

Responsible AI in
Financial Services:

Foundations and Future



accenture



Monetary Authority
of Singapore

ELEVANDI

RESPONSIBLE AI in FINANCIAL SERVICES: FOUNDATIONS & FUTURE

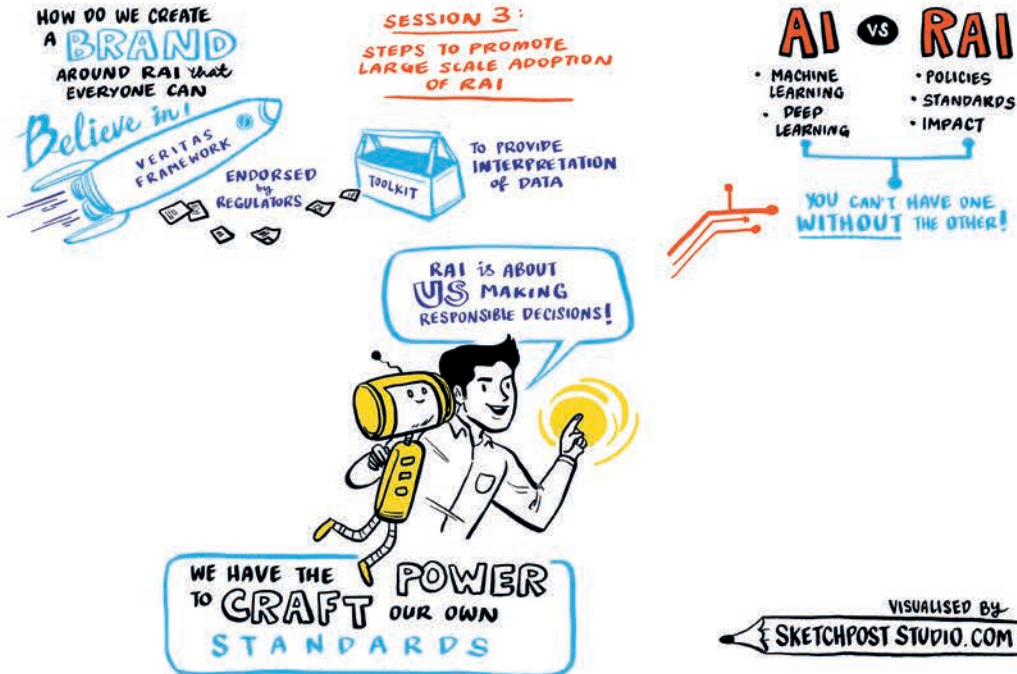


Table of Contents

Introduction	4
Part 1: Context of AI & RAI adoption across financial services	5
What is AI and what is Responsible AI?	5
The regulators' view	7
Principles rather than prescriptions	9
Part 2: RAI Adoption Journey 1:	
Current implementation and challenges	11
Decision-making challenges:	
Balancing risk and reward in an accelerating race	12
Governance challenges: Embedding RAI in the organisation	14
Investment challenges: Explaining RAI in terms of ROI	15
Part 3: RAI Adoption Journey 2:	
Steps to promote large-scale adoption	16
The need for standardisation	17
Don't wait for perfect harmony	18
Education, talent and certification	19
Conclusions, suggested solutions & key findings	21
Acknowledgement	23
References	24

Introduction

The use of Artificial Intelligence (AI) and data analytics in financial services has accelerated dramatically in recent years across a range of business functions, from credit assessment and risk underwriting to customer communications, marketing, recruitment and hundreds of others. As the use of these technologies has become more common, concerns have arisen in many jurisdictions about whether decisions taken with the use of AI are being made with sufficient fairness, ethical consideration, accountability and transparency.

These principles, identified and unified under the acronym “FEAT” by the Monetary Authority of Singapore (MAS) and the financial industry in late 2018,¹ underpin the concept of Responsible AI (RAI). Industry stakeholders around the world have articulated similar concerns about the need for such principles to govern AI and its outcomes. Yet given the relative novelty of the technologies, there has to date been little alignment on what actions are needed to promote a collaborative, consistent and global adoption of RAI standards.

This was one aim of a roundtable on RAI convened at the 2022 Singapore Fintech Festival (SFF), chaired by MAS and Accenture.² This paper is based on discussions at that roundtable between 18 expert participants, representing global banks and insurance companies, technology service providers and regulators. As the event was held under the Chatham House rule,³ comments or points of analysis in this paper are not attributed to specific participants or institutions, except where they refer to information already in the public domain.

While there was considerable agreement between discussants on the need for RAI frameworks, there were also points of difference regarding the challenges of implementing RAI across the industry, the role of different stakeholders in establishing RAI standards, and the steps necessary to promote large-scale adoption. As the meeting did not include a formal mechanism to establish consensus on any specific issue, this paper presents a range of opinions without making concrete recommendations. Conclusions and findings presented at the end of the paper are tentative and should not be held to reflect the opinions of any one participant, or the publishers of this report.

MAS is preparing to address some of the challenges raised during the discussion. Together with Accenture and the Singapore Fintech Association, MAS is working with the industry to promote and develop the RAI ecosystem in Singapore.

