

ARTICLE



IM_

AI in banking

The AI landscape: traditional, generative, and agentic AI

Traditional artificial intelligence (AI) refers to machine learning and statistical models that operate on structured data to automate, predict, and optimise decision-making processes. These systems are designed to identify patterns and improve outcomes in areas such as forecasting, risk management, and operational efficiency.

Generative AI (GenAI) extends these capabilities by using large language models (LLMs) and other generative architectures that are built to handle general inputs and outputs. GenAI can analyse and create content from unstructured data such as text, images, or audio, enabling more interactive and context-aware use of information. An assistant (often used interchangeably with agent) is an LLM configured to perform a specific task through prompt engineering and data access.

Agentic AI systems equip LLMs with a level of autonomy and access to a predefined toolbox allowing models to perform complex tasks such as writing code, drafting slides, querying databases, and accessing systems.



The rise of artificial intelligence is not just transforming how we work; it's reshaping how we live and learn. This shift brings enormous opportunities to enhance productivity, boost competitiveness, and create new value. But it also raises an urgent question: are we ready to adapt to this new reality?

Ana Botín, Executive Chair, Santander Group

The state of AI in banking

AI in banking has shifted from experiments to enterprise capability, but adoption is uneven. Most institutions sit between AI-native challengers and low-maturity incumbents. Where AI is present today, it tends to operate in isolated processes rather than end-to-end journeys, with cost reduction and efficiency as the primary drivers. Legacy mainframes, fragmented integrations, and inconsistent data quality remain persistent barriers to scaling.

Still, momentum is building across the industry. Fraud and financial crime monitoring are well established, and virtual agents are mainstream. Advisor copilots, contract automation, and regulatory tools are scaling into operational use. Risk and explainability assistants are validated but not widely deployed yet. In the near term, we expect continued scaling within customer service and back-office automation where return on investment is clear and risk is manageable.

Regulation is changing the calculus. New requirements for operational resilience, capital, and financial crime raise the bar on explainability, oversight, and reporting under DORA and CRR3/CRD6. Many banking AI applications qualify as high risk under the EU AI Act, making governance a prerequisite for future AI projects in banking.

In recent C-level dialogues at Implement Consulting Group, the focus has moved from whether to invest in AI to how to scale applied, well-governed capability. This document maps twelve high-impact use cases, indicating where adoption is mature, accelerating, or contingent on governance and modernisation.

Research indicates that banks are experimenting broadly, but systematic adoption remains uneven

Current state of AI in the banking industry...

Banks are increasingly adopting GenAI and Agentic AI solutions to automate complex, data-intensive tasks, enhancing customer experiences, and accelerating back-office operations. Key use cases include customer service, fraud detection, personalised marketing, automated document processing and general employee productivity. These areas have demonstrated concrete, measurable impact and are moving from pilot to production at leading institutions

Complex use cases such as virtual agents, regulatory intelligence tools, and similar solutions are in active piloting or have been validated with senior risk leadership. Yet overall maturity remains low to medium.

AI tends to live in isolated processes rather than end-to-end journeys, and AI maturity depends heavily on the digitalisation level of the institution.

... but what is hindering adoption?

The primary barriers are structural, not aspirational. Legacy mainframes, fragmented integrations, and poor data quality consistently emerge as the top constraints to scaling AI.

But technical debt is only part of the story. Banks operate in a highly regulated environment, and legislation such as the EU AI Act classifies many common banking AI applications as high-risk, bringing governance and explainability to centre stage.

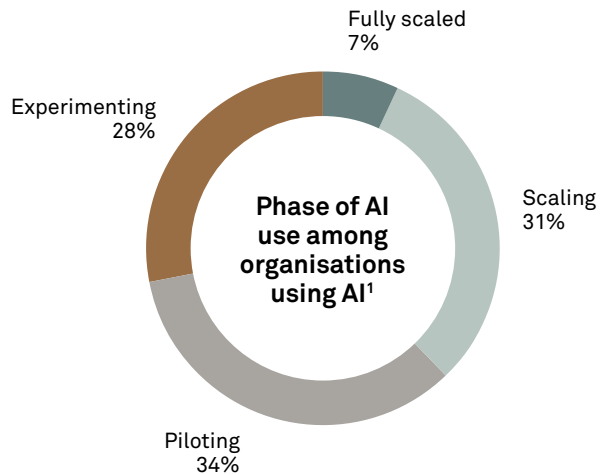
The result is a tension between urgency and caution: banks see clear return on investment across their value chains, but moving to production demands governance maturity, organisational alignment, and reskilling.



