aspx>

17 Business Times, "MAS steers debate on ethics of AI, Big Data; kicks off industry consult," 28 November 2017: <https://www.businesstimes.com.sg/banking-finance/mas-steers-debate-on-ethics-of-ai-big-data-kicks-off-industry-consult>

Definitions¹⁸

Narrow AI: The most widespread form of AI use currently, in which the focus is on a single, particular task such as facial recognition.

Artificial general intelligence (AGI): A futuristic, as-yet-unachieved application of AI which can understand and respond to its environment in multiple ways, just like a human would. While humans may not be able to process data as fast as computers, they can think abstractly and plan its responses, or to invent a product that does not yet exist – in a way computers currently cannot.

Responding to AI

“As the development of AI is a complex process, it is difficult to anticipate any unintended consequences in the design, application and review of AI,” Chng Sok Hui says.

“A priori attempts made to curtail unintended consequences during the design, application and review stages are likely to be counter-productive, depriving society of the more beneficial fruits of AI.

“Top-down oversight is also impractical and not optimal, as the onus ultimately falls on experts in the field, who are required to make judgments on what is good and bad AI for society.”

Rather than to pre-empt the unintended, negative consequences of AI in the design, application and review stages of AI, Chng believes a more sensible goal for society would be to build a system of safeguards. Such a system should seek to ensure that society's values and way of life is resilient and robust to unintended consequences from AI's development.

As an example of measures that could be taken, Chng cites institutional safeguards in the form of tougher penalties for anti-competitive behaviour by large tech firms that would actively prevent concentration of AI algorithms in the hands of a few in the industry. She believes this can be more pro-actively championed by business and community leaders.

Three main themes for conversation

For Chng Sok Hui, there is a need for greater conversation in the public domain around three main themes, as AI develops.

The first relates to the responsibilities of tech firms that employ AI, and the individual designers of AI algorithms – it is not clear who bears ethical responsibilities behind decisions made in AI algorithms.

The second concerns how the value of data can be better distributed between providers and consumers.

The third is about having a more level playing field between incumbent large tech firms – which tend to gravitate towards monopolies naturally – and new entrants in AI.

As to how AI could be more accountable and their algorithm applications more transparent, Chng is of the opinion that market-based solutions are unlikely to achieve this. After all, tech firms which rely on AI for competitive advantage in contemporary capitalist societies gravitate towards natural monopolies due to network effects.

“AI can be made more accountable by setting higher fiduciary standards for tech firms and their key employees, firming and broadening regulatory and governance structures that prevent anti-competitive behaviour by large tech firms and investing persistent time and effort in levelling the playing field between incumbents and new entrants,” Chng says.

Tech companies – Microsoft's AETHER team

Winning the trust of consumers has become ever more important for technology companies, especially given developments such as with the debacle involving Facebook and Cambridge Analytica in early 2018, in which the personal information of social media users was improperly obtained for the purpose of influencing electoral outcomes.

“The development, deployment and on-going nurturing of AI thus comes with a lot of responsibility,” says Richard Koh. “This is why Microsoft has formed a cross-company team of executives and experts called AETHER – AI and Ethics in Engineering and Research – in 2017 to discuss and recommend programmes, policies, procedures, and best practices on issues of AI safety, fairness, transparency, ethics, and questions and challenges arising more broadly at the intersection of AI, people, and society as AI moves into the open world.”

¹⁸ Ben Dickson, “What is Narrow, General and Super Artificial Intelligence,” TechTalks, 12 May 2017: <https://bdtechtalks.com/2017/05/12/what-is-narrow-general-and-super-artificial-intelligence/>

Earlier in 2018, Microsoft published *The Future Computed: Artificial Intelligence and its role in society*¹⁹ – its first book to highlight key ethical, legal and moral issues that must be considered when developing and using AI in our daily lives across society, and most importantly, designing AI to earn trust.