Part 1:

Context of AI & RAI adoption across financial services

The long history of research into artificial intelligence has endured more than one “AI winter”; that is, a steep falloff in investment and reduced academic interest in the field. But in recent years an AI spring has blossomed across many areas, including financial services. Whether for claims processing and underwriting for insurers, credit assessment and customer engagement for banks, or marketing and recruitment across the whole industry, the use of AI is increasingly pervasive. This is having a direct impact on all stakeholders, from customers to shareholders to employees.

What is AI and what is Responsible AI?

The International Organization for Standardization (ISO) defines AI as “an interdisciplinary field, usually regarded as a branch of computer science, dealing with models and systems for the performance of functions generally associated with human intelligence, such as reasoning and learning.”⁴

The term “machine learning”, or ML, is a subcategory of AI whereby, according to one definition from the UK’s Financial Conduct Authority (FCA), “computer programmes fit a model or recognise patterns from data, without being explicitly programmed and with limited or no human intervention.”⁵

As the Bank of England and the FCA noted in a February 2022 report that summarised the findings of a public-private forum on AI in financial services, a “key characteristic of AI systems is their capacity for autonomous decision-making”. It added: “This can have profound implications for how to govern the technology and its outcomes, including ensuring effective accountability and responsibility.”⁶

These implications lead directly to the need for Responsible AI (RAI), defined by Accenture as “the practice of designing, developing, and deploying AI with good intention to empower employees and businesses, and fairly impact customers and society—allowing companies to engender trust and scale AI with confidence”.⁷

Meanwhile, the Veritas Consortium, a Singapore group of industry leaders and the Monetary Authority of Singapore (MAS), notes RAI depends on the application of the FEAT principles in AI technology: Fairness, Ethics, Accountability and Transparency.⁸

The COVID pandemic accelerated the adoption of AI by financial services providers. A study by the Bank of England (BoE) and the Financial Conduct Authority (FCA) in 2019 found that two-thirds of financial firms in the UK were using ML in their businesses, although it noted that many had only a limited number of use cases.⁹

By 2022, though, a follow-up survey found 72% of respondent institutions were using or developing ML applications, and across many more business areas. It also found ML applications are now more advanced and embedded in day-to-day operations, with 79% of ML applications in the latter stages of development. Moreover, the 2022 report found that the trend looks set to continue, with the overall median number of ML applications set to increase by 3.5 times over the next three years.¹⁰



The overall median number of ML applications is set to increase by 3.5 times over the next three years.



Source: BoE and FCA survey, 2022

Many participants at the SFF roundtable remarked on the accelerated speed of AI adoption. Discussants also noted that there has been an acceleration in interest in RAI, particularly given the more rapid deployment of MLOps (that is, the set of practices that aim to apply machine learning to different applications reliably and efficiently).¹¹ Yet as one technology service provider noted, though RAI had been a topic of conversation for many years, financial services (FS) providers were moving faster on AI in general than they were on RAI, for which incentives were still lacking.

To be sure, no financial services firm has ever set out to use AI irresponsibly. Yet the nature of the risks that the new technology presents means industry players now must pay extra attention to ensure AI is being deployed responsibly, in a manner that minimises those risks. That poses a challenge to regulators seeking to provide a level playing field, and a clear framework, for the safe development of AI-based solutions.

The regulators' view

With the acceleration of the adoption of AI, regulators now face an inflexion point. Discussants at the SFF roundtable representing global regulatory authorities reiterated that they recognised the clear benefits of AI technology and were cautious about not blocking innovation, but also that they recognised the need to mitigate potential risks, such as where automated and opaque decision-making reinforces biases or discrimination. Though such risks are applicable across all economies, they are acute in jurisdictions in which vulnerable communities face development challenges, where countries often lack data protection laws, let alone AI regulation.

In short, regulators everywhere face the responsibility to allow innovative firms to reap the benefits that AI promises, provided that they do so without compromising individuals' data rights, perpetuating or reinforcing inequalities, or risking the arbitrary exclusion from financial services of communities through bias and discrimination.

The reality, though, is that regulators operate at a disadvantage, in that they cannot know with absolute certainty how AI and ML technologies are being applied by companies at the cutting edge – and the readiness of the industry for these developments. For its part, industry also cannot predict the direction of travel.

“ **[Innovators are] creating brilliant things. We don't know [everything that] is happening; we don't even know what we don't know. But we know that we don't know a lot.** ”

Panel Discussant

Nevertheless, discussants agreed that clarity on regulation is needed as a catalyst for further innovation — to establish a level playing field that would allow investment in RAI with clarity and certainty. Given the multinational operations of financial services firms, transnational co-ordination and regulatory alignment is also a clear priority.¹²

To be sure, multiple international bodies have already investigated the issue. These include:

- A 2017 landmark investigation by the Financial Stability Board (FSB), considering the financial stability implications of the growing use of AI and ML in financial services;¹³
- A 2020 report from IOSCO on the use of AI and ML by market intermediaries and asset managers;¹⁴
- A 2021 study by the Financial Action Task Force (FATF) on opportunities and challenges in the use of AI for AML/CFT;¹⁵
- A 2022 issues paper published by the International Association of Insurance Supervisors on the use of big data analytics in insurance, focusing on the fair treatment of insurance customers;¹⁶
- The issuance by the Basel Committee on Banking Supervision (BCBS) of a newsletter on AI and ML in the sector.¹⁷

