



ROLE OF ARTIFICIAL INTELLIGENCE IN TRANSFORMING THE INDIAN BANKING SECTOR

Kanishka Yadav¹, Dr. Saurabh Singh²

¹Department of Commerce, Regional College of Professional Studies and Research

²Department of Commerce, Regional College of Professional Studies and Research

ABSTRACT

Artificial intelligence is moving from a supporting technology to a decision infrastructure within banking. In India, this transition is especially significant because AI is being layered on top of a large digital public infrastructure that includes Aadhaar, Jan Dhan accounts, mobile connectivity, UPI, the Account Aggregator framework and emerging digital-credit rails. This paper examines how AI is transforming the Indian banking sector across customer service, onboarding, credit underwriting, fraud detection, anti-money-laundering controls, payments, collections, compliance, operations and managerial decision-making. The study uses a descriptive and analytical research design based on secondary evidence from Reserve Bank of India policy material, NPCI information, Indian data-protection law, international financial-stability guidance and recent peer-reviewed research. The central argument is that the most important effect of AI is not simple automation; it is the shift from transaction processing to continuous prediction, personalisation and risk sensing. At the same time, AI creates new forms of model risk, bias, opacity, cyber vulnerability, third-party dependence and accountability problems. Evidence from Indian banking research also indicates that customers value the speed and accessibility of AI-enabled services but continue to value human assistance for complex or high-stakes interactions, while employee adoption depends strongly on perceived usefulness, ease of use and managerial support. The paper proposes a responsible-AI model in which human accountability, explainability, privacy, model validation, auditability and continuous monitoring are embedded throughout the AI lifecycle. It concludes that India can derive substantial gains from AI in banking only when innovation is combined with institutional trust, strong data governance, inclusive design and proportionate regulatory safeguards.

KEYWORDS: Artificial Intelligence, Indian Banking, Machine Learning, Digital Banking, Credit Scoring, Fraud Detection, Financial Inclusion, Responsible AI, RBI, Digital Public Infrastructure

1. INTRODUCTION

Banking has always been an information-intensive industry. Banks collect information about depositors, borrowers, transactions, collateral, payment behaviour and financial risk, and then convert that information into decisions. Artificial intelligence (AI) changes the speed, scale and granularity of this conversion. Traditional banking systems were designed primarily to record transactions and enforce predefined rules. AI-enabled banking systems can learn from large volumes of structured and unstructured data, detect patterns that are difficult for humans to observe, predict likely customer or risk outcomes, and support decisions in near real time. This changes the bank from a passive record keeper into an institution that increasingly anticipates customer needs and emerging risks.

The Indian context makes this transformation unusually important. Over the past decade, India has created a broad digital public infrastructure around identity, bank accounts, mobile connectivity, real-time payments and consent-based data sharing. In a 2024 address, the Reserve Bank of India (RBI) noted the scale of Aadhaar, Jan Dhan and mobile connectivity, and described the role of UPI, the Account Aggregator framework and the Unified Lending Interface in reducing the frictions associated with payments and credit (Das, 2024). These systems generate digital trails and interoperable data flows that can make AI more useful, but they also enlarge the consequences of poor model design, weak cybersecurity or misuse of personal data.

The current phase of banking transformation therefore differs from earlier computerisation. Core banking solutions automated account records. Internet and mobile banking shifted transactions away from branches. UPI dramatically reduced payment friction. AI now sits above these layers and influences how banks interpret data, allocate attention, detect fraud, assess credit, personalise offers and manage operational exceptions. In this sense, AI is not a single banking product. It is a general-purpose decision technology that can alter nearly every stage of the banking value chain.

The transformation is already visible in official policy discussions. RBI leadership has identified applications such as alternative-data credit assessment, real-time fraud detection, KYC and loan-processing automation, chatbots, document extraction and RegTech, while simultaneously warning about bias, opacity, over-automation and weak accountability (Sankar, 2025). RBI-linked initiatives such as MuleHunter.ai have also demonstrated the use of machine learning against mule accounts; in October 2025 the RBI Deputy Governor reported deployment in about twenty commercial banks and noted continuing work on a Digital Payments Intelligence Platform for real-time transaction risk scoring (Sankar, 2025).