“Making AI technologies widely available and using them alone is not sufficient,” says Koh. “For AI to be adopted on a wider scale, we need to develop AI in a way that people deem trustworthy. An important reason why humans operate in the physical world with confidence and trust is because there is broad consensus about what kinds of behaviour are considered responsible and acceptable, norms that are implicitly understood, consistently applied in everyday social interactions and reinforced in laws. All relevant stakeholders, from business, governments, civil society and research communities must collaborate on shared principles and ethical frameworks.”

The journey and process at Microsoft through which these principles emerged is honed through decades of best practices in software engineering. “When I started my Microsoft career in the product teams in 2006, Trustworthy Computing and the Secure Development Lifecycle had already become the norm for Microsoft,” Koh says.

“Trust has been at the centre of this massive effort. With AI now possessing the abilities to perceive and comprehend, trust remains at the centre of AI being developed and used to serve mankind and all the societal challenges it is facing today. AI should not be misused by a privileged few to pursue unfettered and pervasive privacy-eroding uses and enriching their own personal fortunes.”

Responses from the legal fraternity

Are new, AI-specific laws needed to protect us from privacy intrusions stemming from the use of AI? The response from Matt Pollins was a no.

“AI doesn’t operate in a legal vacuum. We already have laws and regulations that govern matters such as privacy and cybersecurity in the context of AI,” Pollins says.

“In Singapore, the Personal Data Protection Act (PDPA) already governs the collection, use and disclosure of personal data by organisations.” Pollins also highlights

Definitions

Trustworthy Computing is a concept associated with Microsoft, focusing on the key areas of Security, Privacy, Reliability, and Business Integrity, through which security was established as a top priority for the tech industry rather than as an afterthought. Microsoft’s Trustworthy Computing initiative was launched in 2002 by its founder, Bill Gates, in response to massive malware attacks then like Code Red, Nimda, and SQL Slammer which crippled internet networks. Within Microsoft more recently, the Trustworthy Computing team is part of its Cloud and Enterprise Division, and maintains company-wide responsibility for centrally driven programmes such as the Security Development Lifecycle (SDL).²⁰

Microsoft’s Security Development Lifecycle (SDL) is a software development process that helps developers build more secure software and address security compliance requirements while reducing development cost.²¹

the recently implemented European Union (EU) General Data Protection Regulation (GDPR), designed to protect EU citizens from privacy and data breaches. (This law, while passed in the EU, also affects companies based in Singapore that operate in the EU, target EU markets or provide services to companies based in the EU.)

Pollins further points out that the same can be said of laws in other areas, such as competition laws, which would apply, for example, in circumstances where price-setting algorithms of two or more organisations are found to be “colluding” to artificially inflate prices.

“I’d argue that what we need instead is perhaps some guidance, and in appropriate circumstances regulation, to clarify how existing laws apply in the context of AI. It’s more about how our existing laws evolve to address AI, rather than introducing a whole new set of AI-specific laws, which I fear would have the potential to stifle innovation.”

19 Microsoft, “The Future Computed: Artificial Intelligence and its role in society”: <https://news.microsoft.com/futurecomputed/>

20 Scott Charney, “Looking Forward: Trustworthy Computing,” Microsoft Secure, 22 September 2014: <https://cloudblogs.microsoft.com/microsoftsecure/2014/09/22/looking-forward-trustworthy-computing/>

21 Microsoft, “What is the Security Development Lifecycle?” <https://www.microsoft.com/en-us/sdl>



Ethical considerations from the engineering community

On the debate as to whether AI and other such technology should be regulated, the response of Danit Gal, chair of the Outreach Committee of The Institute of Electrical and Electronics Engineers (IEEE) Global AI Ethics Initiative, and Project Assistant Professor at Keio University, Japan, is an emphatic yes.

“If innovation puts society at risk, then governments should work to find a solution and prioritise the mental and physical safety of people over technology,” says Gal. “That’s not over-regulation, it’s prioritising the safety of your citizens over the rapid progress of technology – which should always be the case.”

