



EUROPEAN CENTRAL BANK

BANKING SUPERVISION

AI workshops with banks 2025

Annex

November 2025




Please Note

Please note that this document presents the **anonymised self-reported information of participant banks** as provided to the ECB in the SSM workshops on AI 2025.

The contents of this report are observational and intended for **informational purposes**, and do not constitute a supervisory assessment.

Aggregated overview (AI, API, white labelling, other)

Innovative use cases – adoption rate in percentages



¹year-on-year differences calculated between 2023 ECB Short term exercise (STE) and 2022 Digitalisation Survey results (in red) and 2024 and 2023 STE results (in blue).

2025 AI workshops with supervised banks – background



Purpose of workshops – improve level of understanding of where banks stand in the application of AI, with a focus on the **general aspects of AI** as well as its specific use for **credit scoring** and **fraud detection**. Workshops are run in a **non-prescriptive, exploratory manner, without bank specific follow-up (F&M)**, to identify relevant risks from a microprudential supervisory perspective.



Context – digitalisation and new technologies are a supervisory priority for 2025. This is a continuation of previous work on digitalisation, building on learnings from the working group on machine learning in internal models.



Sample of 13 supervised banks.



Format and timeline – three to four-hour remote/hybrid meetings with the banks between May and August with predefined agenda. A set of questions was circulated to the banks before the meeting, but there were no further interactions or questionnaires outside of the workshops.


AI workshops with supervised banks – high-level sample summary

The sample consisted of **13 banks**, of which:


- **10 banks used AI for credit scoring** (i.e. ~67% of significant institutions reporting use cases in production based on STE templates).
- **10 banks used AI for fraud detection** (i.e. ~26% of significant institutions reporting use cases in production).

Participating banks are headquartered in nine European countries.

Business model



Total assets (€ bn)



Key learnings from the 13 workshops



Governance and risk management of AI systems – governance arrangements are being established mostly by integrating AI into existing policies and risk management frameworks or by creating dedicated AI governance functions (e.g. committees, AI units) and policies. The effectiveness of these functions is yet to be measured.



AI models tend to be developed internally but hosted by cloud service providers – the majority of banks in the sample rely on in-house solutions for both credit and fraud use cases; these are developed by internal centres of competence at larger institutions. Third-party solutions are employed selectively.



A “human in the loop” remains central to banks’ processes – for credit scoring, AI models support decision-making by humans, except for small retail loan originations, which are sometimes automated. Similarly, AI models for fraud detection operate in real time, generating alerts to be followed up by fraud expert investigators.



Intended explainability – AI models are integrated into model cycles: expert reviews of parameters during development and testing are carried out, with regular monitoring of model performance and plausibility of outcomes, leveraging on explainability tools; no self-learning is utilised after deployment, and no uniform understanding of explainability across banks yet.



Data governance integration yet to be established – effective and risk-based application of internal data management standards considering the specificities of AI models are emerging only in a small number of cases.



Progressing towards compliance with EU AI Act ¹ – preparations for the AI Act are in progress, with banks starting to conduct compliance self-assessments and system inventories, and implementing processes for new use cases, despite uncertainty on the steps to take in terms of compliance at some banks.

¹ Most provisions becoming effective in August 2026 for new AI systems or significant changes of existing ones.

Credit scoring and fraud detection: specific learnings from workshops

Credit Scoring



Banks are integrating AI solutions – alongside traditional models – to improve loan decision-making. They are conducting regular assessments of external AI providers, preparing for AI Act compliance, tracking AI performance using varied metrics and fostering explainability.



AI is beneficial in credit granting and monitoring – especially in terms of accuracy and explainability. No banks reported using GenAI in the context of credit scoring, citing challenges such as development time, cost and trustworthiness. Instead, they rely on well-established machine learning techniques.



Regular model monitoring is combined with expert reviews to identify errors and biases – few banks have automated validation tools or explicit fallback procedures for AI failures. Instead, they rely on structured control frameworks as per regular model risk management.

