

Operationalizing AI Compliance Obligations in Financial Services

Operationalizing AI Compliance Obligations Financial Services

Introduction: Financial Services Needs AI, but it Also Needs Help to Govern AI

Financial institutions have been generating statistical and probability models as well as predictive analytics to forecast performance for a long time. The Model Risk Management (MRM) process has reached a level of governance maturity not found in any other industry. The MRM three layers of defense against algorithmic risk has become a paradigm copied across the globe.

But new, advanced AI models have completely changed the governance game. Advanced AI systems incorporate random elements that create inconsistent outputs even without model drift. Instead of the understandable and explainable traditional MRM algorithms, advanced AI models are typically opaque and mysterious, such that even their own creators often cannot explain how they reached their results. Moreover, the intricate nature of advanced AI models requires specialized knowledge to assess their compliance with regulatory requirements – knowledge that may not be within the purview of MRM teams – necessitating new approaches and technologies.


Fortunately for MRM teams that find themselves forced to adapt to the new world of random, opaque, and mysterious AI models, there are new technologies that can help resolve these gaps. One of those new technologies is process mining, a data-driven approach that analyzes event logs from IT systems to visualize and understand business processes. Process mining enables financial institutions to uncover the real processes underlying AI systems by analyzing data from various sources and disparate systems of record.

This whitepaper will explore how process mining, properly applied and supported, can help the MRM team expand their governance of traditional models to even the most complex and mysterious AI systems.

AI Presents New Opportunities - and New Challenges for Financial Services

At least 40% of financial services customer interactions are now AI enabled – from AI-powered chatbots, virtual financial assistants, to automated credit decisioning and real time transaction analysis.

Areas of generative AI (gen AI) application in the credit risk life cycle,¹ % of respondents



A recent McKinsey survey found that 60% of leading global banks intend to deploy gen AI across the entire credit decisioning life cycle.
(graphic courtesy of McKinsey)

Consumers certainly expect to manage their finances via AI systems: 85% of consumers think that robots will soon replace finance professionals. And yet, consumers are inherently apprehensive about trusting AI banking systems: 71% of consumers demand more transparency when they are

interacting with AI-based banking services.¹ This will present challenges for financial institutions relying upon the traditional MRM process.

The Federal Reserve has explicitly recognized that “AI is likely to present challenges in the areas of opacity and explainability.”² (Summarized in Table 1, below)

Table 1: Traditional Models versus Advanced AI Models

	Traditional Models	Advanced AI Models
Approach	Determinative: outputs consistent and predictable	Random: outputs vary; built-in stochastic randomness
Complexity	Low to high: can be complex, but are understandable	Extreme: advanced AI represents the height of complexity
Transparency	Clear: models must be visible to be evaluated in MRM	Opaque: rarely designed to be transparent
Explainability	Full: even if only explainable by those with advanced math degrees	Mysterious: even the AI creators typically cannot explain their actions
Drift Sources	Limited: drift can occur when source data changes over time	Varied: can drift due to data, machine learning or malicious attack

Examining these points in more detail, advanced AI models, such as neural networks and LLMs like ChatGPT, are non linear, leverage multiple input variables, and deliberately incorporate random (i.e., “stochastic”) features – all of which may preclude the ability to provide meaningful interpretability of outputs on any human scale.

Turning to the MRM concepts of transparency and explainability, the OCC’s revised “Retail Lending” booklet of the Comptroller’s Handbook emphasizes their importance in AI governance:

*Transparency and explainability are key considerations that are typically evaluated as part of effective risk management regarding the use of AI (complex) models. . . AI Models applied to significant operations or decisions . . . should be supported by thorough understanding of how the model arrived at its conclusions and validation that it is operating as intended.*³

Yet, despite the importance of transparency and explainability, the OCC Handbook recognizes that many AI systems are inherently difficult to understand and explain:

*There may be challenges with explaining some AI models based on complexity or, in some cases, limited documentation provided for third-party models.*⁴

Further, many advanced AI systems are optimized through continuous machine learning without human interaction, which further complicates interpretability as their models naturally drift as they

¹ Oracle, Money and Machines 2021 Global Study


² Governor Lael Brainard, What Are We Learning about Artificial Intelligence in Financial Services?, November 13, 2018

³ Id.

⁴ Id.

evolve. As well, advanced AI systems, particularly external-facing ones like chatbots, can be subjected to malicious attacks that hijack their outputs in ways that can be impossible to discern from expected model drift without a solid set of performance baselines. Unfortunately, creating such baselines of systems that are massively complex and inherently random requires analyzing results at a vast scale that is rarely achievable without strong analytical technology support. Likewise, AI systems can operate on such a vast scale, processing thousands of transactions in seconds, that the real-time monitoring necessary to protect against model drift is simply not possible without technological support. The result is that advanced AI systems create dangerous gaps in the MRM process. (see Figure 1, below)

Figure 1: AI-driven Gaps in the MRM Process



Despite all of these difficulties, MRM teams cannot ignore the governance challenges posed by AI; the regulators have made it clear that MRM compliance extends to AI.