Then there are complementary national and supra-national initiatives on AI in general that have implications across various sectors. These stem in part from a 2019 OECD report setting out values-based AI principles and recommendations for policymakers that “focus on how governments and other actors can shape a human-centric approach to trustworthy AI”.¹⁸

These tenets were subsequently endorsed by the G20¹⁹ and have been used as a starting point for numerous national-level principles. At the national level a wide range of different types of AI-related publications and guidance notes have been issued, from those covering high-level cross-sector principles to consultation papers specifically on AI in the financial sector. Principles that specifically target the use of AI in the financial sector have been issued, for example, in the European Union, Germany, Hong Kong, the Netherlands, Singapore and the United States.

Notable initiatives on a broader scale, meanwhile, include the European Commission’s adoption of a proposal that would harmonise rules on AI in general and that classifies risk levels of differing AI systems, from low/minimal risk to high risk and unacceptable risk (the last of which would be prohibited).²⁰ The US government, meanwhile, has proposed a blueprint for an “AI Bill of Rights” that would enshrine five principles to guide the design and use of AI systems.²¹ Some countries are instituting brand-new regulatory authorities: in January 2022, Spain became the first European country to create an AI supervisory agency, for instance.²²

Principles rather than prescriptions

Following these extensive preparatory investigations, regulators at the roundtable agreed that it is better at this stage of the development of AI to establish principles in collaboration with industry players, rather than issue prescriptive rules or prohibitions. Two salient public examples of this approach were highlighted.

The first, by the BoE and the FCA, was to establish a public-private forum to discuss the implications of the use of AI in financial services.²³ The AIPPF brought together academics, financial services and fintech firms for a yearlong consultation culminating in a final report, published in early 2022.²⁴ The report explored the barriers to adoption, challenges and risks in three crucial areas of AI development: data, model risk and governance (see figure 1).

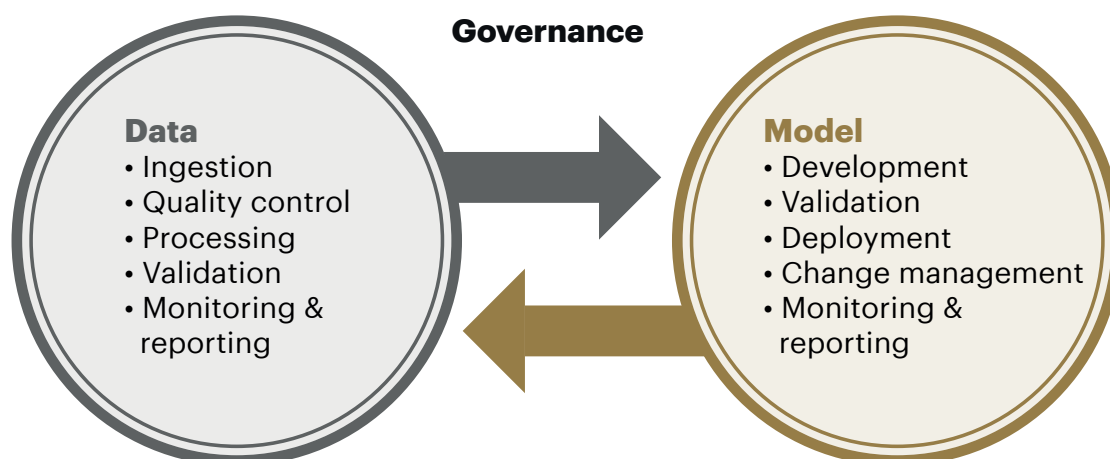


Figure 1: AI system

Source: Final report: Artificial Intelligence Public-Private Forum, Bank of England and Financial Conduct Authority: <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>

The second example discussed at the roundtable was the systematic review of RAI undertaken by an MAS-led consortium of industry stakeholders in Singapore, the Veritas Initiative. This initiative followed the co-creation in 2018 by MAS and the financial industry of the FEAT principles, and has the aim of enabling financial institutions to evaluate their solutions that are driven by AI and Data Analytics (AIDA) against those principles to strengthen internal governance around the application of AI and the management and use of data.²⁵

Conducted in various phases in collaboration with stakeholders from across the financial services industry and academia, Veritas published comprehensive assessment methodologies for each of the four principles in February 2022.²⁶ To accelerate FIs' adoption of the Veritas methodologies and principles, the consortium also developed an open-source software toolkit that enables the automation of the fairness metrics assessment, and allows for visualisation of the interface for fairness assessment and for plug-ins to integrate with an FI's IT systems.²⁷

One clear finding of such initiatives, the roundtable heard, was the need, from the industry's point of view, for regulatory clarity; not only in terms of prospective new rules, but also as regards guidance provided by regulators (such as MAS's FEAT principles) that may have some bearing on the application of AI in financial services.

