KPMG

Modern Risk Management for AI Models

Re-imagining the Model Risk Management function for Artificial Intelligence / Machine Learning models

Whitepaper

Foreword

The concept of Artificial Intelligence (AI) and Machine Learning (ML) dates back to the 1950s with researchers trying to use machines to simulate human intelligence in machines. However, their use in the financial sector has been limited. With enhanced customer service and mounting cost pressures, financial institutions (FIs) have now started to look at AI as a possible solution for improving cost and operational efficiencies.

Regulators in the European Union (EU) and US have recognized the impact AI can create in financial and consumer markets. At the same time, they are also mindful of the inherent risks involved. The 'EU Artificial Intelligence Act 2021' is aimed at creating a risk-based regulatory framework around AI focused on the pyramid of criticality, with a modern, layered enforcement mechanism. In other words, a lighter legal regime would apply to AI applications with negligible risk, and applications with an unacceptable risk are banned.

As more and more banks are adopting AI/ML models in their banking applications, it is important to have an international regulatory guidance on how to handle specific risks arising from these models. Model risk guidance, SR 11-7 has not been adapted to address the specific risk from AI/ML models. There is a widespread difference in the approach banks are taking for handling risks around bias, interpretability and other challenges – some of the global banks are already validating their ML models and some have even invested in AI/ML Centers of Excellence, while for others it is still in a very nascent stage. It is imperative for banks to develop a meaningful understanding of the technology, including its existing and potential uses within their organizations, and take a firm grip on the implications of AI from a risk perspective. Through various stages of the model lifecycle, FIs would need to keep their model risk management (MRM) practices up to date to manage the risks effectively.

In this paper, we aim to discuss some of the common risks and challenges that KPMG firms have encountered across the lifecycle of AI/ML models and how a traditional MRM framework should be adapted with the aim of addressing these challenges.

Rajosik Banerjee
Partner and Head of
Financial Risk Management,
KPMG in India

Matthias Peter
Partner,
KPMG in Germany

Contents

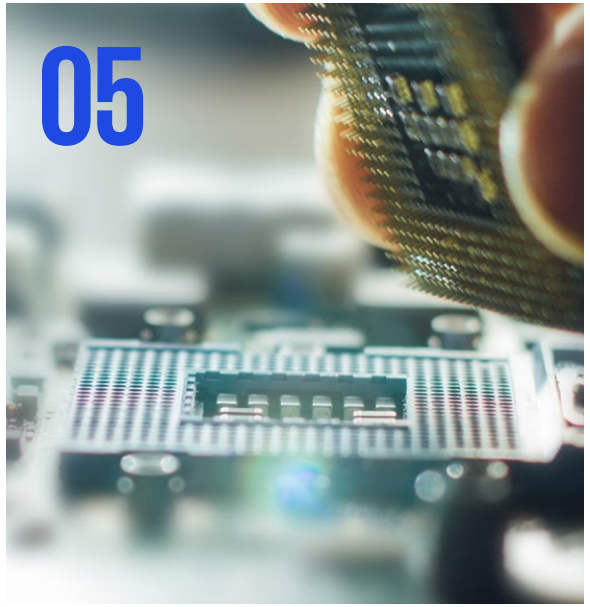

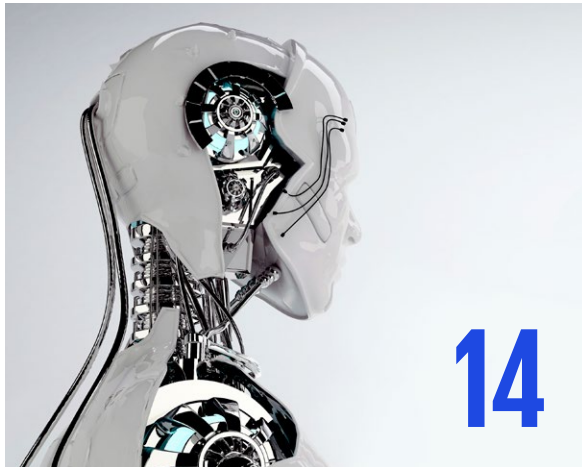

Introduction

As we enter a new era of unprecedented customer experiences, the demand for faster and higher-quality financial analysis and forecasting has increased significantly. The clear benefits of AI in finance, banking, and business analytics are easy to gauge. It is difficult to define AI, but broadly speaking it is the theory and technology around development of smart and intelligent computer systems or algorithms that are not explicitly programmed for, but "self-learn" to perform tasks which otherwise would require human intelligence. Machine Learning, Deep Learning (DL), Speech Recognition, Natural Language Processing (NLP) and visual recognition technologies all belong to this class.