Fraud Detection



AI benefits fraud detection – banks reported tangible benefits from using AI in fraud detection, such as higher accuracy, resulting in reduced fraud losses and fewer false positives.



Flexible sourcing, well-established techniques with human in the loop – banks use internally developed models using established ML techniques to detect types of fraud, while maintaining human involvement for oversight and decision-making.



Compliance with EU AI Act and key practices – banks expect fraud-detection AI to be classed as low-risk under the AI Act. Key practices include clear business cases, testing, explainability, oversight and governance.



Varied approaches regarding specific policies and governance structures

The majority of banks are **integrating the governance of AI risks into their existing risk management frameworks**, while some have made adjustments to cater for the specificities of AI¹ or are considering doing so by establishing dedicated policies and/or committees. Some larger institutions have established dedicated 1LoD teams to centrally provide AI models.

Effectiveness of emerging governance patterns for risk management and compliance is yet to be measured.


Emerging practices

A number of banks are setting up a dedicated AI policy, dedicated committee(s) and/or have established a Chief AI Officer.

Clear link to digitalisation strategy, and (publicly shared) KPIs.

Dedicated senior roles in the 1LoD (e.g. Chief AI Officer) and staff in 1 and/or 2 LoD with veto/approver rights for AI use cases.

Organisation of AI risk management in banks



¹ For instance, levels of transparency and explainability, mandatory use of specific models or data sources and AI Act requirements.



Models often developed internally but hosted by cloud service providers

Within the use cases of credit scoring and fraud detection, **most banks reported using in-house solutions**, often leveraging on external risk engines, data or open-source libraries. Some banks opted to source AI models via partnerships with specialised vendors, tailoring solutions for use cases or a “model-as-a-service solution”.

As banks increasingly rely on external providers, they are becoming more aware of the associated risks, including data privacy, operational resilience and regulatory compliance. Follow-up could be required regarding deficiencies in operational resilience frameworks with regard to cybersecurity and third-party risk management capabilities as a prioritised vulnerability.

Emerging practices

Self-hosted (LLM) models on private cloud or with distant back-up locations. Use of **EU-based LLM model provider**.


An internally developed model with external AI models from separate providers, used to balance workload and ensure continuity in event of failure.

Detailed compliance assessments for chain outsourcing.

Sourcing of use cases



Dependencies on external providers





Human in the loop remains central in banks' processes

For credit scoring, AI models support decision-making by **humans**, except for small retail loan originations which are often automated.

Similarly, AI models for fraud detection operate in real-time, generating alerts to be followed-up by fraud expert investigators.


Interpretation of human in the loop is still varied, in some cases limited to explaining the credit decision to customer on demand.

Emerging practices

Higher degree of **human oversight for higher-risk, larger applications**, more complex products or higher-impact decisions.

Establishing a **feedback loop**, where experts who act on AI outputs provide input back to the model to support continuous improvement and fine-tuning during model testing.

Human in the loop





Explainability techniques focused on monitoring results

Most banks state that they can provide a detailed explanation of how their AI systems reach decisions, especially for self-developed models, **leveraging expert review of parameters** during development and testing, and **regular monitoring of model performance and plausibility of outcomes**, leveraging on explainability tools and quantifying input variable contributions. While some could explain the output based on specific risk drivers selected but not the detailed decision process.


No unified understanding across banks of what explainability means in practice. No bank allows for self-learning after deployment.

Emerging practices

Industrialised monitoring with dashboards set up by central team provided together with the models (or modules) using standardised KPIs and data visualisation.

Partnership with external parties to break down model's overall output into **individual feature-level** (input variable) **contributions**.

Banks' views on explainability of AI systems





Data governance frameworks yet to adapt AI aspects

Data Governance frameworks¹ apply to data used for AI models, and in some cases, are **being adopted for the specificities of data used for AI** (large quantity, less structured).