In terms of challenges facing FIs in implementing RAI, the UK investigation also found, *inter alia*, that data governance is a particular challenge, given the complex patchwork of data-related regulations and the rising use of unstructured or "alternative" data; that model risk is a rising issue given the complexity of the inputs and the intricacy of the models themselves (and that being able to explain model outputs is vital); and that for governance, the issue of accountability is of paramount importance, with the need for a centralised body within firms that should set AI governance standards.²⁸

These issues — data quality, building trust and fairness into increasingly complex AI models, and the need for effective governance — form the core of the challenge that financial services firms face in adopting and implementing RAI, which was the focus of the second part of the roundtable and Part 2 of this paper.



Part 2:

RAI Adoption Journey 1: Current implementation and challenges

As noted in Part 1, the adoption of AI in general in financial services has accelerated dramatically in recent years. Banking industry discussants at the SFF roundtable described the increasing adoption of AI and ML in three principal areas. First, in delivering innovative products and services, and in speeding up the delivery of those services (through, for instance, the use of technologies like facial recognition). Second, in cybersecurity — protecting customers’ money and data, and preventing financial crime. And third, in helping benefit society through applying AI in reaching sustainability goals — helping close data gaps, for example.

Within these broad categories, the number of functional use cases is growing quickly: recent research by the BoE and FCA found customer engagement and risk management to be the two areas with the most ML applications (accounting for 28% and 23% of all applications respectively). In third place, though, was the “miscellaneous” category, covering a broad range of functions from HR to legal departments — illustrating the breadth of potential use cases across any business.²⁹ One roundtable attendee spoke of having 260 use cases across their organisation.

When it comes to the adoption of Responsible AI, as opposed to AI in general, the picture is less clear. A straw poll conducted at the roundtable found that 27% of attendees’ institutions were at “Stage 1” in the process of implementation; that is, they were at the “wait and see” stage, having investigated RAI though with no applications yet rolled out.

A further 19% were at Stage 2, which meant they were undertaking internal discussions on use cases and budgeting. Another 30% were at Stage 3, in that they were preparing to implement RAI and had set up the necessary policies and governance frameworks to do so. Finally, 24% were at Stage 4, which meant they have already implemented and enforced RAI.




Figure 2: RAI rollout stage
Source: Audience poll, SFF roundtable, November 2nd 2022

Decision-making challenges: Balancing risk and reward in an accelerating race

The sheer range of potential applications underscores, too, the first challenge of RAI adoption: deciding where to act, and how to balance the opportunity and risk of doing so in a rapidly changing environment. Each AI use case is different, as are the risks of failing to ensure RAI principles are adequately factored in.

At the heart of any decision is an ethical deliberation about whether data should be used, even given consent and compliance with applicable regulations. In many cases the potential benefits for end-users will justify the risks of allowing the data to be used: anyone using Google Maps, for example, is prepared to sacrifice location data privacy for the convenience of the service it offers.

Even with such implied consent, in each use case, discussants said, questions must be answered around whether the institution is using the correct data; whether that dataset contains biases and, if so, how that might affect outcomes; whether the outcomes are being analysed fairly; and in what manner ethics should be applied in each of these considerations.

The result is the need to judge the opportunities and risks finely: mistakes risk undermining customers' trust, while inaction might jeopardise the interests of other stakeholders, such as shareholders. This means assessing the materiality of the use case versus the effort required in mitigating associated material risks: naturally the higher the risk, the higher the effort required, but balancing these for investment decision-making is not straightforward.



This is partly because the risks of error may have society-level implications. An incorrectly calibrated AI underwriting model, for instance, could lead to premiums rising for everyone — and in the extreme, the breakdown of the concept of sustainable risk-pooling. Or, to take a notorious example, gender inequities have been reinforced by the skewing of many datasets towards men in fields as diverse as economics, healthcare, education and public policy.³⁰ The need for a clear and responsible understanding of risks like these is therefore enormous.

The issue is exacerbated by the fact that though regulators in different jurisdictions might profess similar RAI principles, in practice they may look at different things as they are applied in specific use cases. Definitions vary across jurisdictions, too, not just for AI and ML but for what responsibility means — and to whom.

Financial firms face responsibility to customers primarily, specifically by providing services equitably across society, but also employees, including the data scientists who might be held responsible if their solutions lead to unintended adverse consequences or contravene rapidly changing regulations. FS firms are also responsible to their business partners in the AI service supply chain, which rely on the data provided to them but might not have sufficient control over its provenance or quality.

In short, standards of responsibility are evolving almost as quickly as the technology. This amplifies concerns about whether there is a level playing field for financial firms compared to those employing AI in other industries — if the thresholds in finance are too high to allow for innovation.

“It’s happening faster than we can keep up with,” one industry executive said, summing up these crosscutting issues. “It’s this uncertainty, that we get too far ahead of this new regulatory regime. And then [regulators] come in and say... ‘We decided that we shouldn’t do it that way.’ That’s one of the big hesitations. But I don’t think anybody can afford to wait, we just have to go.”

“

It’s this uncertainty, that we get too far ahead of this new regulatory regime. And then [regulators] come in and say... ‘We decided that we shouldn’t do it that way.’ That’s one of the big hesitations.