Even though we use the terms AI and ML interchangeably, there are a few differences among them, but at the heart of it, the primary aim of these algorithms is to create intelligent systems. For example, ML applications like chat-bots have now enabled banks to automate time-consuming, repetitive processes offering a far more streamlined and personalized customer experience. They have also allowed banks to work more productively with large databases, unstructured information, significantly improving the quality of asset valuation, forecasting financial performance, and solving key issues around data security.

In addition to new models being implemented, the more widespread use of AI/ML applications can also lead to the replacement of manual processes or simple models with AI/ML models.

Hence the classical model definition needs to be extended affecting MRM directly as it is based on the model inventory.

Traditionally, the model inventories had been focused on statistical models for credit, market, liquidity, technology and operational risks but with the advent of the ML models, the model inventory scope has also expanded. Many banks are now inducting more ML models in their model inventory – in areas of surveillance, fraud detection, text analytics, customer service, digital marketing, trading, underwriting, customer behavior and pricing predictions. The vast amount of data consumed by these models make them represent real life problems and human behavior more aptly and enable analysis of multiple dimensions, thus offering an edge over traditional models.

However, ML models can come with their own set of risks and challenges. Often many of these models are black box in nature and hence it is difficult to ascertain if they are performing without bias. In such situations, FIs can find themselves in violation of antidiscrimination laws which not only would result in huge fines from regulators but could also significantly damage the FI's reputation.

Another important aspect in addition to mitigating bias that is essential for the use of AI/ML models is ensuring their explainability.

Specific risks in AI/ML models

Specific risks in AI/ML models

One key difference between an AI/ML and traditional model is that an ML model is expected to continuously learn from the data, identify patterns and refine its decision-making process. In other words, the quality of an AI/ML model is as good as the data which has been used to develop it. If the data quality is compromised, so will be the decisions made by the model.

Perhaps the most important topic for AI/ML practitioners is to **ensure fairness**. Real life data always has inherent biases which when used in modeling gets carried forward in the decision making (refer to Figure 1). There are multiple ways in which the models may start to behave unethically in terms of providing different opportunities, resources, information or quality of services to specific groups of people. Use of biased data would in turn also affect the future data that will get used for subsequent model training. For example, a biased credit scoring model would impact the customer selection which in turn impacts the constitution of a future portfolio and hence the future input data. This creates a loop, and the bias keeps on propagating and gets enlarged over time.

The **increased model complexity** is another key driver of the amplified risks associated with ML models. These models are built on large data sets – structured as well as unstructured – and use complex quantitative algorithms. Some of the ML algorithms like neural networks and gradient descent are opaque making them difficult to interpret. There are multiple hidden layers of decision making which impact the final output.

Results based on AI/ML algorithms need to be **explainable to all stakeholders** – i.e. customers, management and regulators – at all times. For example in credit scoring it is essential to be able to explain which criteria lead to a rejection (e.g. salary, savings) and what would be possible options to improve a customers score. Additionally, transparency and explainability are requirements to ensure the algorithm to function properly.

Another key constraint follows from the **lack of ML experts** who understand the methodology and software, the data and the outputs generated by the models, interpret them, and assess if there are any inherent biases in those outputs, as this skill set is in short supply.

A distinctive feature of AI applications is that they should **continuously be re-trained on new data**. This re-training may change essential properties of the model and the parametrization, i.e., lead in effect to a model change, so that renewed validation and adequacy checks are required. It is therefore necessary for validation to closely monitor the model in production.

To make use of the advantages of an AI/ML application, compared to conventional applications, it is in general necessary to train it on an extensive set of data. Here it is not uncommon that extensive use of personal data is made. Therefore, it is crucial to **ensure data protection, abide by privacy laws** and a proper legal basis for the use of such data in AI training.

The **technical implementation** of AI/ML applications is usually not build up from scratch but relies on program packages either from professional third parties or from non commercial open source projects. In addition to the obvious need for proper licensing this poses an additional challenge as the proper functioning of these packages needs to be tested before their initiale use and after each update.

