

AI model risk: What the current model risk management framework can teach us about managing the risks of AI models

Received (in revised form): 22nd April, 2022

Catarina Souza

Head of Model Development and Review Division, Bank of England, UK

Catarina Souza is Head of Model Development and Review Division at the Bank of England, Prudential Regulation Authority (PRA). Prior to joining the PRA, Catarina was a Senior Expert for Model Risk Management at ING, where she led the implementation of the model risk management framework globally for the bank, with a significant focus on the adoption of Artificial Intelligence (AI) technology. Catarina has over 16 years of working experience, including leading multicultural international teams successfully delivering cross-border projects in the areas of credit risk and model risk management. Prior to joining ING, Catarina worked with S&P Global Market Intelligence, overseeing numerous clients across EMEA and the Americas, managing the development, implementation, maintenance and validation of credit risk models.

ABSTRACT

The rapid adoption of Artificial Intelligence (AI) among financial institutions in recent years creates several opportunities, but also presents significant risks that require adequate risk management. Despite advances in recent years, AI regulation remains fragmented. This creates a challenge for financial institutions when looking for guidance on how to address the emerging risks presented by the use of AI. Given the complexity and speed of revision, AI models tend to propagate and amplify existing model risk. This grants them the potential to be more harmful, and raises important model ethics concerns. This paper discusses how the existing model risk management framework can offer important lessons for financial institutions on how to tackle these emerging risks. Additionally, the paper explores possible enhancements to the model

risk management framework in order to address the unique challenges posed by AI models. These include adapting governance and policies, including model ethics considerations; enhancing model risk identification and classification; and updating model life cycles, with an emphasis on data management, model development, validation and monitoring. While the author agrees that AI risks are diverse in nature, the focus of the paper is on the risks derived from the use and development of AI models.

Keywords: model risk, SR 11-7, artificial intelligence (AI), machine learning (ML), model risk management, model life cycle, model ethics

INTRODUCTION

The use of Artificial Intelligence (AI) among financial institutions has increased dramatically in recent years. AI technology, particularly Machine Learning (ML) and Natural Language Processing (NLP) capabilities — both a subset of AI —, are now commonly used by financial institutions¹ for different applications. These applications include credit rating, credit decisioning, fraud detection and prevention, marketing, chatbots, suspicious activity monitoring and customer due diligence, to name a few.^{2,3}

The use of AI/ML/NLP (the author acknowledges that the terms are different, but for the purpose of this paper these terms will be used interchangeably⁴) brings significant benefits, namely improving back-office operations and customer experience, as well as enabling the use of a larger set of data,



Catarina Souza

Bank of England,
20 Moorgate,
London, EC2R 6DA,
UK
Tel: +44 20 3461 5822;
E-mail: Catarina.Souza@
bankofengland.co.uk

Journal of Financial Compliance
Vol. 6, No. 2 2022, pp. 103–112
Henry Stewart Publications
2398-8053 (2022)

including non-traditional data, which potentially can lead to better decisions.

While the adoption of this new technology by financial institutions is estimated to generate benefits of around US\$250 billion in the banking industry,⁵ it also poses emerging risks that require adequate management. There are several cases highlighting these risks. Examples include raising discrimination/bias concerns — the ‘famous’ case of the Apple/Goldman Sachs⁶ payment card or Amazon’s hiring algorithm.⁷

In light of the perceived failures exhibited by AI models, regulators have started to issue guidance surrounding the use of AI, aimed at addressing some of the emerging risks and unintended uses. The European Commission (EC) proposal for AI regulation or the Model AI Governance Framework proposed by Singapore,⁸ are relevant examples of attempts by regulators to address the evolving risks from the use of AI.

Despite recent regulatory advancements, regulatory guidance remains fragmented, which can be challenging for financial institutions. At the same time, the fast pace of adoption of AI models requires financial institutions to understand and manage fully the risks of these models. The lack of a comprehensive regulatory guidance means financial institutions are left to answer the following questions: How to deal with AI model risk? How to keep the bank safe and compliant, while reaping the benefits AI technology can generate?

Fortunately, banks do not need to start from ‘zero’ when answering these questions. Existing model risk management practices, already adopted by financial institutions, can offer important lessons on how to deal with the emerging risks of AI models. And although the new technology brings specific risks that need to be accounted for, banks can build on their model risk management framework to find comprehensive ways to address them.