Often, data quality aspects are incorporated by **applying the bank's model risk management framework** to the AI models.

In some cases, data governance aspects are **explicitly incorporated into the AI use case funnel**.

Only a few banks reported to effectively applying data management standards in practice, and to adjusting them to the specific requirements of AI models. This is critical, as poor data inputs will inevitably lead to unreliable results. This also stresses the need to follow-up on the ECB Banking Supervisory priority on risk data aggregation and risk reporting (RDARR).

Emerging practices

Chief Data Officer appointed as AI Officer in dual role.

AI models use **golden sources** incorporated in centralised AI module repository, with exceptions clearly defined and requiring approval.

¹ See [ECB Guide on Risk Data Aggregation and Reporting](#).



AI Act preparations are in progress, at differing stages – some issues on definitions and interpretations remain

All banks preparing to implement the AI Act¹ have established a specific process to put AI models into production (AI funnel) and built up **AI systems inventories**.


Most banks have performed a self-assessment of the compliance of their AI systems **with the AI Act**, typically based on developers' self-assessment or on their policies' coverage of AI Act requirements. There are **varying degrees of certainty** regarding definitions and implementation.

Emerging practices

AI Act compliance is supported by **automated tools for all AI use cases**, monitoring inventory and workflow.

Mapping of AI act definitions and categorisations to internal ones.

Self-assessment performed



¹ Most provisions take effect in August 2026 for new AI systems or significant changes to existing ones. Please note that the ECB will not assess AI Act compliance, as it is not the market surveillance authority.

Key learnings on use of AI for credit scoring



AI is used in a range of applications across credit granting and monitoring, with a consistent focus on accuracy – most banks in the sample use well-established AI techniques in credit scoring and observed increased model performance, together with the explainability of model outcomes.



Prevalence of in-house development when sourcing AI credit-scoring models – this is related to existing in-house model development capabilities and portfolio-specific considerations such as available data inputs, but also to confidentiality. External cooperation for model deployment (platform operation of self-built model) is observed occasionally.



AI models coexisting with legacy tools dominate the sample – banks in the sample tend to replace or complement their existing loan origination or monitoring tools (at least for an initial period) with AI-powered ones. This symbiosis of both traditional and innovative approaches is viewed as providing additional assurance in terms of explainability while improving accuracy in predictions.



Slow uptake of GenAI/GPAI – banks across the sample reported using established AI techniques rather than GenAI for credit scoring, mainly owing to challenges related to development time, expense and trustworthiness. Tailored creditworthiness-related AI systems are also preferred over GenAI/GPAI owing to the specificity of loan contexts.



AI Act readiness picking up speed – almost 80% of the sample reported having started performing AI Act readiness- or compliance assessments for their credit scoring AI tools, while 20% are still working on implementing the assessment procedures. Some banks reported challenges associated with the AI Act's implementation; these were related to definitions and possible overlap with existing regulations.

Key learnings on use of AI for fraud detection



The benefits are real – despite some initial challenges, banks report tangible benefits from using AI in fraud detection, most importantly higher accuracy, resulting in reduced fraud losses and fewer false positives.



Broad usability and flexible sourcing – across banks, AI systems are used to detect various types of fraud, including fraudulent transactions, account takeover, identity fraud and loan fraud. Most banks use internally developed solutions, though some also use third-party providers.



AI in fraud detection is not (yet) a revolution – banks apply well-established machine learning techniques (e.g. gradient boosting) in fraud detection, often starting from a specific process step, channel or geography. New and revolutionary techniques like GenAI have not yet been adopted in fraud detection use cases by any bank participating in the exercise.



Humans are staying in the loop – even with the high levels of automation and varying degrees of autonomy of AI systems, all banks maintain some human involvement in the process, either for final decision-making or for ex-post oversight (e.g. in the case of client complaints).



Lower impact of AI Act – although the assessment is still ongoing in many banks, it is not considered likely that AI systems used for fraud detection will be classified as high-risk in most cases.