The Current State of AI Regulation in Financial Services

While SR 11-7 was promulgated before the prevalence of AI, MRM experts, including David Palmer, the chief architect of SR 11-7, see AI as subject to the SR 11-7 framework.⁵ In fact, SR 11-7 parallels the requirements of AI governance, particularly when it comes to testing and validation:

An integral part of model development is testing, in which the various components of a model and its overall functioning are evaluated to show the model is performing as

⁵ See, e.g., Jacob Kosoff, Model Risk Management (available at https://www.kansascityfed.org/documents/8040/Kosoff_Feb_2021_MRM_External_Version.pdf); the authors of this whitepaper have personally heard David Palmer at live events confirm that SR 11-7 is meant to cover AI.

*intended; to demonstrate that it is accurate, robust, and stable; and to evaluate its limitations and assumptions.*⁶

As well, SR 11-7 envisions MRM as an ongoing process that does not stop with development or deployment of models.⁷ This viewpoint was confirmed by the OCC’s issuance of a revised “Retail Lending” booklet of the Comptroller’s Handbook that extends the SR 11-7 MRM models to AI:

*Regardless of how AI is classified (i.e., as a model or not a model), the associated risk management should be comparable with the level of risk of the function that the AI supports.*⁸

There is however another and perhaps more consequential compliance issue: AI explainability focuses on outputs, but it does not address *ex ante* (i.e., “before the event”) adverse inferences and their justification for using AI in the first place. In a 2018 paper “A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI, a pair of researchers at the Oxford Internet Institute make a compelling argument that:

*Explanations can provide an effective ex-post [i.e., “after the event”] remedy, but an explanation can be rendered only after a decision has been made. . . . explanations of a decision do not equal justification of an inference or decision and governance requires both effective ex-ante and ex-post remedies.*⁹

The lack of potential *ex-ante* remedies in AI, whether explainable or otherwise, is troubling, because new AI laws require that AI deployers, such as financial institutions, provide exactly that to prevent the release on the market of potentially-harmful AI systems.

Financial institutions need to find additional ways to bring MRM disciplines to AI, including, as we will see, through Process Mining. But first, we need to discuss why there is a new urgency to operationalize AI governance – because the first AI regulation, the EU AI Act, is here, and it demands rigorous *ex ante* compliance measures.

The EU AI Act Requires Strict AI Governance

The EU AI Act (AIA) was officially adopted on May 21, 2024. The AIA, the first of its kind globally, aims to ensure that AI systems used within the EU are safe, transparent, and respect fundamental rights. High-risk systems, including systems used for “evaluating creditworthiness or establishing credit scores,”¹⁰ face rigorous *ex ante* compliance measures:

- **Risk management** shall be established, implemented, documented, and maintained.¹¹

⁶ Id.

⁷ Id.

⁸ OCC Bulletin 2021-39, Model Risk Management: New Comptroller’s Handbook Booklet, August 18, 2021

⁹ Sandra Wachter and Brent Mittelstadt, Oxford Internet Institute, A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI

¹⁰ EU AIA, Annex III (5)(b)

¹¹ Id., Art. 9

- **Data governance** relating to data preparation processing operations (e.g., annotation, labelling, cleaning, enrichment, aggregation, and possible biases).¹²
- **Record-keeping** that enables the automatic recording of events relating to the operation of high-risk systems.¹³
- **Transparency processes** to enable users to interpret the system's output and include concise, complete, correct, and clear information that is relevant, accessible, and comprehensible to users.¹⁴
- **Human oversight** that prevents or minimizes the risks of AI adverse outcomes.¹⁵
- **Quality management** established and documented in a systematic and orderly manner in the form of written policies, procedures and instructions.¹⁶

To navigate this complex regulatory landscape effectively, financial institutions are mandated to implement robust data governance practices. This includes ongoing auditing and assessment of AI system performance and detection of model drift. However, as discussed above, the randomness, and complexity of AI systems will require monitoring and assessment at a scale simply not achievable through manual, humans-only processes.

Moreover, many institutions operate legacy systems that were not designed to accommodate AI technologies. Integrating AI into these systems can create gaps in traceability, as the existing infrastructure may not support the necessary data logging and monitoring capabilities required for effective auditing. This integration challenge can lead to inconsistencies and gaps in the audit trail, making it even more difficult to ensure compliance.