This paper aims to offer a possible approach to deal with AI models risk, learning from and leveraging on existing model risk management practices. And while the author agrees that AI risks are diverse in nature, the focus of the proposed approach is on the risks derived from the use and development of AI models.

MODEL RISK MANAGEMENT FRAMEWORK AS A STARTING POINT FOR MANAGING AI MODEL RISK

The Supervisory Guidance on Model Risk Management (OCC 2011-12/SR11-7) remains a reference when it comes to managing model risk, which is, ‘[. . .] the potential for adverse consequences from decisions based on incorrect or misused model outputs and reports. Model risk can lead to financial loss, poor business and strategic decision making, or damage to a bank’s reputation’.⁹

According to this document, the following key components are required for the successful model risk management:

- Governance, which includes the establishment of a comprehensive model inventory and documentation;
- Model development, implementation and use, which includes important considerations for model testing; and
- Model validation.

For the past decade, financial institutions have developed their internal model risk management frameworks considering the components above. A proposed adoption is illustrated and described in Figure 1.

The model risk management framework aims to proactively identify model risks to enable their control and mitigation, considering the bank’s tolerance for model risk. The level of model risk tolerance — also known as model risk appetite — must be set

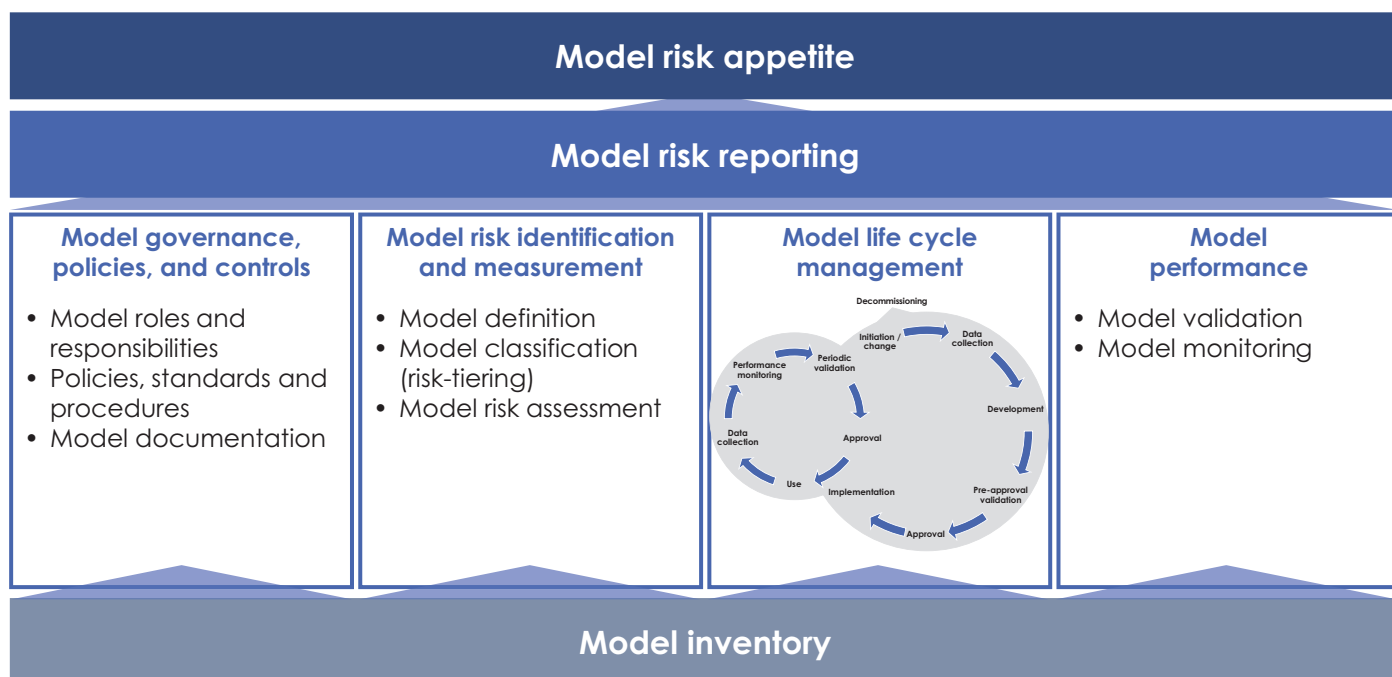


Figure 1 The model risk management framework

by the bank's board and is cascaded and drives all model risk management activities.