Nevertheless, she believes that knowledge sharing between the government and the private sector and among ministries “is the best shot we have at effectively regulating AI without losing our competitive edge or stifling innovation.”

With regard to how we should regulate AI, Gal believes governments should understand two things. “They will need to consult experts that understand the technology and its risks on an ongoing basis to ensure they regulate effectively,” she says. “Secondly, AI, just like other technologies upon which it is built, is relevant across the ministry portfolios that the work of governments is typically divided into.”

On her second point, Gal explains that she is not calling for a special “ministry” just to regulate AI – just like how the Internet or the semi-conductor industry does not need one.

“It will require more governance flexibility and collaboration to ensure the technology is safely applied across various usage areas. This type of flexible regulatory model will help more industries introduce AI technologies safely and without needing to launch full-fledged investigations every single time.”

“Preparing for system failure”

To pre-empt the kinds of scenarios of autonomous vehicles being remotely hijacked through cyber attacks, we would have to anticipate the unintended consequences of AI in the design, application and review stages of AI. This is not straightforward, because of the “unknown unknowns” – we do not know what we do not know.

Nevertheless, Gal says, “There are things that we can do to prepare for system failure that could harm users. This is a key part of the work we do at the IEEE P7009 Standard working group on the Fail-Safe Design of Autonomous and Semi-Autonomous Systems, which I chair.” This standard seeks to establish a practical, technical baseline of specific methodologies and tools for the development, implementation, and use of effective fail-safe mechanisms in autonomous and semi-autonomous systems.²²

“In addition, data-sharing is critical to prevent repeating accidents by ensuring that system failure would not occur across different systems,” Gal adds. “This exists in aviation systems, but not in drones or autonomous vehicles – and it should.”

Finally, Gal suggests people can create likely projections and scenarios based on past usage patterns and known physical and virtual vulnerabilities in both humans and technology to try and estimate what cases they would like to avoid and how to best do so. “This is a tactic often utilised in military and security planning, and should be made available to engineers developing the technology,” she says.

With every incident that occurs, the database of projections and scenarios grows and becomes more accurate. While we need to understand that accidents are bound to happen and that we cannot always prevent them, we should actively work to understand which accidents can be minimised and, in time, prevented before we deploy the technology any further.

“This requires a constant feedback loop between users and designers, as well as regulators – technological advancements should be a conversation,” Gal concludes.

The accountancy profession: professional scepticism

Human accountants are still needed for the professional scepticism they bring to businesses and financial reporting today, Enoch Ch’ng points out.

That is because AI is ultimately dependent on the quality and integrity of the data that is fed into it – the “garbage in, garbage out” concept of computer science. Ch’ng believes that human accountants are still needed to check for data manipulation, such as for evidence of inconsistencies or outright fraud.

Data in the form of PDF files, such as an invoice or a receipt, could be easily doctored for malicious purposes. Such doctored data may still elude the AI and data systems of today, particularly the less sophisticated ones, which are what smaller companies and organisations can afford to implement. It is not uncommon for macros embedded in PDF files, excel spreadsheets or even word documents to contain malware, which can cause data to be distorted.

Given the current state of the technology, the human brain is still formidable in terms of processing speed, processing power, and particularly in the ability to learn multi-functional skills quicker and faster. Human



accountants do have some advantage in allowing themselves to be ahead of “AI accountants,” unless a situation such as that in the tale of the Tortoise and the Hare from Aesop’s Fables occurs.

With the growing complexity of business and financial reporting, in part due to the more widespread use of estimates and management judgment, as well as business model changes, professional skepticism in audit is already becoming more highly valued than before.

In 2015, the International Auditing and Assurance Standards Board (IAASB), International Ethics Standards Board for Accountants (IESBA), and the International Accounting Education Standards Board (IAESB) convened the Professional Skepticism Working Group – a small, cross-representational working group — to formulate views on whether and how each of their sets of international standards could further

contribute to strengthening the understanding and application of the concept of professional skepticism in the realm of audit.