This paper analyses the role of AI in transforming Indian banking from five connected perspectives: operational efficiency, customer experience, risk management, financial inclusion and institutional governance. It also examines a central tension: the



same capability that makes AI powerful—its ability to make rapid inferences from very large datasets—can create unacceptable outcomes if data are biased, the model is opaque, responsibility is unclear or customers cannot contest important decisions.

Therefore, the study treats AI transformation not simply as a technology-adoption question but as a problem of organisational design, risk governance and public trust.



Figure 1. Flow of AI-led transformation in Indian banking

2. CONCEPTUAL FOUNDATION: WHAT AI CHANGES IN BANKING

Artificial intelligence in banking includes a family of technologies rather than one tool. Machine learning identifies relationships in historical data and produces classifications or predictions; deep learning processes highly complex patterns; natural language processing interprets text and speech; computer vision can examine documents and images; and generative AI can produce or summarise text, assist employees and support conversational interfaces. Robotic process automation is not always AI in itself, but it becomes more intelligent when combined with classification, extraction and decision models.

The transformative feature of AI is that it changes the economic unit of banking from the transaction to the prediction. A conventional rule may flag every transaction above a threshold. A machine-learning system can instead estimate the probability that a transaction is fraudulent by combining amount, device, location, merchant, sequence, customer history and network relationships. A conventional loan policy may assign fixed points for income and repayment history. An AI-assisted model can incorporate cash-flow patterns, spending behaviour and other permitted data to estimate default risk with greater granularity. Similarly, a conventional customer-service process waits for the customer to complain, while predictive systems can identify probable service failure, churn risk or product need before a complaint occurs.

This shift creates three forms of value. First, AI can lower the marginal cost of analysing additional data, making mass personalisation economically feasible. Second, it can increase the frequency of risk assessment from periodic to continuous. Third, it can redirect human labour away from repetitive screening toward exception handling, judgement, relationship management and governance. These gains, however, depend on data quality and on the ability of the organisation to validate, explain and monitor model behaviour. AI does not eliminate banking judgement; it redistributes judgement between algorithms, employees, managers, auditors and regulators.

3. REVIEW OF LITERATURE

Research on AI in banking has moved from descriptive accounts of automation toward empirical studies of customer trust, employee readiness, model governance and systemic risk. Kalyani and Gupta (2023), in a systematic review of 734 relevant articles, showed that AI and machine learning had expanded across multiple banking functions, while also identifying scalability, trust, cybersecurity and implementation challenges. Their review is important because it confirms that AI adoption is not concentrated in one front-office use case; it is becoming embedded across the architecture of banking.

In the Indian customer context, Shaikh, Kumar, Mishra and Elahi (2024) surveyed 189 customers of State Bank of India, Axis Bank, Punjab National Bank and HDFC Bank. Their findings indicate that AI-enabled banking can improve accessibility, enjoyment and time efficiency, but they also report that human service remains more satisfying in important situations. This is a critical result for Indian banks. It challenges the assumption that a successful AI strategy should aim at maximum replacement of human contact. Instead, it suggests a hybrid service model in which machines handle speed-sensitive and repetitive interactions while trained employees manage ambiguity, emotion, complaints and consequential decisions.

Employee acceptance is equally important because AI systems fail if front-line and managerial users do not trust or use them appropriately. Kumar, Rani, Rani and Rani (2025) studied 389 banking employees in India's National Capital Region using a technology-organisation-environment framework. Perceived usefulness, perceived ease of use, top-management support and competitive pressure were significantly associated with behavioural intention to adopt AI. The finding implies that AI adoption cannot be treated as an IT department project. Employees need understandable tools, workflow integration, managerial sponsorship and clear evidence that the technology improves work rather than merely adding monitoring or complexity.