The last challenge comes from the **lack of specific regulatory guidance around MRM** practices for AI/ML models. SR 11-7 has not been updated to address the specific challenges for ML models. Additionally, it is unclear what the requirements of future regulations will be of banks. While there is industry research and some broad framework defined (refer to the section below) banks have been left on their own to adapt their MRM framework to these models. Some banks have started incorporating additional tests for ML models, but others are still contemplating. Lack of prescriptive regulatory guidance and laws might create disparity among the practices followed by different banks.

Figure 1: Different forms of bias that can be rooted in the real world, manifest in the data and can be magnified by AI/ML algorithms and human intervention lead to potentially harmful actions.

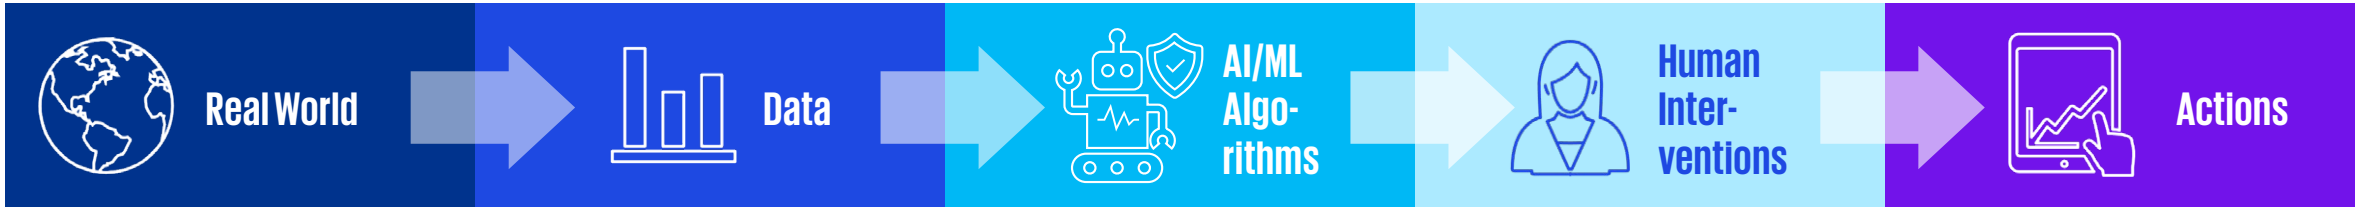

Different forms of bias

- Historical Bias
- Representation Bias
- Measurement Bias
- Temporal Bias
- Omitted Variable Bias
- Algorithmic Bias
- Evaluation Bias
- Aggregation Bias
- Popularity Bias
- Ranking Bias
- Emergent Bias
- Linking Bias
- Behavioral Bias
- Presentation Bias
- Content Production Bias
- Social Bias

Source: KPMG International, 2022

Regulatory guidance

In United States' regulation **SR 11-7** acts as a high-level guidance for practitioners to control model risk and is still largely applicable to validation of AI/ML models. In March 2021, the five largest federal financial regulators in the US released a **request for information[1] on how banks use AI**, including ML as well as on governance and risk management framework, controls and any challenges in developing, adopting and managing AI. Future updates of SR 11-7 can be expected to address AI/ML models and their treatment.

Following the call from the European Parliament, the **EU** has carved out a **'human-centric' approach** to AI that strives to ensure that human values are central to the way in which AI systems are developed, deployed, used, and monitored, by ensuring respect for fundamental rights, including those set out in the Treaties of the European Union and Charter of Fundamental Rights of the EU. The EU published its guidelines on ethics in AI in April 2019, some key requirements of which **include human oversight, robustness and safety, privacy and data governance, transparency, diversity and fairness, societal and environmental well-being and accountability**. These are also the cornerstones laid out by the Bank of International Settlements (BIS as part of its emerging regulatory expectations around use of AI in financial sector domain (refer to Figure 2)).

The Bank of England has recently published a consultation paper on MRM that explicitly refers to the risks from the use of AI/ML models.