To sum this all up, AI has become a cornerstone in modernizing various banking operations, notably in credit decisioning. However, due to complex workflows and disparate data silos it is challenging for compliance operations teams to gain granular visibility to potential compliance gaps in AI-driven processes such as credit decisioning. Process mining empowers compliance operations teams to gain real-time insights to process performance, identify bottlenecks, isolate their root causes and remediate potential compliance risks.

Role of Process Mining in Enhancing AI Auditability and Traceability

Why Process Mining?

Process mining is a data-driven approach that analyzes event logs from IT systems to visualize and understand business processes. By transforming event and transaction logs into meaningful representations of processes, process mining bridges data science and process science, providing valuable insights into organizational workflows.


¹² Id., Art. 10

¹³ Id., Art. 12

¹⁴ Id., Art. 13

¹⁵ Id., Art. 14

¹⁶ Id., Art. 17; See *also*, Arts. 19, 21, and 23

Figure 2: The Process Mining Workflow

Process mining enables financial institutions to uncover the real processes underlying AI systems by analyzing and visualizing vast output volumes from various sources and systems of record in a way that no human team could accomplish. This discovery helps identify deviations from intended workflows, allowing for corrective actions to ensure compliance. By providing a clear picture of how AI systems operate within the organization, process mining facilitates the identification of potential compliance risks. This clarity is crucial for both internal audits and external regulatory reviews, as it allows stakeholders to understand the rationale behind AI-driven outcomes. Enhanced explainability fosters trust in AI systems and supports compliance efforts by ensuring that decision-making processes are more accessible.


Process mining directly addresses the challenges faced by financial institutions over AIA high-risk AI system obligations, especially for AI-based credit decisioning processes – as summarized on the chart on the next page:

Table 2: Auditability and Traceability of Compliance Obligations with Process Mining

	Process Discovery	Conformance Checking	Enhancement
Risk Management	Visualize event logs, documentation for risk mitigation and control	Identify process execution deviations from prescribed compliance provisions	Remediate potential process bottlenecks and root causes
Data Governance	Identify data biases and gaps to ensure process fairness	Trigger alerts for processes that fail to meet data governance protocols	Continuous monitoring to trigger faster/immediate remediation
Record-keeping	Visualize data collection, data labelling, traceability	Audit AIA conformance of documentation operations	Monitor collection and documentation processes
Transparency process	Create transparency for responses to adverse decision objections	Confirm transparency levels required by AIA and other regulations.	Systematize ad hoc objection responses into repeatable processes
Human oversight	Discover where human oversight is necessary	Verify the effectiveness of human oversights.	Real-time monitoring for human oversight gaps
Quality management	Create “digital twin” from multiple disparate back-end processes	Confirm AIA compliance	Proactive monitoring for reporting serious incidents

With process mining, compliance operations teams can continuously monitor AI-driven processes to ensure they adhere to regulatory requirements. Process mining can even capture the human element, whether it is humans-on-the-loop, humans-in-the-loop, or humans-in-command, because their decisions are also logged for auditability. This real-time oversight allows for the early detection of non-compliance issues, enabling prompt remediation. Continuous, real-time monitoring also helps maintain the integrity of AI systems by ensuring they operate within the defined compliance parameters. And, again, the continuous monitoring enabled by process mining can provide the only viable solution when the vast scale of AI system operations overwhelms even the best-resourced compliance team. In short, process mining gives financial institutions the ability to fill the gaps in their MRM processes created by advanced AI models. (see Figure 3, below)

Figure 3: Process Mining Extends the MRM Process to Fill the Gaps



Auditability and Traceability of AI compliance Risks


AI has become a cornerstone in modernizing various banking operations, notably in credit decisioning, fraud detection and AML/KYC processes. Credit decisioning in particular is an AI driven business process that may potentially adversely impact eligibility for credit while at the same time has the potential to extend credit to underserved communities. By analyzing vast datasets, AI systems can assess creditworthiness more efficiently than traditional methods. However, the integration of AI introduces compliance risks, including issues related to transparency, fairness, and accountability.

Process Discovery

Auditing of compliance processes in financial services remains fragmented, with the scattering of data pieces across multiple storage locations, databases, or even different formats. In the context

of compliance, fragmented data can mean that crucial information about a customer's transaction history might be stored in one system, while risk assessments or ID verifications could be in another.

Process mining helps compliance operations teams to make data driven decisions as to potential credit decisioning compliance risks, surface their root causes and remediate regulatory risk. Process mining also empowers MRM teams to conduct the necessary *ex ante* testing demanded by the EU AIA at the scale necessary for complex, random, opaque, and mysterious AI systems. The model on the next page depicts three phases of the credit decisioning process, each with its own subprocesses that relate to auditing obligations of model risk management processes:



Process mining enables compliance operations team visualize the execution of credit decisioning processes and make data driven decisions to mitigate AI compliance risks.