These activities include the establishment of the governance and policies for model risk management, the identification and measurement of model risk (including the definition of what is a model), and the model life cycle management (comprising model performance).

The foundation of all model risk management activities is the establishment of a comprehensive model inventory, with complete and up-to-date information for each model across the model life cycle. The inventory enables an effective and efficient reporting, which is essential to ensure senior management oversight and accountability for model risk management activities.

The question is: how can financial institutions leverage on the model risk management activities mentioned above to proactively address emerging AI model risks? Before attempting to address this question, it is

useful to recap the most common risk types for models, particularly in the context of AI models.

Model risk types and what changes with AI models

The literature has thus far identified several risks types¹⁰ for models. These can be summarised as follows:

- **Input risk:** Caused by inadequate or insufficient data/data quality, data assumptions, modifications, feeder models.
- **Design risk:** Caused by inadequate, incomplete or mistaken user requirements, modelling choices, assumptions, biased output.
- **Adaptability risk:** Caused by a mismatch between the model's ability to be timely adapted, eg by way of parameter setting, and the level of change in the model environment.

- **Implementation risk:** Caused by inadequate implementation choices, technical design or execution of the implementation.
- **Use risk:** Caused by misuse, misinterpretation of the output, misunderstanding of the model.
- **Transferability risk:** Caused by the lacking ability to transfer model knowledge, eg design, implementation, limitations, to a third party such as users, developers, validator, regulators.
- **Regulatory risk:** Caused by non-compliance with applicable regulation, law and internal banks' policies.

A less discussed model risk type that has been identified and gained prominence recently is the '*inventory risk*'. This can be described as 'risks associated with incomplete or inaccurate model inventories, use of unvalidated models or models that have been retired or failed validation, opacity of model usage, orphan models in inventory'.¹¹

As with 'traditional' models, all the above risks apply to AI models. However, the scale of use of AI models, as well as their speed of revision and complexity, means that these models tend to propagate and amplify model risk. This grants AI models the potential to be more harmful.

Therefore, AI models are generally associated to give rise to the following emerging model risk:

- **Potential to proliferate or augment bias (scale of bias):** models may produce unfair outcomes to individuals or some groups of people. The data might include prohibited information or the data can be used as proxy or is correlated with information that may capture protected attributes that were not included during data collection. Examples of protected attributes include racial or ethnic origin, disability, sexual orientation.
- **Higher complexity can lead to 'black box' models:** the higher complexity

generally associated with AI models — they learn autonomously, at great speed, relying on empirical data as opposed to fundamental theory — can lead to lower transparency, explainability and interpretability of results. Patterns identified by models may be non-intuitive and it may not be possible to explain model results.

- **Information leakage and security:** models may mistakenly conceal or reveal information that they use and produce (personal or critical data). The data used by models and respective outcomes is open to manipulation and utilised for unintended purposes.

When developing or using models, financial institutions are required to comply with the relevant legal and regulatory requirements. These include, among others, compliance with data protection laws — eg General Data Protection Regulation (GDPR)¹² — as well as laws preventing unlawful discrimination.

Given the particular risks posed by AI models, compliance with such laws and regulations can be further challenging. The 'black box' nature of these models, together with the greater complexity, 'can affect the ability to understand and predict the behaviour of models, making it more difficult to identify and fix problems that may undermine compliance'. Additionally, '[. . .] the effects of compliance violations can be more severe' given these models' capacity to propagate and amplify model risk.¹³

Current regulatory guidance continues to evolve in an effort to address the specific risks posed by AI models, such as The European Commission (EC) proposal for AI regulation.

Having clarified what is AI model risk, the remaining sections of this paper discuss what can be learnt from existing model risk management practices and what enhancements are required to the model risk

management framework to address the model risk as described above.

LEVERAGING THE EXISTING MODEL RISK MANAGEMENT FRAMEWORK WHEN ADDRESSING AI MODEL RISK

Model governance, including oversight structures, remains essential for managing AI model risk. Enhancements should include AI model ethics considerations

A comprehensive model governance promotes a clear definition of roles and responsibilities for managing model risk across the model life cycle. The role of model owner remains at the core of the model governance for AI models. Model owners are ultimately accountable for their models and should explain and understand the associated risks. Given the diverse background of many of the stakeholders involved with AI, who may not be familiar with model risk (compared to stakeholders in risk or finance areas of the bank), it is important to establish a clear governance for managing AI. Banks should devote efforts to increase model risk management awareness, particularly clarifying stakeholder's responsibilities across the model life cycle.