Figure 2: Key Regulatory Expectation around AI/ML models as suggested by BIS[1]

	Principles				
	Reliability/Soundness	**Accountability**	**Transparency**	**Fairness**	**Ethics**
Regulations/Guidance	– Similar expectations as those for traditional models (e.g. model validation, defining metrics for accuracy, updating/retraining the models, ascertaining quality of data inputs) – For AI models, assessing reliability/soundness of the model outcomes viewed from the perspective of avoiding causing harm e.g. to consumers	– Similar expectations as outlines in general accountability or governance requirements, but human involvement is viewed more as a necessity – For AI models, accountability includes "external accountability" to ascertain that data subjects (i.e. prospective or existing customers) are aware of AI-driven decisions and have channels for recourse	– Similar expectations as those for traditional models, particularly as they relate to explainability and auditability – For AI models, external disclosures to data subjects is also expected (e.g. data used to make AI-driven decisions and how the data affects the decision)	– Stronger emphasis in AI models (although covered in existing regulatory standards, fairness expectations are not typically applied explicitly to traditional models) – Expectations on fairness relate to addressing or preventing biases in AI models that could lead to discriminatory outcomes however, fairness is typically not defined in FS	– Ethics expectations are broader than "fairness": and relate to ascertaining that customers will not be exploited or harmed either through bias, discrimination or other causes (e.g. AI using illegally obtained information)

Source: KPMG International, 2022

1 https://www.federalreserve.gov/newsevents/pressreleases/bcreg20210329a.htm (Download: July 6, 2022)

Regulatory guidance

The publication of the draft of the **EU's Artificial Intelligence Act** (EU AI Act) represents an important milestone in regulation of Artificial Intelligence in Europe. Besides a definition of Artificial Intelligence, it contains a classification of AI applications in terms of risk. **Four kinds of AI applications** are distinguished:

- Prohibited applications
- High-risk applications
- Applications with special requirements (e.g. bots)
- Low-risk applications

Credit-scoring by banks is explicitly named as a high-risk application. High-risk AI applications must fulfill comprehensive requirements (e.g., a risk management system, data and its governance, transparency & information). Prior to the introduction of such an application, it is necessary to produce documentation and proof of compliance with these requirements by means of a so-called **Conformity Assessment**. The Conformity Assessment must be updated in case of changes. In addition, if substantial errors occur during operation, the responsible authority must be notified.

To ensure compliance of all AI applications with the requirements of the EU AI Act, a bank **needs to establish processes** which are quite like MRM Frameworks, or which can be integrated into them. These are:

- Identification/integration of AI methods into the model inventory
- Assessment of the materiality/risk of the AI models
- Conformity Assessment/Validation and adequacy checks
- Monitoring of model performance

Hence, integration of AI into the Model-Risk framework can lead to useful **synergies** with respect to ensuring regulatory compliance.

Another important regulatory publication which gives hints about future requirements is **EBA's Consultation Paper on "machine learning for internal ratings-based models",** which deals with the use of ML models in internal risk models. Furthermore, European Insurance and Occupational Pensions Authority (EIOPA) published an interesting report: In "Artificial Intelligence Governance Principles", it formulates comprehensive requirements for AI applications, including risk assessment for AI applications and principles for "Fair Machine Learning".

With a few exceptions like in the European Union (EU) where a legislative proposal to harmonize the rules for AI already exists, **most frameworks are still in their early stages of development** and range from application of existing principles-based corporate governance requirements in an AI context to practical non-binding supervisory guides on how to manage AI governance risks.

An **overview of current developments** can be found in a 2021 paper by the Bank of International Settlements (BIS) Financial Stability Institute (FSI)[2].

2 Financial Stability Insight (FSI) Insights on policy implementation, August 2021 "Humans keeping AI in check – emerging regulatory expectations in the financial sector", Bank of International Settlements

Implications for the MRM framework

Traditional MRM processes are often not capable to address **specific risks of AI/ML models** (refer to Figure 3). There are also specific challenges around feature selection, hyper parameter tuning (i.e. technical parameters intrinsic to the AI/ML model, e.g. the number of layers in a neural network), data biases that might pose additional challenges for validators (refer to Figure 4).

During development, it needs to be ensured that the models that are consistent with the company's values and risk appetite. Due to the wide range of possible uses it might be necessary to expand the **model tiering** so that risks specific to AI/ML, e.g. on reputation or social impact, are accounted for properly. This might result in the need for new measurement approaches for model risk. Tools such as model explanation, bias detection, and performance monitoring are built in so that there is a constant and consistent oversight. This should be embedded within AI development activities right from the start. The standards, testing, and controls need to be embedded into various stages of the **model's life cycle**, from development to deployment and use, unlike for traditional models where risk managers come in usually only towards the end of the development process. The most common approach should be to keep these models under a **"constant monitoring process"** over and above what is done for the traditional models.

