Minutes

Artificial Intelligence Public-Private Forum - First meeting

12 October 2020

Attendees

Co-Chair	Organisation
Ramsden, Dave	Bank of England
Mills, Sheldon	Financial Conduct Authority

Moderator	Organisation
Saporta, Victoria	Bank of England

Member	Organisation
Barto, Jason	Amazon Web Services
Browne, Fiona	Datactics
Campos-Zabala, Javier	Experian
Christensen, Hugh	Amazon Web Services
Dewar, Michael	Mastercard
Dorobantu, Cosmina	Alan Turing Institute
Gadd, Sarah	Credit Suisse
Kellett, Dan	Capital One UK
Kirkham, Rachel	Mindbridge AI
Kundu, Shameek	Standard Chartered
Lennard, Jessica	Visa
Moniz, Andy	Putnam Investments
Morris, Owen	Aviva PLC
Morrison, Gwilym	Royal London
Rees, Harriet	Starling Bank
Rosenshine, Kate	Microsoft UK
Sandhu, Jas	Royal Bank of Canada
Shi-Nash, Amy	HSBC
Tetlow, Phil	IBM UK
Treleaven, Philip	University College London
Tromans, James	Google Cloud

Observer	Organisation
Buckley, Oliver	Centre for Data Ethics and Innovation
Dipple-Johnstone, James	Information Commissioner's Office
Mountford, Laura	HM Treasury
Yallop, Mark	FICC Markets Standards Board

Item 1. Opening remarks by co-chairs and moderator

As co-chairs, Dave Ramsden and Sheldon Mills, welcomed the members and observers to the first meeting of the Artificial Intelligence Public-Private Forum (AIPPF). They emphasised the importance of the AIPPF, namely the ability for stakeholders to collaboratively identify issues and remedies, and its potential to help the public and private sectors better understand the key issues, clarify the priorities and determine what actions may be needed to support the safe adoption of Artificial Intelligence (AI) in UK financial services.

Victoria Saporta also welcomed participants as moderator for the meeting, noting that the AIPPF launch video had been published the same morning, introduced the observers and set out the objectives and procedures of the meeting, as well as the overall purpose and scope of the AIPPF, as set out in the terms-of-reference.

The moderator noted that the Forum was being held under the Chatham House Rule.

Item 2. Roundtable discussion

The moderator explained how the first Forum meeting should help frame the key themes and agendas for the subsequent three meetings. The members then responded to a wide-ranging set of questions circulated prior to the meeting (see below), followed by a short discussion. After the discussion, the Observers were invited to share their reflections. They emphasised that the conversation had been very interesting and useful, especially in providing direction for the next Forum meeting and future work of the AIPPF.

Question 1: AI is developing rapidly and the COVID-19 crisis may accelerate some trends and slow the pace of others. What do you see as the key short and long-term trends in use-cases and capabilities of AI?

Several members expressed the view that Covid-19 has accelerated the pace of automation and adoption of AI in financial services. As a result, it was argued that firms need to keep up with appropriate controls and focus on the resilience of their AI systems in the short-term. That may suggest a need for further work on auditing AI algorithms, which was already a key area of interest prior to Covid-19. In the longer-term, firms will also need to think about AI and data management in a more holistic manner and within the context of their wider technology infrastructures, as well as adjusting risk management processes accordingly.

Covid-19 has also accelerated the shift towards online society and existence. AI will likely become an increasingly integral part of this in the short-term. Therefore, AI needs to be considered as a socio-technical issue. Other jurisdictions are already using AI as part of the national digital infrastructures which integrate government, regulation, and financial services. Similar approaches and best practices may be beneficial for the UK services sector. In the long-term, one of the members suggested, if AI adoption were to follow a path similar to that of the internet, then there would be a gradual increase in the benefits associated with AI, interspersed with periods of relative inertia.

Question 2: AI technologies are widely used in homes and workplaces, in healthcare and media, and in everyday devices from phones and tablets to cars. What are the lessons to be learned from outside the financial sector including best practices?

The members acknowledged that AI can bring significant benefits to consumers and businesses. The use of AI in homes, workplaces and other aspects of daily life are good examples of this. However, members also pointed out examples in these settings that may have gone against principles of accountability, fairness, privacy and transparency.

In order to harness the benefits and mitigate any potential risks, members suggested that financial services could look to digital-native companies in other sectors to see how they design, test, implement, and monitor AI applications. For example, and as required by ICO guidance, embedding data considerations from the start to avoid bias, including by ensuring data is truly representative of the target audience. Or using cross-functional teams and adversarial inputs to test performance against multiple criteria, rather than limited teams testing for a single output that may miss an unintended consequence. Another suggestion was for financial services firms to explore how best practices from other sectors could be integrated into AI systems. By taking such approaches, some of the members argued that firms will help build trust in AI and this is fundamental if AI is to deliver beneficial innovation.

The members also discussed how major advances, benefits and challenges in other sectors are primarily related to data collection and analysis, rather than the AI itself. From navigation apps to e-commerce platforms, these companies often use simple time-series prediction, facial recognition and recommendation algorithms. The reason they are so successful and bring so many benefits is because they have large volumes of high-quality data rather than highly complex and sophisticated algorithms. Therefore, fair and equal access to data can be more of an enabler (in terms of competition and innovation) for firms in the financial sector than the AI algorithms and models themselves.

Question 3: What are the main barriers to adoption and use of AI systems, especially for smaller or more resource-constrained firms, and how do you see these barriers developing in a post-COVID economy?