84% of business leaders believe that AI will have a significant impact on their business¹

78% of organisations report using AI in at least one business function¹

However, only 7% of the organisations using AI have fully scaled it, whereas the others are either piloting, scaling, or experimenting.



Sources: ¹⁾ Implement Analysis (2026), Rand (2024), Stanford HAI AI Index Report (2025), MIT CISR Enterprise AI Maturity Model (2024)

AI across the value chain of banking

AI is already transforming banking operations and service delivery, creating competitive advantages for early adopters.

	Use case	Description	Industry adoption ¹			
			Emerging	Low	Medium	High
01. Sales and servicing (front/middle office)	Personalised next-best-offer outreach	Machine learning surfaces product opportunities per customer; AI crafts proactive messages	○	○	●	○
	Banking advisory assistant (advisor co-pilot)	Meeting prep, note summarisation, and credit application assistance before/during/after client interactions	●	○	○	○
	Customer service virtual agent uplift	Augment vendor chatbots with bank data / orchestration to improve resolution and customer satisfaction	○	○	○	●
02. Operations (back office)	Contract drafting and review automation	GenAI drafts standard contracts and flags risky clauses for operational / commercial teams	○	●	○	○
	Advanced churn prediction (traditional machine learning + LLMs)	Analyse contracts for risks, compliance, and financial terms, improving negotiation efficiency	○	●	○	○
	Complaint analysis with natural language processing	LLMs to analyse, categorise, and monitor customer complaints to improve efficiency and quality	○	○	●	○
03. Support functions	Legacy code modernisation with AI	GenAI-assisted code understanding / refactoring to reduce technical debt and decouple integrations	○	○	○	●
	Real-time steering and regulatory reporting	Improve pipelines and AI-assisted analytics for near real-time KPIs (T+1 → real-time)	○	○	●	○
	AI assurance layer for explainability	Assurance tooling for explainability, monitoring, and controls aligned to EU AI Act high-risk needs	●	○	○	○
04. Risk and compliance	Regulatory shield (regulatory change intelligence)	Link internal policies / procedures to external regulations / sanctions; monitor changes and trigger updates	○	●	○	○
	Transaction monitoring with AI and GenAI (AML / KYC, ODD)	Anomaly detection for AML / fraud; GenAI supports periodic reviews at scale	○	○	●	○
	GenAI-assisted enterprise risk management	Risk identification, bow-tie analysis, and conversational access to risk registers	○	○	●	○

Highlighted AI use cases in banking

Several advanced AI use cases are implemented at full industrial scale and can deliver significant returns on investment from day one.

01. Personalised next-best-offer outreach	Why:	Banks primarily serve existing customers and are not efficient in targeting new customers or products. Proactive advice can uplift customer experience and engagement at scale.
	Examples:	In-app proactive messages generated by GenAI, leveraging customer holdings and behaviours. Existing use cases with Santander UK, KBC Belgium, HSBC, and others.
	Impact:	Higher digital engagement, satisfaction, and spending; more clicks on tailored offers; stronger customer loyalty and switching intent driven by AI-powered insights.
02. Contract drafting and review automation	Why:	Manual, repetitive drafting with error risk; long documents are often signed without thorough review.
	Examples:	GenAI produces standard drafts and flags risky clauses for operations. Examples seen with BNP Paribas, UBS, and Deutsche Bank.
	Impact:	Faster reviews and complex analysis with high lawyer adoption; shorter processing times in banking workflows; significant untapped potential in negotiations.
03. Transaction monitoring (AML/KYC/ODD)	Why:	Regulatory obligation to monitor transactions for AML / fraud. Periodic KYC reviews require scalable intelligence to reduce manual workloads.
	Examples:	AI-enabled transaction monitoring with PKR, combining anomaly detection and GenAI summarisation. HSBC has deployed third-party software, while Danske Bank collaborated with Feature Space's ARIC framework to reduce fraud.
	Impact:	More suspicious activity detected with fewer false positives; major cost and headcount reduction potential; stronger compliance posture and avoidance of remediation and penalties.
04. Regulatory shield (regulatory change intelligence)	Why:	Regulatory volume inflation from the EU AI Act, DORA, CRR3/CRD6, SFDR, and ESAP is demanding explainability and assurance.
	Examples:	Custom solutions to link policies and external regulations have been built at Nordea and UBS. Various banks utilise third-party platforms to handle regulatory changes (Barclays, BNP Paribas, Citi, J.P. Morgan, etc.).
	Impact:	Rapid regulatory analysis at scale with strong tier-1 adoption signals; improved operational risk management and audit readiness; essential capability as concurrent regulations make manual tracking infeasible.