More recent research has placed transparency and institutional safeguards at the centre of AI adoption. Lakshmi,



Benny and Sonny (2026), using survey data from 885 employees across public and private financial institutions in India, found that perceived transparency and competence of AI systems significantly predict trust, and that trust positively affects conflict-resolution effectiveness. The regulatory environment strengthens the relationship between transparency and trust. This suggests that regulation need not be interpreted only as a constraint on innovation. Well-designed regulation can increase adoption by providing legitimacy, standards and confidence.

The policy literature reaches a similar conclusion. NITI Aayog's Responsible AI documents emphasise safety, equality, privacy, transparency and accountability (NITI Aayog, 2021a, 2021b). The Financial Stability Board (FSB, 2024) identifies operational efficiency, advanced analytics, customisation and compliance as benefits, but warns of third-party concentration, model risk, data-governance failures, cyber risk and correlated behaviour. In 2026 the FSB further proposed a set of sound practices for organisation-wide responsible AI adoption, reflecting the shift from abstract AI ethics to lifecycle governance (FSB, 2026).

Taken together, the literature establishes three gaps that remain especially relevant for India. First, there is a gap between technological capability and customer trust. Second, there is a gap between model accuracy and explainable, contestable decision-making. Third, there is a gap between isolated AI pilots and enterprise-wide governance. The present paper addresses these gaps through an integrated analysis of the Indian banking value chain rather than by examining a single tool such as a chatbot or credit model.

4. RESEARCH OBJECTIVES AND METHODOLOGY

The study has five objectives:

- (1) to examine the major applications of AI across the Indian banking value chain.
- (2) to analyse how AI changes customer experience, credit delivery, fraud management and operational efficiency.
- (3) to evaluate the implications of AI for financial inclusion and banking employment.
- (4) to identify ethical, legal, cybersecurity and model-governance risks.
- (5) to propose a practical responsible-AI framework suitable for Indian banks.

The research adopts a descriptive-analytical design and is based on secondary data. Sources include RBI speeches and policy material, NPCI information on the UPI ecosystem, the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025, NITI Aayog's Responsible AI framework, FSB guidance and peer-reviewed studies focused on banking and Indian users or employees. The method is qualitative synthesis rather than statistical estimation. Evidence is organised by banking function and evaluated using four analytical questions: What decision is AI improving? What value does that improvement create? What new risk does the use of AI introduce? What control is required to make the use case trustworthy?

This method is appropriate because the Indian AI-banking landscape is evolving rapidly and institutions disclose different levels of implementation detail. A secondary analytical study avoids inventing bank-level performance numbers and instead connects verified evidence with established banking principles. The limitation is that the study does not estimate causal effects for the entire banking sector. Its contribution is conceptual integration: it shows how multiple AI use cases collectively alter the architecture of Indian banking and why governance must be designed at the same enterprise scale.

5. WHY INDIA IS A DISTINCTIVE ENVIRONMENT FOR AI-DRIVEN BANKING

AI adoption in Indian banking is shaped by a combination of scale, heterogeneity and digital infrastructure. India has a very large customer base that ranges from highly digital urban consumers to customers who require assisted banking in local languages. This creates an unusual design requirement: the same bank may need advanced personalisation for a mobile-first user and simplified, voice-enabled or employee-assisted services for a customer with limited digital literacy.

India's digital public infrastructure provides the second distinctive condition. Aadhaar supports digital identity; Jan Dhan widened formal account access; mobile connectivity created a distribution channel; UPI changed the economics of retail payments; the Account Aggregator framework enables consent-based financial data sharing; and emerging lending infrastructure aims to reduce information friction (Das, 2024). AI can use these digital rails to improve fraud detection, service design and credit analytics. Yet the availability of more data does not automatically justify more data use. Purpose limitation, consent, security and proportionality remain essential, particularly after the Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025 strengthened India's data-governance environment (Government of India, 2023; MeitY, 2025).

The third condition is competitive diversity. Public-sector banks, large private banks, small finance banks, payments banks, cooperative institutions and fintech partners have different legacy systems, data quality and technical capacity. Therefore, AI transformation will not occur at one speed. Larger institutions may build or fine-tune models internally, while smaller banks may depend more heavily on cloud and third-party vendors. This creates a policy challenge: the benefits of AI should not become available only to institutions with the deepest technology budgets, while the risks of vendor concentration and black-box dependence accumulate across the sector.