To ensure senior management oversight, model risk committee(s) are established, typically comprising members representing compliance, legal, data and AI functions within the bank. Banks should consider expanding existing model risk committee(s) to ensure appropriate technical knowledge for AI models is available. This will help ensure that AI models are not rejected due to lack of understanding of their capabilities and risks.

In parallel, it is useful to develop an oversight group of experts within the bank dedicated to what is frequently referred to as 'AI model ethics', that is, the ethical implications of using AI models by the bank. This oversight group should include a diverse group of stakeholders (eg compliance, legal,

risk), and can have the role of advising senior management on AI decisions.

When enhancing the governance and policies to include AI model risk, it is useful to follow a two-stage approach by developing overarching (principal) standards that are applicable across the bank and additional operational-level standards that can be more frequently adapted to meet AI model developments.

The importance of model documentation

Model documentation remains a key mitigant for model transferability risks. Sound model documentation enables a third party to understand how the model operates, its limitations and key assumptions, and allows the replication of the model by said third party.

AI models are typically more complex and less transparent due to their 'black box' nature. Special attention is therefore required to ensure that all choices during model development — 'theory, choice of sample design [. . .], numerical routines, selection of inputs and exclusions, estimation, and implementation in different information systems'¹⁴ — are clearly stated in the documentation.

Banks should follow a risk-based approach when establishing documentation requirements, particularly when determining the effort regarding model explainability. The higher the model importance for the bank, the stricter must the requirements be to ensure the model is explainable and understood by stakeholders.

Model documentation should clearly state the AI model's limitations to mitigate the risk of misuse by business users, a risk that tends to be exacerbated by the lack of model interpretability and transparency.

Finally, it is crucial that model documentation supports the dynamic nature of AI models. Banks can explore the benefits of advances in IT and electronic documentation, for example, with the automation of documentation. This will bring agility to

the model life cycle, particularly model development.

Model inventories should support model risk management reporting and oversight

Model inventories should facilitate a holistic understanding of the use of models across the bank and enable a comprehensive overview of the models throughout the model life cycle — ie from inception to decommission.

According to a survey recently conducted by Deloitte, ‘small and medium banks store on average approximately 30 data fields on the models, where large banks on average 58 data fields in their model inventory’. Although most banks report storing information related to the model owner, model materiality and model quality, banks report less availability of information related to model dependencies or limitations.¹⁵

Banks should therefore enhance their inventories with information to enable the appropriate understanding and reporting of AI model risk. This includes having ‘an inventory of AI uses’ across the bank, together with the ‘identification of the level of risk associated with each AI use’.¹⁶

It is useful to include in the inventory information regarding the model technique used by model development. If the model is deemed to be an AI model, the inventory should highlight the specific model risk associated — eg model ethics.

Although these aspects seem intuitive, their practical implementation is challenging. In the first instance, because it requires banks to define what AI models are, which in itself is not a straightforward concept.

Considerations for model definition and the importance of classification/risk tiering to determine the intensity of model risk management activities

The question of whether a model candidate qualifies as a model or not, is one that is

extensively debated by the industry, mostly because there is no common accepted definition of what constitutes a model.

As recommended by the Prudential Regulation Authority (PRA),¹⁷ financial institutions ‘should establish their own definition of a model’. The above-mentioned survey, recently conducted by Deloitte,¹⁸ concluded that the definition from the SR 11-7 document remains the most widely used regulatory model definition.

[...] the term model refers to a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates. A model consists of three components: an information input component, which delivers assumptions and data to the model; a processing component, which transforms inputs into estimates; and a reporting component, which translates the estimates into useful business information. [...] The definition of model also covers quantitative approaches whose inputs are partially or wholly qualitative or based on expert judgment, provided that the output is quantitative in nature.¹⁹

The above regulatory definition will, in general, apply to most AI candidate models. There are, however, instances where the outputs generated by an AI tool might not be quantitative in nature, although these will still use statistical or mathematical techniques. For example, a chatbot might suggest certain products to clients based on their preferences, or a marketing model might select different target groups for different marketing campaigns.