Reference cases

Implement reference cases

The two cases illustrate how Implement Consulting Group has embedded AI into core business processes with clients, delivering tangible results and meaningful business impact.

Accelerating regulatory compliance and reducing review effort with generative AI **Regulated enterprise // Compliance GenAI**

A client in a highly regulated industry struggled to keep pace with frequent regulatory updates across jurisdictions. Manual consolidation and assessment made compliance slow, error-prone, and difficult to scale. We designed an AI-driven compliance workflow that automatically screens regulatory updates, identifies relevant changes, and supports faster, more consistent regulatory adaptation with human oversight.

Impact

- 50% faster regulatory review processes
- 50%+ reduction in early compliance labour costs
- Improved accuracy and consistency, enabling amore proactive and risk-resilient compliance approach

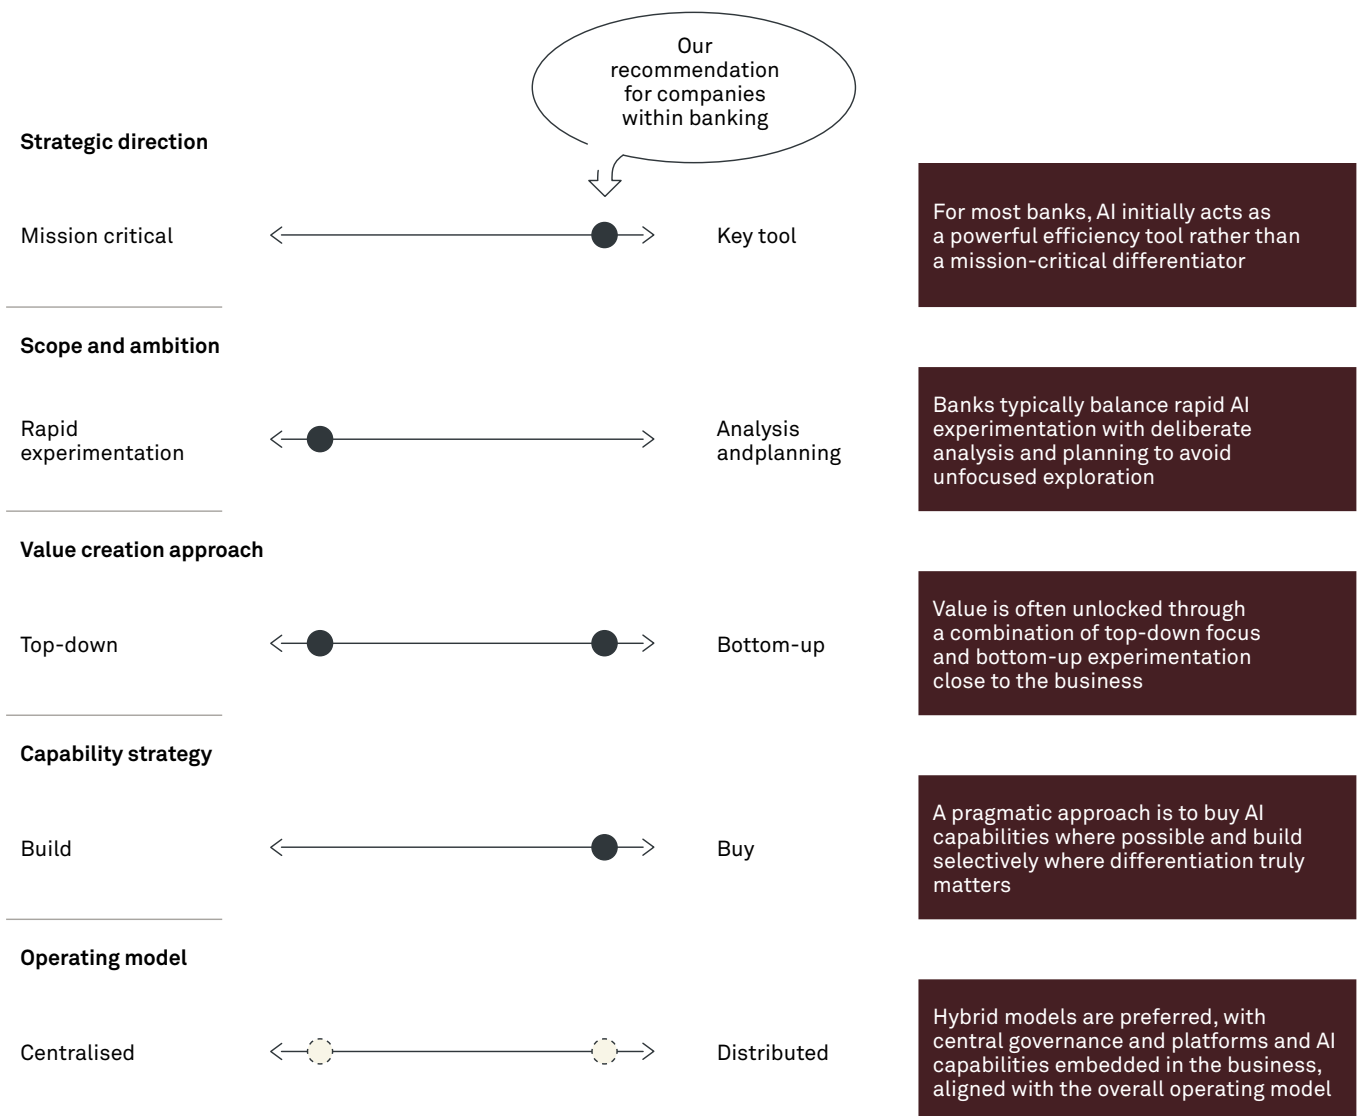
State-of-the-art agentic AI solution with potential to fully automate service requests **Customer operations // Agentic AI**