With respect to fair AI/ML and bias mitigation FIs need to **define what fair means** for them, as that is a prerequisite to test models for fairness. Additionally, the organization needs to establish a body responsible for validating fairness in MRM taking into account that new skills, methods and tools are needed.

Figure 3: Challenges of using traditional MRM processes for ML model validation

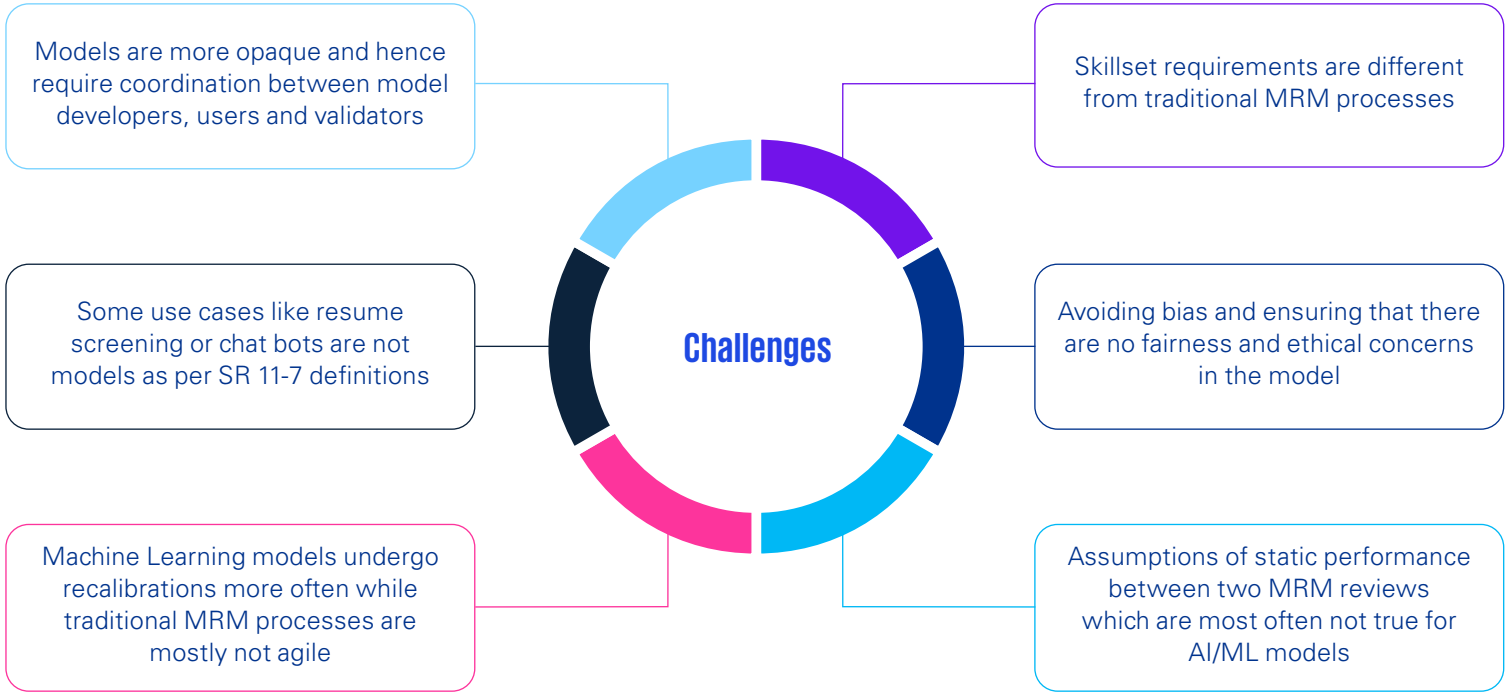

Source: KPMG International, 2022

Figure 4: Key challenges/potential risks in specific AI/ML algorithms

Decision Trees

- Variable selection is often not intuitive
- Striking the appropriate balance between prediction accuracy (less trees) and interpretability (more trees)
- Inappropriate selection of optimization algorithm – ID3, C4.5, CART etc.
- Model might be overfitted