As the model landscape evolves to include AI tools, banks continue to build on the regulatory definition to develop internal (and enhanced) model definitions, in many cases, using decision trees or scorecards to support the assessment. According to the

same Deloitte's survey, around 53 per cent of banks use either a decision tree or a scorecard/questionnaire to assess model candidates, with 30 per cent using 'another assessment methodology than a decision tree or scorecard/questionnaire'.²⁰

Perhaps a useful approach is for banks to ask what risks such AI tools pose for the bank and whether there are controls in place to manage them. This approach will require banks to rethink the risk management for 'non-models' and the respective impact for the model inventory. It is beyond the scope of this paper to elaborate on this aspect, but it is relevant for banks to shift the discussion from the model definition towards adequately assessing and managing the emerging risks from AI technology.

This is in line with regulatory expectations: 'Regardless of how AI is classified (ie as a model or not a model), the associated risk management should be commensurate with the level of risk of the function that the AI supports'.²¹

Classification or risk-tiering within model risk management is used to determine the intensity of model risk management activities for different models (or groups of models). Financial institutions typically consider criteria such as materiality/exposure, model use or reputation risks to determine the model classification.

These aspects remain relevant for AI models, but their application for AI models requires enhancements. For example, reputational risks should be expanded to take into account non-compliance with GDPR regulations or potential discriminatory consequences from model biases. Additionally, determining materiality for an AI model might involve considering new aspects, such as number of interactions performed by the model during a certain time period or the number of customers affected by the model.

Finally, model complexity should be included as an important criterion when determining the model classification and

can be a very important driver when determining their risk tiering vis-à-vis other criteria.

Model life cycle: Considerations for data management, model development, model validation and monitoring

Data management

Sound data governance is an important tool in mitigating input risk, as it contributes to ensure data availability, quality, suitability, as well as whether assumptions and limitations are clear.

AI models use large amounts of data, both structured and unstructured, to drive the parameters/features selection, which increases the importance of data validation and controls. Given the volume and dynamic nature of the data used by AI models, assessing data integrity should entail confirmation that the data is suitable for the models' purpose and use in the light of existing legal requirements, privacy and security aspects. Additionally, special attention should be devoted to ensure that the data is labelled correctly over time and database, is monitored frequently, and an assessment of biases embedded or inherent in the data is performed.

In general, data governance is beyond the scope of the model risk management framework. That said, given the interdependencies, it is important that the requirements are clear and understood by all stakeholders across the model life cycle.

Model development

The model risk management framework should be enhanced by setting specific AI model development standards. These model development standards should address model explainability and transparency, as well as fairness — eg establishing a clear taxonomy of protected attributes (eg racial or ethnic origin).

Validation and monitoring for AI models

The use of AI contributes, to a great extent, to an exponential increase of the number of models in the inventory, with circa 70 per cent of banks in Deloitte's survey indicating that machine learning techniques are used in their model development.²² This 'model creep'²³ phenomenon is compounded by the velocity at which these models are modified or updated.

Validation continues to be an important approach to mitigate AI model risk. Sound AI risk management includes an 'effective process to validate that AI use provides sound, fair, and unbiased results'.²⁴ However, given the dynamic nature of AI models, and since financial institutions cannot rely on unlimited resources, the application of the existing static validation framework can be challenging.

According to the SR 11-7 'The range and rigor of validation activities conducted prior to first use of a model should be in line with the potential risk presented by use of the model'.²⁵ Banks should therefore continue to follow a risk-based approach when determining the extend of pre-approval validations for AI models (before initial use). Moreover, 'If the bank has not fully validated models before implementation, examiners should assess the bank's compensating controls and other measures to mitigate risks'.²⁶

In parallel, banks must invest in automating, where possible, the validation activities to increase efficiency. Advances in cloud computing offer important gains for validation and enable the Second Model Line of Defence (MLoD) to match the pace of the First (model development).

Once the model is in use, banks continue to attest that the model is fit for its purpose until the model is decommissioned. According to the SR 11-7, 'Banks should conduct a periodic review — at least annually but more frequently if warranted — of each model'.²⁷

Since AI models are continuously learning from data, which grants them a dynamic nature, they will typically require more frequent reviews. In many cases, these models require continuous monitoring to ensure they remain fit for purpose. In fact, when it comes to AI models, '[. . .] monitoring is the new validation'.²⁸

Minimum standards for model monitoring should be enhanced with key performance metrics and thresholds, to understand model performance in real time and gain trust that the AI model performs as intended. There are several practical considerations when developing the monitoring standards. For example, it is important to define material model changes and when these may warrant a validation (eg in which cases does new data imply a re-validation of an AI model?). It is also important to include in the monitoring activities not only an assessment of the technical performance of the model, but also monitor regulatory developments and business output uses from the model.