A client sought to move beyond static FAQs by deploying an autonomous, agentic AI solution. Implement supported the client by deploying specialised AI agents that collaborate directly to resolve complex queries in real time, avoiding central bottlenecks. The decentralised architecture enables low-latency, personalised responses and establishes a scalable foundation for future voice-enabled interfaces and internal productivity workflows.

Impact

- Enabled low-latency resolution of complex service requests, even during peak demand
- Delivered personalised responses by leveraging live core system data
- Established a scalable foundation for 24/7 automated service operations

In our experience, banks must make explicit strategic choices to unlock the value of AI



In our experience, banks must make explicit strategic choices to unlock the value of AI – deep dive

From potential to value

How do we translate AI capabilities into sustained customer and shareholder value? There is no universally right answer. In fact, there is not even a single right question. Banks are instead faced with a set of interdependent strategic choices that must be navigated over time. Each choice involves trade-offs. Not choosing is also a choice, often with a tangible cost. Together, these choices form the foundation of an organisation's AI strategy.

These considerations are rarely best addressed as a linear sequence. They are better understood as a 'choice cascade', where decisions in one area shape what is realistic in others. Many AI initiatives stall, not because the technology is immature, but because underlying choices are not formally made leaving the organisation in a state of uncertainty.

Strategic direction

An AI strategy, like any divisional strategy, should enable the organisation's overall strategic direction. A central consideration is whether AI is viewed as mission critical to long-term competitiveness, or primarily as a powerful tool to improve and optimise existing operations.

When AI is viewed primarily as a tool, investments tend to be more cautious and closely tied to near-term business cases and measurable returns. Timing matters as well. Some banks expect AI to become mission critical over time, while remaining a useful optimisation lever in the short to medium term.

In our experience, most banks today lean towards this latter position. AI is often approached as an efficiency and productivity lever first. This can be a sensible starting point, provided the choice is explicit and revisited as technology and competitive dynamics evolve. Strategy, is as much about choosing what not to pursue –as it is about selecting where to invest.