Random Forest

- Model process is a black box and not interpretable
- Variable importance measure may be misleading in presence of categorical variables or numerical values that vary in the scale of measurement
- Selection of hyper parameters like tree size, number of trees is critical
- Long run time leading to lag in real time prediction

ADA Boost

- As a sequential algorithm, based on re-weighting of misclassified samples from the previous weak learner, the model can be prohibitively slow if the data size (n) and/or the number of features (d) are too large
- Might overfit in the presence of data with outliers
- Variable selection process is a black box

Gradient Boosting

- Inappropriate selection of hyper parameters leads to deterioration of model performance
- Extremely sensitive to missing data, data outliers
- Trees are built sequentially leading to longer run time
- Slow if data size or no. of features are higher

K Nearest Neighbor

- Create an abstraction from specific instances and hence lacks clarity and interpretability
- Non standardized independent variables lead to errors in the distance computation

K-means Clustering

- Optimal selection of number of clusters is a challenge
- Automatically determined clusters might have uneven data distribution
- Sensitive to input data like outliers and missing values
- Large number of features may lead to higher number of iterations which can generate high computational load on the network.

Quadratic Discriminant Analysis

- Not able to describe effects of higher than quadratic order
- Expensive use of matrix operations
- Cannot be used as a dimensional reduction technique

Anomaly Detection Algorithm

- Algorithm does not reveal the cause of the anomaly; may therefore not be intuitive if dealing with large data sets
- May not produce optimal results if there are multiple anomalies or if anomalies are homogeneous in nature
- Covariance matrix may be non-invertible for smaller data sets

Neural Networks

- Inappropriate selection of hyper parameters would lead to non-optimal results, higher run time, model overfitting
- Sensitive to noise in the data
- Lack of interpretability of the outputs, i.e., Black Box model

Recurrent Neural Networks

- Significant computational power and time is required which increases cost
- Structure selection needs to be in line with business requirement; incorrect structure selection would result in model underperformance or significantly higher run time

Source: KPMG International, 2022

Re-imagining the MRM framework for AI/ML models

MRM Enhancement for AI/ML Models – seven key pillars

MRM for AI/ML models can be integrated into existing MRM frameworks. This is advantageous as synergies arise from using proven processes, methods and IT tools.

In addition, the integration helps with fulfilling regulatory requirements, e.g. those from the EU AI Act. However, the integration requires small changes in order to take AI/ML model's specific risks and challenges into account. The most significant changes can be summarized in seven key pillars and include aspects like model definition, risk appetite or AI/ML specific tests.

1. Establish a definition of AI/ML models
Banks would establish an 'enterprise-wide definition' of what AI/ML models comprise, over and above the traditional definition of a model. Accordingly, the model inventory needs to be expanded to include these models.

2. Updating the model tiering definition
The model tiering parameters specifically the criteria around materiality, criticality and uncertainty might need to be improved upon to correctly identify the inherent risks of the AI/ML model. Here, a new approach is necessary, as for example the 'risk of harming the customer', such as the danger of discrimination must be addressed. The level of risk can be assessed in terms of likelihood and severity of the harm, or a combination of the two.

3. Establish an appropriate risk appetite
Traditional risk appetite statements will not work for ML models. There is a lack of regulatory guidance in this area and hence banks need to leverage their peer networks to gather industry intelligence and establish the first draft of a risk appetite statement and associated thresholds. This is an evolutionary process and banks would need to continually update this as understanding grows in this area.

4. Identify accountability
Given that multiple independent risk management functions will be involved – MRM, Compliance, Data Management and Controls, Operational Risk Management (ORM) teams – it is imperative that there is a cross functional governance framework established with clear definitions of roles and accountabilities.

5. Invest in skill enhancements
Banks would need to develop the skill set inhouse or bring in external experts. External Subject Matter Experts (SME) can help them benchmark themselves with peer banks around risk management, controls and governance framework enhancements as well as leading edge model development and validation techniques for AI/ML.