From a content point of view, assessing the AI model's conceptual soundness remains a key challenge. This assessment is described as 'assessing the quality of the model design and construction. It entails a review of documentation and empirical evidence supporting the methods used and variables selected for the model'.²⁹

As mentioned previously, it is difficult to judge model explainability and its impact on model documentation. Validators can use tools to help explain the relation between model variables, such as 'Shapley values'³⁰ and 'LIME'³¹ tools. If this is not possible, banks can explore alternative approaches to ensure trust in the model. There are useful lessons from the use of third party 'traditional' models, where the emphasis is to understand the use (outcome analysis and stability testing) and perform sensitivity analyses or benchmarking. Banks can focus on either particular elements of the AI model's

use — eg the use of human judgment or the quality of the data — or perform sensitivity analyses on the AI model variables to gain further understanding of the outputs.

One additional important consideration for model validation is assessing unintended bias or discrimination by the model. Even if AI models do not have access to protected attributes, such as racial or ethnic origin, their outputs can still have a disproportionate impact on certain groups. In addition to reviewing the data used for model development, the model output should be tested for model disparity. Pro-actively, banks should install a process for when models do not perform as expected to act swiftly to detect and take corrective action.

CONCLUSIONS

As the model landscape continues to expand and the use of AI technology increases, new questions arise regarding how financial institutions can proactively assess and address the emerging risks that accompany the use of these models.

These risks are not necessarily new, but rather enhanced by the specific characteristics of these AI models. These risks can be summarised as: (a) potential to proliferate or augment bias (scale of bias); (b) higher complexity can lead to ‘black box’ models; and (c) information leakage and security.

While the existing model risk management framework offers a useful starting point to identify, understand and mitigate AI model risk, there are important enhancements required to accommodate AI model risk.

Financial institutions should consider:

- The model governance remains essential when managing model risk. The expansion of the model landscape to AI models requires a continued focus on increasing stakeholder’s knowledge on model risk management. In parallel, oversight structures should be enhanced to ensure

adequate knowledge of AI is available, particularly for model ethics.

- Since AI models are typically more complex and less transparent due to their ‘black box’ nature, model documentation merits special attention, as it is a key mitigant to model transferability risk.
- The model inventory should be improved to continue to facilitate adequate reporting and oversight of AI model risk. Emphasis should be placed on AI model use.
- Internal model definitions may require revision, as some AI models will not meet the existing regulatory definitions. When doing so, banks should stress the importance of understanding the risks posed by these models and what controls and mitigations are in place. The proportionality principle remains crucial to managing AI model risk.
- Model standards for data management, model development, validation and monitoring should be expanded to address specific AI models risk, including model explainability, transparency and fairness. This is important to ensure there is trust in the AI model.

AUTHOR’S NOTE

This paper was written by the author prior to joining the Bank of England and while working for ING. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of ING, the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.

© Catarina Souza, 2022

NOTES AND REFERENCES

- (1) Throughout the paper, the terms financial institutions and banks are used interchangeably.