Scope and ambition

One of the most practical challenges in any AI strategy is deciding how to identify opportunity areas to pursue. Many banks balance between rapid experimentation on the one hand and more deliberate analysis and planning on the other.

Rapid experimentation can accelerate learning and quickly surface opportunities with high business impact. At the same time, without clear oversight and coordination, it risks leading to fragmented initiatives with limited scalability. Analysis and planning provide structure and coherence, but can slow learning in a fast-moving field if taken too far.

For most banks, the question is not which approach is right, but how to combine them. The aim is to enable experimentation within a clearly defined scope, while ensuring that successful initiatives can be industrialised and scaled.

Value creation

Closely linked to scope and ambition is how specific AI use cases are identified. A top-down approach, where leadership defines focus areas, can work well where problems and outcomes are relatively well understood. A bottom-up approach, where teams are empowered to explore new applications, can support adoption and uncover new value.

Each approach has strengths and limitations. In practice, banks that are currently focus on using AI as an efficiency and productivity lever, tend to lean towards a top-down approach. Though, all banks are likely to benefit from a deliberate combination, aligned with their ambition and maturity.

Capability strategy

Once priorities are clearer, the question of build versus buy becomes pertinent. Off-the-shelf solutions typically offer faster time to value at lower cost, but limited differentiation. Custom solutions offer greater control and potential competitive advantage, but require sustained investment and ongoing ownership.

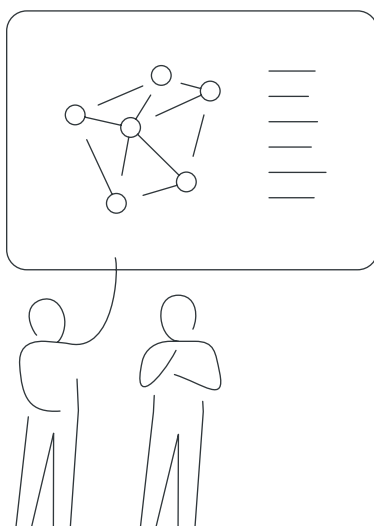
A pragmatic stance for many banks is to always buy where possible and build selectively where off-the-shelf solutions are not viable, particularly close to the core of the business where competitive advantage matter more. Given the pace of AI development, patience can sometimes be as important as speed.

When AI becomes part of everyday operations rather than a parallel structure, value realisation becomes more likely.

Operating model

Regardless of sourcing choices, AI capability will be developed in-house. That capability can be centralised or distributed across the organisation. Centralised models support scale, consistency, and governance. Distributed models support business relevance and adoption but are often harder to scale.

As a result, many banks move towards hybrid models, with central governance and platforms combined with AI capabilities embedded in the business. The optimal hybrid is often where the AI operating model aligns with how the organisation already runs the business.



A concluding thought for leadership teams

Despite significant investment, few organisations have yet demonstrated sustained bottom-line impact from AI. Ironically, this is pushing many organisation towards rapid time to market initiatives often with more promising short term ROIs – in turn limiting longer term scalability and operational efficiency.

Realising sustained AI impact is driven by a small number of deliberate strategic trade-offs, rather than isolated technology investments.



Future outlook

The future is impossible to predict. But what is certain is that AI can no longer be treated as something to explore “when the time is right” – as that time is now.

Across industries, including banking, organisations are moving in a common direction. Not because the destination is known, but because standing still is no longer an option. The early leaders are not those with perfect roadmaps, but those who take AI seriously, invest early, and remain curious and experimental.

In the near term, progress often starts with pragmatic use cases such as customer service augmentation, advisory support, and back-office automation. These are not ends in themselves, but entry points, safe spaces to experiment while building governance, explainability, and regulatory readiness under the EU AI Act.

The harder work lies inside the organisation. Technology matters, but the real challenge is understanding where AI can fundamentally reshape value creation across the organisation and value chain – and where experimentation, not optimisation, is required.

At Implement, we see early leaders investing not just in technology, but in AI factories, data foundations, assurance, and, most importantly, organisational learning. Societal adoption will influence how AI evolves, but lasting impact will depend on organisations’ ability to disrupt themselves from the inside out.