(screenshot from ABBYY Timeline)

- **On boarding:** Involves the use AI-based algorithms – including increased use of biometric identification and intelligent OCR for capturing and extracting data from trailing documents received from loan applicants.
- **Credit Adjudication:** One scenario is that following credit checks (based on FICO and increasingly other proxies) the credit is approved. The other scenario is that credit is denied. In the event of credit denial, there are specific obligations imposed on banks to mitigate unfair credit outcomes which may result in adverse impact on protected groups.
- **Post Credit Denial Auditing of Credit Decisioning:** There are prescribed compliance protocols that must be followed, must be documented, adverse events logs must be maintained. Federal Reserve Board specifically requires financial institutions to identify, monitor and mitigate possible adverse consequences of lending decisions based on models that are incorrect or misused through active model risk management.




Compare the discovered processes against established compliance requirements and internal policies. Identify any deviations, bottlenecks, or patterns that may indicate compliance risks, such as potential biases or data quality issues.

(screenshot from ABBYY Timeline)

Compliance Analysis

Compare the discovered processes against established compliance requirements and internal policies. Identify any deviations, bottlenecks, or patterns that may indicate compliance risks, such as potential biases or data quality issues.



Protocol item	Violation type	Count	Timeliness	
Anomalies Detected	Missing step	283 (34.98%)	4.79%	Apply
Evaluate Training Data Set	Wrong position	1 (0.12%)	0.02%	Apply
Foreseeable Risks Identified	Wrong position	1 (0.12%)	0.02%	Apply
QA Process Documented	Wrong position	4 (0.49%)	0.07%	Apply
Results Reviewed by Compliance Officer	Missing step	283 (34.98%)	4.79%	Apply
Results Reviewed by Compliance Officer	Time violation	120 (14.83%)	2.03%	Apply
Risk Mitigation Control Measures Verified	Wrong position	4 (0.49%)	0.07%	Apply

Protocols include sets of rules or procedures for events in a process. Many processes have protocols or rules that must be followed. The Protocol analysis tool uncovers areas where protocols are potentially not followed in the credit decisioning process and identifies root causes and frequency.

(screenshot from ABBYY Timeline)

Process Enhancement

Based on the insights gained, implement necessary adjustments to the AI models and decisioning processes. This may involve retraining AI algorithms with unbiased data, enhancing data validation procedures, or modifying decision rules to align with compliance standards.

Bottleneck analysis Filtering tools

Select analysis: Time from the event to the next

Total time saving: 243,423.95 hour

Event	Count	Per timeline	Average time (h)	New time (h)	Total time (h)	Total time
Loan Documents Processed	7,786	1.32	1.53		11,926.20	0.47%
Verification of Loan Documents	7,737	1.31	62.25		481,626.73	19.06%
Loan Application Submitted	5,910	1.00	1.55		9,169.05	0.36%
Credit Check	5,541	0.94	103.93	60	332,460.00	13.16%
Risk Analysis	5,463	0.92	143.56		784,252.50	31.03%

The Bottleneck analysis surfaces inefficient steps in the process by analysing the time spent between events and their associated costs. Bottlenecks are often caused by unnecessary or inefficient processes; being aware of the extra losses is the first step to eliminate them. The Bottleneck analysis enables compliance operations teams to isolate and model out the impact of changes to the process flow and then make data driven decisions to identify and implement process improvement efficiencies.

(screenshot from ABBYY Timeline)


Ongoing Monitoring

Establish ongoing monitoring mechanisms using process mining to ensure that the credit decisioning processes remain compliant over time. Regularly update the process models to reflect any changes in regulations or internal policies.

Monitor process execution in near real-time to give you continuous performance insight, prevent process deviations, set and customize alerts for process execution, whether it's performing as planned or not and trigger an email to users, calls to other Business Process Management (BPM) processes, or new Robotic Process Automation (RPA) bots.

Executive Dashboard

The dashboard provides at-a-glance views of the KPIs relevant to the credit decisioning process metrics and can be used to drill in further on individual timelines.



Conclusion: Process Mining Makes Explainable AI Possible for MRM

Process mining helps financial institutions comply with MRM requirements by enhancing transparency, accuracy, and efficiency in AI model governance. Process mining maps end-to-end processes, identifying inefficiencies, deviations, and compliance risks. By providing real-time monitoring, process mining ensures adherence to regulatory requirements, reducing operational risks. It strengthens model validation by offering data-driven insights into model performance, usage, and limitations. Automated auditing improves documentation and traceability, aiding regulatory reporting. Additionally, process mining optimizes model lifecycle management, supporting risk mitigation and control mechanisms. This leads to better decision-making, reduced compliance costs, and improved regulatory alignment in financial institutions.

Andrew Pery: Andrew.Pery@abbyy.com

Michael Simon: Michael.Simon@LawPlusData.com