6. Enhance the compensatory control framework
The existing risk and control frameworks including MRM, data management (including privacy), compliance and operational risk management (IT risk, information security, third party, cyber) do not explicitly address risks as envisaged in the AI/ML risks and thus need to be enhanced by:

- **Designing additional compensatory controls** around benchmarking, feature selection, bias elimination, data point inspections and others to account for lack of transparency
- **Enhancing existing data management framework** to assess the scope of data sources (specifically unstructured and third-party data used in the model development), improve data quality programs to profile inbound data, embed data privacy requirements and strengthen data monitoring processes. Developers and validators need to put special emphasis on the suitability of underlying data and associated risks from data sourcing, data filtration, feature engineering and data bias/representativeness while developing the models and validating their use.
- **Building control frameworks around compliance and operational risks,** especially for consumer applications, third-party assessments, technology risks, etc. to address risks related to conduct, fair lending, data privacy and underlying technology infrastructure.
- **Conducting enterprise-wide training programs** to train all the stakeholders including the senior management on key aspects of AI/ML, including applications, ecosystems, risks, and controls, such that they can gauge the risks better and are able to challenge during the model approval process

7. Develop additional tests and procedures for AI/ML models
There are key elements which need to be specifically tested during the model's life cycle, including e.g. during design, implementation, operation and validation **(refer to the following pages for details)**.

Additional Testing for ML models – the key elements

As shown in Figure 4 "Key challenges/risks in specific ML algorithms", AI/ML algorithms differ from classical models, which requires new approaches for the entire model life cycle, including e.g. design, implementation, operation and validation. The new validation approaches address a wide range of elements.

These elements mainly cover the input data, selection of parameters, model calibration and improvement in interpretability and bias elimination, developing an ongoing monitoring framework, model implementation and designing challenger models as an effective alternate.

Input data

- The choice of data source, data selection, cleansing, mining, and transformation process needs to be examined
- It needs to be verified that the data distribution of each feature matches the expectations, especially for algorithms like AdaBoost or Support Vector Machines that can be sensitive to outliers and deviations from the expected
- Additional testing for various data biases – labelling bias, exclusion bias, measurement and design bias – need to be performed. Scores such as KL Divergence and Wasserstein may be used for capturing data deviations
- If there are any imbalances in the data, it needs to be adjusted before use for subsequent processing. Sensitivity analysis of imbalanced data to be covered as part of validation testing

Feature engineering

- Feature engineering is the process of finding the relevant features (i.e. input variables), in particular identifying those that have large influence on the model output
- Banks need to determine the level of support required to establish the conceptual soundness of each feature which can vary with the model use – e.g., a credit decision model might require that every individual feature in the model be assessed while for lower-risk models, banks might choose to review the feature-engineering process only for data transformation and feature exclusion
- Sensitivity analysis of output to changes in the inputs is an important step. Another approach can be to insist that developers also have a challenger model using alternate traditional algorithms to benchmark performance

Hyper parameter selection

- Given that an ML model's performance and stability are highly dependent on the selection of the hyper parameters, validators need to ensure no overfitting or underfitting occurs
- Hyper parameters can be set for example using Spearmint (either with Gaussian processes or Hyperopt) by means of tree-based estimators
- It needs to be ensured that the selection maps the entire parameter space, and the reasonable range is chosen
- The calibrated equalized odds postprocessing technique may be used as an alternative to optimize over calibrated classifier score outputs to find probabilities with which to change output labels with an equalized odds objective

Interpretability

- Mostly handled by bank wide policy based on the risk appetite which determines the extent to which the results need to be interpretable – whether it should hold all ML models to the same standard of interpretability or differentiate based on the model's risk
- For algorithms that incorporate automated feature selection (e.g., Support Vector Machines), it is important to assess if the model integrity is undermined by such lack of interpretability
- Explainable AI (XAI) programs are essential if banks need to effectively manage an emerging generation of artificially intelligent models
- Methodologies such as SHAP (Shapley Additive Explanations) or LIME (local interpretable model-agnostic explanations) that approximate any black box machine learning model with a local, interpretable model can be leveraged

Bias elimination

- Embedding fairness in the modeling process is a key requirement
- Exclusion or label bias can be eliminated by ensuring SMEs opine on excluded features
- Measurement bias needs to be reduced by checking for outliers and computing their degree of influence on outcome variables using metrics like 'Cook's Distance' or 'Mahalanobis Distance'
- Reweighing based on frequency counts can be another technique used to reduce bias in the data. Classification with 'Fairness Constraints', 'Prejudice Remover Regularizer' and 'Adversarial debiasing' can be used to reduce bias during in-processing step
- Bias mitigation techniques should be used by the developer in data pre-processing as well as the in-processing steps and it should be checked that the technique has effectively dealt with the bias in the 'Through the Door (TTD)' population
- Design bias can be eliminated through down sampling or oversampling methods (using algorithms like SMOTE in Python)
- One possible measure to evaluate fairness in AI/ML models is 'disparate impact' that compares proportion of positive outcomes received by two groups – privileged and unprivileged. Typical industry standards would be 80%. The validation team can set acceptable threshold limits depending on the model type