”

Financial Services Executive

Governance challenges: Embedding RAI in the organisation

Assuming ethical and opportunity-cost considerations are resolved on whether to proceed, discussants at the forum agreed that implementing RAI carries operational challenges throughout project lifecycles.

Consequently, two related points garnered broad agreement: that RAI principles needed to be embedded in the organisation and that the concept of RAI in general could not be a bolt-on, separate activity. As one discussant noted, “It needs to be part and parcel of [the roles of] everyone who’s involved, whether it’s a product person coming up with a use case, a data scientist, or a governor; all of them have to have this view of their responsibility.”

Throughout, this requires clear oversight at senior levels, which in turn raises questions about the kind of background and training required of people in decision-making roles: an understanding of data science and the technology to be deployed is necessary to grasp dependencies and operational risks. Senior executives and board members still need education in this regard, while a pipeline of relevant talent is needed to be able to incorporate RAI tools as part of the “operating DNA” of the firm.

At the more granular level, the practical aspects of rolling out AI solutions face model and governance challenges of the kind identified in the consultations carried out by the BoE and the FCA, as well as by MAS. Without adequate maintenance in human and AI operations, AI models cannot be used safely: in the pre-production phase, risks need to be clearly explained to enable informed sign-off.

Concerns over the RAI principles of interpretability and explainability down the chain is another concern. The characteristics and parameters of models that take decisions or make predictions must be clear to internal teams in the compliance and risk functions, who must be able to interrogate them.

Difficulties explaining AI solutions to end-users is also giving some financial services executives pause. It is not simply a matter, for instance, of revealing the degree of precision or confidence with which AI models might make predictions. Something that is right 73% of the time will be wrong 27% of the time and the impact of being wrong is not clearly understood, since the risk of a false positive may outweigh that of a false negative. Weather reports are a good analogy: though a model might report only a 10% chance of rain, the human forecaster might report this as 25%, since end-users’ anger at getting wet would probably exceed their irritation at having carried an umbrella for no reason.

As one executive put it, “I think the hesitation is that we’re not sure that our end-users are as prepared as they need to be to really understand what these models are telling them.”

Investment challenges: Explaining RAI in terms of ROI

While the impetus for RAI is moral and ethical, companies must still justify business decisions in terms of return on investment. As such, one key challenge is ensuring data scientists and business decision-makers can speak the same language.

Currently, according to a technology services provider discussant, decision-makers are still likely to see RAI as an additional step in the workflow or an extra piece of technology. Ultimately, though, the explanation can be made in terms of improving customer experience — by ensuring, for instance, that no groups of potential customers are unfairly excluded. Or that by being open and transparent as a matter of “hygiene” will yield benefits in terms of increased trust scores, and the greater likelihood of appealing to privacy- and data-rights-conscious millennial or Gen Z customers. That, ultimately, could be a differentiating commercial factor, rather than a net expense.

Ultimately, the kinds of risks and concerns that discussants noted must be addressed for RAI to be adopted on a large scale. The means to do that was the subject of the third part of the roundtable, and Part 3 of this paper.



Part 3:

RAI Adoption Journey 2: Steps to promote large-scale adoption

Ensuring the consistent, global adoption of RAI in FS requires that players undertake a mix of short-term and long-term steps across the following pillars:

- Building RAI foundations through standardisation (and potentially incorporating an open-source ecosystem for toolkits);
- Training and accreditation schemes;
- Widespread advocacy.

For the FS sector itself, participants in the SFF roundtable heard that the first step is to set out the principles or general rules that factor in the specific uses that apply to this industry. After all, FS firms and regulators cannot use the same set of RAI principles as might apply with, for instance, medical information or HR processes.

To be sure, there is as yet no global consensus on the best way to ensure the adoption of these RAI principles. Nevertheless, firms and regulators in some jurisdictions have concluded that the next step is to translate those principles into standards that articulate what is required in areas like internal governance, machine-learning operations and communicating with customers, and that the final step would then be to develop the tools to implement those standards.



The need for standardisation

A central challenge is that it will prove highly difficult to get the world's regulators to agree to a single set of rules. For that reason, discussants said, it would be logical to create an overarching framework that covers, say, 90 percent of an RAI approach, with individual regulators and central banks working with industry to refine the remaining 10 percent to match local needs. That would culminate in a local standard.

In this way, regulators in different countries that have different skills and needs can sign up to an overarching framework, which they can tailor to a standard that best fits their requirements. The benefit of such an approach is that it would see the majority of countries follow a similar RAI path.

That is not the case currently. There are various frameworks under discussion in different forums, and about 60 countries have made more than 700 AI policy changes in recent years, one discussant estimated. Additionally, new metrics are devised on a regular basis, which adds to the complexity. Managing this across a range of regulators that are operating bespoke, unharmonised systems is extremely challenging and leads to fragmentation.

While it makes sense, then, to harmonise regulations (including with other relevant data-related standards), doing this successfully requires considering how to translate from principles to practice.

It is useful to consider this starting point: RAI will at some point use personal data as inputs, one regulatory participant noted, so adopting a customer's perspective is optimal when considering the best approach, with this underpinned by principles, standards and verification.