The members discussed a range of barriers for small and large firms alike. Some of the biggest barriers relate to data, including poor access to data and a lack of transparency (both in terms of providence of datasets and potential bias within them). At the model level, members discussed the challenges associated with explaining and documenting complex models, as well as ensuring appropriate governance overall. All of this applies to both small and large firms.

Additional barriers included a lack of skills and buy-in from both senior managers and business areas. Members recognised the lack of data science skills at board level and the difficulty of hiring individual data scientists with all of the relevant skills, although firms can hire teams of data scientist to ensure the right combination of skills. This is more of a challenge for small firms, which may lack the resources and therefore turn to techniques such as automated machine learning to help bridge the

gap. However, these techniques can introduce additional complexity and 'black box' risk as it is harder to document the model process.

The members also said that identifying appropriate use-cases for AI can be a barrier for firms when first starting to use AI but they should look to small-scale use cases where they already have data and buy-in from the business.

Question 4: AI systems are used in financial firms to perform a range of different functions; each of these will fit into risk management frameworks in some way. Do we need an AI-specific overlay or framework to manage AI risks across the firm?

The members observed that although risk management in financial services has improved over time, AI introduces new challenges and existing frameworks (such as model risk management) may need to evolve to cope with them.

One of the new challenges identified by members is that AI models are dynamic, meaning they continuously learn from live data and their outputs can change. This is a particular issue when the underlying data inputs change (data drift) or statistical properties of the data change (concept drift). Members noted how Covid-19 has changed consumer behaviour, which in turn has led to data drift and changes in model performance as the new behaviour is not included in the training data. Therefore more continuous monitoring and validation (and potentially documentation) is required to mitigate this risk, compared to static validation and testing methods.

Another challenge is that AI models can introduce non-financial risk that is less well understood, such as those associated with data privacy and protection, cyber security and transfer learning. The members also noted that firms' use of ever larger data sets could lead to unexpected outcomes, as well as challenges around the use of multiple internal data sets that were previously kept in silos. Lastly, the members highlighted how the speed and scale of automation can potentially lead to risks.

There was no consensus on the need for an AI-specific overlay or framework. However, there was broad agreement that risk management frameworks need to be proportionate to the complexity of the AI models and the business applications or processes they are applied to. There was some discussion regarding the merits of principles-based approaches but also the need for more granular controls.

Question 5: Accountability and informed decision-making are themes that come up often in discussions of AI and governance. How should financial firms reconfigure or adapt existing governance structures to accommodate AI?

Lots of AI-governance principles have been produced around the world, both for financial services and other sectors. However, the members said these can be difficult to execute without more appropriate and practical internal governance frameworks that provide minimum standards or best practice guidelines. Members also said that further clarification on how best to apply principles would be welcome.

The members acknowledged that existing governance structures provide a good starting point for AI models and systems. This is because AI models will interact with

other risk and governance functions, such as operational resilience, data and model risk management. Similarly, many of the risks associated with AI-models apply to more traditional models and existing governance structures deal with those issues.

Where firms do adapt existing governance structures, the members also emphasised the need for clear definition of roles and accountability, from individuals responsible for the design and development of model, to the business areas that use models, up to senior managers and board members. The members also raised the need to incorporate customer outcomes and the impact on customers into AI governance frameworks.

Question 6: Firms often note the fragmented regulatory landscape with respect to AI. There is a lot of existing regulation and policy that may apply to AI, as well as published AI principles, guidance, and best practice. What is your top ask of policy-makers and where would additional clarification, guidance, or policy interventions be useful?

There was broad consensus among the members regarding the value of AI principles that have been produced by regulators and governments around the world, most of which address a similar set of topics. However, the difficulty for firms is in understanding how to apply the various requirements, in translating them into effective internal practices and in helping those developing AI systems and those accountable for them within firms to make decisions. Thereafter, the discussion focused on the merits of more practical guidance.

For example, data protection is a common theme for most of the global AI principles but there is less guidance about how to apply these principles in practice and what regulators expect from firms. One member said that providing 'real life' examples of what could have been done to prevent negative outcomes and risks that have occurred would be extremely helpful. The members highlighted that relevant standards and guidance have been produced in other sectors that may be applicable to financial services, such as software development.

The members also commented on the lack of any standardised definitions for key terms, such as ethics, fairness and transparency, and the lack of any standardised measurements for such principles. Both of these can make it more difficult for firms to implement high-level principles.

Lastly, members agreed on the need to avoid regulatory fragmentation, both domestically and internationally, and between different sectors. Harmonising regulation at these levels would help ensure accountability and mitigate risks without stifling innovation.

Item 3. Work plan and next steps

The co-chairs and members agreed on the proposed work plan and next steps, which included the following items:

- Second Forum on Data will take place in Q1 2021.
- Third Forum on Model Risk Management in Q2 2021.

- Fourth Forum on Governance in Q3 2021.
- Each Forum meeting will be followed by a workshop two weeks later.
- Preparatory material will be circulated in advance of all meetings and workshops, with members providing input

The moderator noted that the next AIPPF meeting on Data would likely be held digitally, given the current circumstances regarding Covid-19.

Item 4. Closing remarks by co-chairs and moderator

The co-chairs and moderator closed the session and thanked the AIPPF members and observers for their contributions. They emphasised how fundamental the issues discussed are for firms and regulators alike and welcomed the opportunity to collaborate closely over the next 12 months.