With technological developments such as the emergence of AI, the growing consensus is that professional skepticism does indeed lie at the heart of a quality audit.²³

However, this is merely an issue of AI and other technology not being sophisticated enough yet to detect fraud. Even in this day and age of ever pervasive technological use, incidents of fraud and unethical practices have not gone away – rather, their complexities have grown. To be ahead of the curve, the guardians of data – such as accountants and auditors, in the area of financial statements and the like – would always be needed for their exercise of professional skepticism.

Conclusion

As with any new trend, a healthy perspective towards AI would be to embrace the possibilities it can bring, while also being alert to the potential risks.

Adopting a balanced perspective is key for Richard Koh. Having shared about the risks and threats AI may pose, Koh concludes, “The AI advances in recent years have been powered by massive improvements in computing power, larger scale data storage capabilities and even more complex algorithms in machine learning and deep neural networks. With hyper-scale cloud services platforms, the capabilities of AI and data analytics have become more democratised

among businesses and organisations including government, education and healthcare.”

“The current use of AI in Singapore is at an early adopter’s stage with a cautiously optimistic view towards the hype around AI. In other countries like the US and China, the sheer number of businesses purportedly using AI in their businesses would create an impression of overwhelming use and adoption, but the optics distortion must be tempered with the real positive impact that AI is bringing to the businesses and organisations who have integrated AI into their businesses.”

23 International Auditing and Assurance Standards Board (IAASB), International Ethics Standards Board for Accountants (IESBA), and the International Accounting Education Standards Board (IAESB), “Toward Enhanced Professional Skepticism”: <https://www.ifac.org/publications-resources/toward-enhanced-professional-skepticism>

With thanks to:

Chng Sok Hui, Chief Financial Officer, DBS Bank

Enoch Ch'ng, former Associate Professor of Information Systems (Education) at Singapore Management University, and Senior Advisor, Data Science Rex

Danit Gal, Chair of the Outreach Committee of The Institute of Electrical and Electronics Engineers (IEEE) Global AI Ethics Initiative, and Project Assistant Professor at Keio University, Japan

Richard Koh, Chief Technology Officer, Microsoft Singapore

David Leow, Managing Director, Thaler Group

Matt Pollins, Partner and Head of Telecom, Media & Technology Practice, CMS Singapore

Yeong Zee Kin, Assistant Chief Executive (Data Innovation and Protection Group) of the Infocomm Media Development Authority of Singapore, and Deputy Commissioner, Personal Data Protection Commission (PDPC)

This publication has been prepared for general informational purposes only and is not intended to be relied upon as accounting, legal or other professional advice. Please refer to your advisors for specific advice.

ABOUT THE INSTITUTE OF SINGAPORE CHARTERED ACCOUNTANTS

The Institute of Singapore Chartered Accountants (ISCA) is the national accountancy body of Singapore. ISCA's vision is to be a globally recognised professional accountancy body, bringing value to our members, the profession and wider community. There are over 32,000 ISCA members making their stride in businesses across industries in Singapore and around the world.

Established in 1963, ISCA is an advocate of the interests of the profession. Possessing a Global Mindset, with Asian Insights, ISCA leverages its regional expertise, knowledge, and networks with diverse stakeholders to contribute towards Singapore's transformation into a global accountancy hub.

ISCA is the Administrator of the Singapore CA Qualification and the Designated Entity to confer the Chartered Accountant of Singapore - CA (Singapore) - designation.

ISCA is a member of Chartered Accountants Worldwide (CAW). CAW brings together 12 chartered accountancy bodies connecting and representing the interests of over 1.7 million members and students globally.

For more information, visit www.isca.org.sg.



60 Cecil Street, ISCA House,
Singapore 049709
TEL +65 6749 8060 FAX +65 6749 8061
www.isca.org.sg