Importantly, standardisation covers more than techniques: standardisation around processes is also crucial — particularly embedding this in areas like, for instance, the data-science workflow.

An additional complication is that creating an RAI approach means adapting legacy systems. While it is logical to bring in a data science team to build best-in-class RAI models, these might then run on a mainframe that operates a much older, less flexible programming language. Consequently, firms need to know how these interact — for instance, how the RAI model ingests data from the database. Ultimately, this requires demystifying often-complex legacy systems.



Don't wait for perfect harmony

In short, firms face an uphill battle to create RAI approaches that are fit for purpose, which is why discussants expressed doubt that this could wait for harmonisation. As one put it, although “regulation is absolutely necessary, we do not have a level playing field, and many of the countries in which we operate don't have this regulation”. The inevitable question is how, in the interim, firms can deploy AI tools that are responsible.



[Although] regulation is absolutely necessary, we do not have a level playing field, and many of the countries in which we operate don't have this regulation.



Financial Services Executive

One route could be to take the open-source approach, on the grounds that — as Professor Lawrence Lessig wrote two decades ago — “code is law”.³¹ In this way, the RAI industry could piggyback on developments in the open-source community to ensure more robust systems are made more broadly available, with the principles and standards embedded in the code. Broadening the scope for people to contribute to RAI systems could drive more widespread adoption.

Positive though that could be, regulators say the problem is that open source cannot be adapted for hundreds of different use cases. Instead, each must be considered separately, which means open source does not provide a short cut to success.

And while regulators acknowledged the need for greater cooperation with the FS industry, they have their own views about which use cases are more important from a supervisory perspective. Because a standard cannot cover every one of the dozens of use cases, both sides could start by working together to identify the most important.

Additionally, there needs to be agreement on various parameters and their relative importance. For instance, what fairness metrics should be extracted for this particular use case? Or what is the stress test for a particular kind of externality? Currently, regulators and academics lack answers to these aspects, and deriving those will require a joint effort by all parties.

Understanding the impact on attempts to standardise rules of different use cases, one industry discussant said, demonstrated the importance of starting by prioritising the most crucial uses. “Clearly, we can’t do hundreds of use cases, although with overarching principles we could find some common ground,” they said.

In addition, the discussant said, there should be a greater focus on fairness, which is often missing from existing metrics which focus more on business KPIs. Linked to this, those involved should also take responsibility for figuring out how to translate data science performance into business and risk implications that affect the end-customer and other stakeholders — which, while it might not be done most practically as part of a standard, could at least be emphasised as a principle.

Such an approach chimes with the need to balance the risk and reward of the steps taken, and to avoid being overly prescriptive. Ultimately, as another discussant said, success will stem from the skills of the individuals who develop, maintain and assess the AI solutions, and the talents they bring to that task.

Education, talent and certification

In an area as new as RAI, it is inevitable that there is a shortage of skills and talent. Closing this starts by recognising the difference between AI and RAI, noted one industry discussant. The skills needed for AI, they said, typically revolve around analytics, machine-learning and deep-learning, whereas those needed for RAI also include policies, standards, governance, its business impact and its impact on the real world — all of which must be factored in.

In short, success in the RAI sphere requires that you must be able “to write Python and also write policy — and you need everything in between to be effective in RAI”. And while the industry is not yet facing a crisis in RAI skills, discussants said (not least down to the relatively low scale of AI adoption compared to other historical technologies), that could yet arise.

“

Success in the RAI sphere requires that you must be able “to write Python and also write policy”.

”

Financial Services Executive

Closing these gaps requires an array of different training criteria, the SFF audience heard. First, data scientists need to learn the process of engineering where responsible metrics are being used. And more is needed to deal with the lack of translators: those staff who understand the data scientist's world and the business world, and who can bridge the two. One shortfall is a lack of materials and education for translators, whether they are product managers or those who work alongside data scientists.

Over time, as the various business metrics employed evolve, there will also be a need to translate those into measurements that are meaningful for non-specialist staff, including those in the C-suite. There will also be a need for education on how techniques are applied and on how standardisation fits into the broader picture.

Additionally, firms need to establish ongoing responsibility and create a risk culture that fits RAI models. One industry discussant said there is an erroneous belief that RAI models are already compliant, pre-production. Yet responsibility does not end at production, they said, and models must be assessed for gaps that can affect performance, including whether the code is written to comply with production architectures.

Creating the right risk culture starts at the pre-production stage by examining the true dependencies for success. Although the last decade has been about experimentation, today, as RAI moves into practice, real business value and risks must be sustainable. Teams that build RAI models must ensure they are not misled into thinking naively that what they have created is production-ready.

Away from those directly involved in crafting RAI solutions, it is also crucial that others – from the C-suite to staff, and from customers to the general public – understand the implications of RAI. Education is vital for mass-scale roll-out, discussants noted, including at the university level to ensure professional trainers at tertiary institutions have the relevant training knowledge.

Ultimately, all stakeholders will need to be involved to ensure education is sufficient in this regard. Financial institutions called for RAI to be part of academic training, “[so] we would be able to raise the next generation of people who actually embed Responsible AI into the way that they do things”. Regulators, for their part, said they would need to hear from financial institutions about the internal training they require, “then we will find the experts to design the curriculum for you”.