Dynamic calibration of models

- Banks need to decide on a case-by-case basis when to allow a dynamic recalibration of the model – without appropriate control, as otherwise, short term patterns in the data may affect the model's long-term performance
- Once the policy is set, validators need to ensure that the dynamic calibration is in line with the intended use of the model and the monitoring plan is in sync with the periodicity of the dynamic recalibration
- Any threshold breach of a Key Model Performance Indicator (KMPI) which indicates a material shift in the model's performance needs to trigger a review/revalidation process

Implementation

- It needs to be ensured that volume of data does not impact the operational stability of the system
- The computational performance in speed, capacity, and efficiency, through metrics such as latency, throughput and RAM usage need to be determined, especially for models that use recursive algorithms for optimization, receive real-time requests and generate instant outputs
- In the validation also methods and implementations from third parties or other external resources need to be taken into account
- In addition to the technical requirements, the general appropriateness of the models is to be checked

Ongoing monitoring

- The performance of ML models may change over time due to 'data drift' and/or 'feature drift' or 'model drift'
- It is important to ensure the monitoring plan is comprehensive and robust
- Basic statistical techniques (mean, standard deviation, range, quantiles) along with distance and divergence measures (KL statistic, Hellinger distance), chi-square and entropy, outlier detection as well as model performance indicators (Gini, ROC, MSE, MAD, MAPE) offer a large variety of key metrics which can be included in the monitoring plan along with respective thresholds
- The model change policy needs to be adhered to, in particular for pillar I models

Conclusion

AI/ML models can offer significant added value in the delivery of financial services – personalizing the customer experience, automating routine and repetitive tasks, improving productivity as well as improving the precision of risk assessment (including fraud detection) practices. But these models come with their own set of risks and hence traditional MRM practices need a facelift to deal with these models.

Ensuring fairness, accountability and transparency should be the key guiding principles when designing policies and processes for AI/ML models. During development, it needs to be ensured that the models are consistent with the company's risk appetite. Tools such as model explanation, bias detection, and performance monitoring need to be built in so that there is constant and consistent oversight. This should be embedded within all AI development activities right from the start. The standards, testing, and controls need to be embedded into various stages of the analytics model's life cycle, from development to deployment and use. Model definitions and tiering principles, governance framework, compensatory controls as well as validation methodologies need to be adapted to address specific risks from AI/ML models. Lastly, ML models unlike traditional models need to be constantly monitored and the validation process needs to take this into account. Upskilling across all levels – from board and senior leadership to model owners, validators and users – is also a critical factor for success.

Over the next few years, the regulatory scrutiny around AI/ML models is expected to grow as banks increasingly start adding more AI/ML models into their inventories. As of 2021, there are detailed proposals from the EU as well as Federal Trade Commissions (FTC) for stricter AI regulations. It remains to be seen how this area evolves and regulations shape up over the next few years, but international guidance or standards in this area will be helpful in setting the minimum benchmark for MRM practices across jurisdictions.

Contact

KPMG in India

KPMG Assurance and Consulting Services LLP
Lodha Excelus, 2nd Floor, Apollo Mills Compound,
N.M. Joshi Marg, Mahalaxmi, Mumbai 400 011

Rajosik Banerjee
Partner and Head of
Financial Risk Management
rajosik@kpmg.com

Amitava Mukherjee
Partner,
Financial Risk Management
amitava@kpmg.com

Kinshuk Pal
Associate Partner,
Financial Risk Management
kinshukpal@kpmg.com

Sreya Paul
Associate Director,
Financial Risk Management
sreyapaul@kpmg.com

KPMG in Germany

KPMG AG Wirtschaftsprüfungsgesellschaft
Klingelhöferstrasse 18
10785 Berlin

Matthias Peter
Partner,
Financial Services
matthiaspeter@kpmg.com

Janek Gallitschke
Senior Manager,
Financial Services
jgallitschke@kpmg.com

Dr. Christoph Anders
Manager,
Financial Services
christophanders@kpmg.com

https://home.kpmg/xx/en/home.html