- (2) EBA Report on Big Data and Advance Analytics, January 2020, (EBA/REP/2020/01), available at https://www.eba.europa.eu/sites/default/files/document_library/Final%20Report%20on%20Big%20Data%20and%20Advanced%20Analytics.pdf (accessed 24th August, 2022).
- (3) DeNederlandscheBank (2019), 'General Principles for the Use of Artificial Intelligence in the Financial Sector', available at <https://www.dnb.nl/media/voffsrc/general-principles-for-the-use-of-artificial-intelligence-in-the-financial-sector.pdf> (accessed 24th August, 2022).
- (4) It should be noted that both ML and NPL are a subset of AI. ML includes: supervised, unsupervised and deep-learning algorithms. NLP covers content extraction, classification, machine translation, question answering and text generation algorithms.
- (5) McKinsey & Company (2019), 'Derisking Machine Learning and Artificial Intelligence', available at <https://www.mckinsey.com/business-functions/risk-and-resilience/our-insights/derisking-machine-learning-and-artificial-intelligence> (accessed 24th August, 2022).
- (6) Newberger, E. (10th November, 2019), 'Wall Street Regulator Probes Goldman over Allegations of Sexist Credit Decisions at Apple Card', CNBC, available at <https://www.cnbc.com/2019/11/10/wall-street-regulator-probes-goldman-over-allegations-of-sexist-credit-decisions-at-apple-card.html> (accessed 24th August, 2022).
- (7) Dastin, J. (11th October, 2018), 'Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women', *Reuters*, available at <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G> (accessed 24th August, 2022).
- (8) Singapore Digital (2020) 'Model Artificial Intelligence Governance Framework' (second edition), available at <https://www.pdpc.gov.sg/-/media/files/pdpc/pdf-files/resource-for-organisation/ai/smodelaigovframework2.ashx#:~:text=Singapore%20is%20proud%20to%20launch,OECD%20Expert%20Group%20on%20AI> (accessed 24th August, 2022).
- (9) Board of Governors of the Federal Reserve System Office of the Comptroller of the Currency (2011) Supervisory Guidance on Model Risk Management (OCC 2011-12/SR11-7).
- (10) Hill, J. R. (2019), 'The Top 14 Challenges for Today's Model Risk Managers: Has the Time Come to Think about Going Beyond SR 11-7?', *Journal of Risk Management in Financial Institutions*, Vol. 12, 2.
- (11) *Ibid.*
- (12) Regulation (EU) 2016/679 of the European Parliament and of the Council, available at <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679> (accessed 24th August, 2022).
- (13) Ostmann, F. and Dorobantu, C. (2021), 'AI in Financial Services'. *The Alan Turing Institute*. <https://doi.org/10.5281/zenodo.4916041>.
- (14) Comptroller's Handbook Model Risk Management, version 1.0, August 2021, available at <https://www.occ.gov/publications-and-resources/publications/comptrollers-handbook/files/model-risk-management/pub-ch-model-risk.pdf> (accessed 24th August, 2022).
- (15) Deloitte (2021), 'From Validation to Optimisation, Tackling a Growing Model Landscape EMEA Model Risk Management Survey', available at <https://www2.deloitte.com/content/dam/Deloitte/nl/Documents/financial-services/deloitte-nl-fsi-model-risk-management-survey.pdf> (accessed 24th August, 2022).
- (16) Comptroller's Handbook Model Risk Management, ref. 14 above.
- (17) Supervisory Statement SS3/18, 'Model Risk Management Principles for Stress Testing', April 2018, available at <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2018/ss318.pdf> (accessed 24th August 2022).
- (18) Deloitte, ref. 15 above.
- (19) Board of Governors of the Federal Reserve System Office of the Comptroller of the Currency, ref. 9 above.
- (20) Deloitte, ref. 15 above.
- (21) Comptroller's Handbook Model Risk Management, ref. 14 above.
- (22) Deloitte, ref. 15 above.
- (23) Hill, J. R. (2019) 'A Smarter Model Risk Management Discipline Will Follow from Building Smarter Models', *Journal of Risk Management in Financial Institutions*, Vol. 13, 1.
- (24) Comptroller's Handbook Model Risk Management, ref. 14 above.
- (25) Board of Governors of the Federal Reserve System Office of the Comptroller of the Currency, ref. 9 above.
- (26) Comptroller's Handbook Model Risk Management, ref. 14 above.
- (27) Board of Governors of the Federal Reserve System Office of the Comptroller of the Currency, ref. 9 above.
- (28) Quell, P., Bellotti, A. G., Breeden, J. L. and Martin, J. C., 'Machine Learning and Model Risk Management', Tech Report 2021-01, 8th March, 2021, Version 1.0, MRMIA.
- (29) Board of Governors of the Federal Reserve System Office of the Comptroller of the Currency, ref. 9 above.
- (30) Lundberg, S. M. and Lee, Su-In (2017), 'A Unified Approach to Interpreting Model Predictions', in (eds) Guyon, I. et al., *Advances in Neural Information Processing Systems 30* (NIPS).
- (31) Ribeiro, M. T., Singh, S., and Guestrin, C. (2016) "'Why Should I Trust You?'" Explaining the Predictions of Any Classifier', available at <https://www.kdd.org/kdd2016/papers/files/rfp0573-ribeiroA.pdf> (accessed 24th August, 2022).