Conclusion, suggested solutions & key findings


As the adoption of AI accelerates, and as interest in RAI climbs, it is timely that institutions around the world are working to ensure AI is used responsibly – that AI systems, which are characterised by their “capacity for autonomous decision-making”, as the BoE and FCA stated in their 2022 report³², operate fairly, ethically, and with accountability and transparency.

As Accenture has pointed out, this requires taking the route of “designing, developing, and deploying AI with good intention to empower employees and businesses, and fairly impact customers and society—allowing companies to engender trust and scale AI with confidence”.³³

However, aligning on how best to drive large-scale adoption has to date proved challenging for a range of reasons, including a lack of regulatory clarity, insufficient understanding across the board about the steps to take, a plethora of competing frameworks, the roles different stakeholders should play in establishing RAI standards, and a lack of incentives for RAI adoption.

While regulators recognise the benefits of AI technology, they are also cognisant of the need to tackle potential risks, including that automated and opaque decision-making can reinforce biases and discriminate against people – and particularly against vulnerable communities in jurisdictions facing development challenges, which might lack more basic laws for related areas like data protection. It is also the case that the industry is unable to predict exactly how organisations will apply AI and ML technologies, making it harder to devise suitable rules.

The balancing act for regulators is to allow innovation in AI and RAI while ensuring individuals’ data rights are respected, and that people and communities are not discriminated against or excluded from access to financial services.



To that end, the roundtable heard that the following aspects can help to increase RAI adoption:

- Regulations that are clear and aligned: this would act as a catalyst for further innovation, create a level playing field, and ensure RAI-related investments could be made with certainty.
- Coordination across jurisdictions, and greater cooperation among stakeholders.
- Integrating RAI into corporate operations (rather than considering it as a bolt-on), and creating a culture that ensures Responsible AI by design — including embedding it in the curricula of tertiary institutions, and educating corporate staff on RAI (including the C-suite and all business lines) and provide relevant training for decision-makers.
- Ensuring that data scientists can link with the business more broadly so that their work is harmonised with corporate needs and goals, and meets guidance such as the FEAT principles.

One element currently lacking in terms of a successful roll-out of RAI is regulation. A clear message from discussants at the 2022 SFF roundtable is that industry requires guidance in terms of specific standards and protocols that will allow it to map a path forward.

In this way, firms could strike a balance between governance needs and innovation, and apply their innovation skills to ensure they practice Responsible AI from design — and even prior to design. This could and should include the use of sandboxes to design and test solutions effectively and safely.

Given the global nature of RAI and its challenges, and given the variety of needs in different jurisdictions, a framework agreed by all (or most) participants would constitute an important step. Such a framework could be finessed by regulators to take account of local requirements, ensuring that while specifics varied between countries, the overall approaches and tenor would be broadly similar.

In this way, standardisation would be placed front and centre of a global RAI approach. That would also help to propel a linked issue: improved collaboration between industry and regulators for specific use cases, which is highly likely to happen and particularly if there can be agreement on basic open-source standards.

Last and not least, much more education is needed in firms and training establishments, including institutes of higher learning, to ensure staff working in the many areas where RAI will increasingly be applied are sufficiently skilled. In this way, the work done on RAI will be integrated into a machine-learning lifecycle – rather than it being added as an afterthought – with each person involved fully aware of their responsibilities.

Acknowledgement



Discussants

Sameer Gupta

Managing Director,
Chief Analytics Officer
DBS Bank

Vijay Jairaj

Executive Director,
Responsible AI and Data Ethics
Standard Chartered

Sunita Kannan

Global AI Sales and
Strategy Lead
Microsoft

Donald MacDonald

Head, Group Data Office
OCBC Bank

Johannes Ehrentraud

Senior Advisor
Bank for International Settlements

Dr Mara Balestrini

Digital Transformation
Lead Specialist
Inter-American Development Bank

Dr Ramine Tinati

Managing Director,
Applied Intelligence
Accenture

Andreas Viljoen

Head of Reporting, Disclosure,
Data Strategy and AI Policy
Bank of England

Manisha Khanna

Head, AI & Analytics,
Asia Pacific
SAS Institute Pte Ltd

Anup Kumar

Distinguished Engineer
& CTO for Data & AI
IBM

Lee Sarkin

Regional Chief Analytics Officer,
APAC, Middle East,
Africa (Life & Health)
Munich Re

Reuben Lim

Chief Operating Officer
Singapore Fintech Association

Hays McCormick

Data Science Fellow
BNY Mellon

Dr Chinmay Parikh

Innovation Lead
Financial Conduct Authority

Charmaine Wong

Head of Group BI & Analytics,
Group Head of ESG Data & Analytics
HSBC

Johnson Poh

Head, Group Enterprise AI and Data
United Overseas Bank

Yeong Zee Kin

Assistant Chief Executive (Data
Innovation and Protection Group),
IMDA & Deputy Commissioner, PDPC
Infocomm Media
Development Authority

Xuchun Li

Head of AI Development Office,
Fintech & Innovation Group
Monetary Authority of Singapore

Moderator

Damien Pang

Executive Director, (Data and
Technology Architecture) and
Deputy Chief FinTech Officer
Monetary Authority of Singapore

Irene Liu

Managing Director
Accenture

References

- ¹ <https://www.mas.gov.sg/schemes-and-initiatives/veritas>
- ² <https://www.fintechfestival.sg/sff-agend>
- ³ <https://www.chathamhouse.org/about-us/chatham-house-rule>
- ⁴ As cited in <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>
- ⁵ <https://www.fca.org.uk/publication/research/research-note-on-machine-learning-in-uk-financial-services.pdf>
- ⁶ <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>
- ⁷ <https://www.accenture.com/nz-en/services/applied-intelligence/ai-ethics-governance>
- ⁸ <https://www.mas.gov.sg/news/media-releases/2022/mas-led-industry-consortium-publishes-assessment-methodologies-for-responsible-use-of-ai-by-financial-institutions>
- ⁹ <https://www.fca.org.uk/publication/research/research-note-on-machine-learning-in-uk-financial-services.pdf>
- ¹⁰ <https://www.bankofengland.co.uk/Report/2022/machine-learning-in-uk-financial-services>
- ¹¹ <https://en.wikipedia.org/wiki/MLOps>
- ¹² As also noted in <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>
- ¹³ <https://www.fsb.org/2017/11/artificial-intelligence-and-machine-learning-in-financial-service/>
- ¹⁴ <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD658.pdf>
- ¹⁵ <https://www.fatf-gafi.org/media/fatf/documents/reports/Opportunities-Challenges-of-New-Technologies-for-AML-CFT.pdf>
- ¹⁶ <https://www.iaisweb.org/uploads/2022/01/200319-Issues-Paper-on-Use-of-Big-Data-Analytics-in-Insurance-FINAL.pdf>
- ¹⁷ https://www.bis.org/publ/bcbs_n127.htm
- ¹⁸ <https://oecd.ai/en/ai-principles>
- ¹⁹ <https://oecd.ai/en/wonk/documents/g20-ai-principles>
- ²⁰ <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=celex:52021PC0206>
- ²¹ <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- ²² <https://www.lexology.com/library/detail.aspx?g=d2d01036-bbdf-451a-b3e7-485f51148b76>
- ²³ <https://www.bankofengland.co.uk/research/fintech/ai-public-private-forum>
- ²⁴ <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>
- ²⁵ <https://www.mas.gov.sg/schemes-and-initiatives/veritas>
- ²⁶ <https://www.mas.gov.sg/news/media-releases/2022/mas-led-industry-consortium-publishes-assessment-methodologies-for-responsible-use-of-ai-by-financial-institutions>
- ²⁷ <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>
- ²⁸ <https://www.bankofengland.co.uk/Report/2022/machine-learning-in-uk-financial-services>
- ²⁹ As explained by Caroline Criado Pérez in her 2019 book “Invisible Women: Data Bias in a World Designed for Men”: <https://carolinecriadoperez.com/book/invisible-women/>
- ³⁰ <https://www.harvardmagazine.com/2000/01/code-is-law-html>
- ³¹ <https://www.harvardmagazine.com/2000/01/code-is-law-html>
- ³² <https://www.bankofengland.co.uk/-/media/boe/files/fintech/ai-public-private-forum-final-report.pdf?la=en&hash=F432B83794DDF3F580AC5A454F7DFF433D091AA5>
- ³³ <https://www.accenture.com/nz-en/services/applied-intelligence/ai-ethics-governance>

About Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Technology and Operations services and Accenture Song — all powered by the world's largest network of Advanced Technology and Intelligent Operations centers. Our 721,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities. Visit us at www.accenture.com.

About MAS

The Monetary Authority of Singapore (MAS) is Singapore's central bank and integrated financial regulator. As a central bank, MAS promotes sustained, non-inflationary economic growth through the conduct of monetary policy and close macroeconomic surveillance and analysis. It manages Singapore's exchange rate, official foreign reserves, and liquidity in the banking sector. As an integrated financial supervisor, MAS fosters a sound financial services sector through its prudential oversight of all financial institutions in Singapore — banks, insurers, capital market intermediaries, financial advisors and financial market infrastructure. It is also responsible for well-functioning financial markets, sound conduct, and investor education. MAS also works with the financial industry to promote Singapore as a dynamic international financial centre. It facilitates the development of infrastructure, adoption of technology, and upgrading of skills in the financial industry.

About Elevandi

Elevandi is set up by the Monetary Authority of Singapore (MAS) to foster an open dialogue between the public and private sectors to advance FinTech in the digital economy.

We work closely with governments, founders, investors, and corporate leaders to drive collaboration, education, and new sources of value at the industry and national levels.

Our initiatives have convened over 350,000 people since 2016 to drive the growth of FinTech through events, closed-door roundtables, investor programmes, educational initiatives, and research. Our flagship product is the Singapore FinTech Festival alongside fast-rising platforms, including the World FinTech Festival and Point Zero Forum.

Disclaimer

This content is provided for general information purposes and is not intended to be used in place of consultation with our professional advisors. This document may refer to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.