6. AI APPLICATIONS ACROSS THE BANKING VALUE CHAIN

The impact of AI becomes clearer when examined across the bank's end-to-end value chain. At the customer-acquisition stage, models can segment customers, predict likely product needs and improve the timing of offers. Unlike broad demographic marketing, AI allows behaviour-based personalisation. However, banks must distinguish useful personalisation from manipulative profiling. An offer that helps a



customer manage cash flow is different from exploiting behavioural vulnerability to push unsuitable credit.

During onboarding, AI-supported document processing can extract information from forms, identify inconsistencies and accelerate KYC workflows. Computer vision and pattern recognition can support identity verification and document checks, while natural language systems can guide customers through account-opening steps. The productivity benefit comes not only from speed but from reducing manual re-entry and freeing employees to investigate exceptions. The control requirement is strong because false rejection can exclude genuine customers and false acceptance can facilitate fraud.

Credit underwriting is one of the most consequential uses. AI can combine traditional bureau information with cash-flow behaviour and other legally available data to estimate repayment risk. This is particularly relevant for thin-file customers, new-to-credit borrowers, small businesses and informal enterprises whose creditworthiness is not fully represented by conventional collateral or bureau history. RBI leadership has explicitly noted the potential for alternative data to improve assessment of unbanked customers (Sankar, 2025). The inclusion opportunity is substantial, but so is the fairness risk: historical data may encode geographic, occupational, gender or income-related disadvantage. A model can be statistically accurate and still produce socially undesirable exclusion if it learns structural inequalities without appropriate governance.

In payments and transaction banking, AI can detect abnormal behaviour at a speed impossible for manual review. Fraud systems can analyse transaction sequence, device fingerprints, beneficiary networks, customer history and timing. Network models are especially useful for mule-account detection because suspicious behaviour often emerges from relationships among accounts rather than from a single transaction. MuleHunter.ai illustrates this shift from static rules toward pattern-based detection in the Indian environment (Sankar, 2025).

AI also changes collections. Rather than applying identical reminder cycles to all overdue accounts, predictive models can

estimate contactability, repayment likelihood and suitable resolution channels. A responsible system can prioritise early assistance and restructuring for temporary stress while identifying deliberate evasion. An irresponsible system could become an instrument for excessive pressure or discriminatory treatment. Consequently, collections AI must be subject to conduct rules and human review, not judged only by recovery rates.

Customer service is the most visible application. Chatbots and virtual assistants can answer routine questions at any time, retrieve transaction information, guide users through common processes and reduce call-centre workload. Generative AI can further help employees summarise conversations, search policy knowledge and draft responses. Yet a banking chatbot should not be treated like an ordinary consumer chatbot. It operates in a high-trust environment where hallucinated information about fees, loan obligations or dispute rights can cause financial harm. For this reason, high-risk responses should be grounded in approved bank knowledge, auditable and escalated to humans when uncertainty is high.

In compliance and internal audit, AI can screen large transaction datasets, identify unusual patterns, summarise regulatory changes and prioritise review queues. RegTech tools can improve suspicious-transaction monitoring and reduce the volume of false alerts. At the same time, compliance officers must be able to understand why a case was flagged and what evidence supports the alert. Automation should reduce routine workload but must not weaken professional scepticism.

Finally, AI can support treasury, liquidity, branch planning, workforce allocation, complaint analytics and management reporting. These back-office uses are less visible to customers but may generate large productivity gains. The cumulative result is that AI does not transform one department; it links previously separate activities through a common analytical layer. The strategic task for banks is therefore to build reusable data, model and governance capabilities rather than a collection of disconnected pilots.




Figure 2. AI application map across the banking value chain

Table 1. Major AI applications, value creation and principal risks

Banking function	AI application	Potential value	Principal risk
Customer service	Chatbots, GenAI assistants, intent detection	24×7 response, faster resolution	Hallucination, privacy, weak escalation
Onboarding/KYC	Document AI, vision, anomaly detection	Faster verification, fewer manual errors	False rejection, identity fraud
Credit	ML scoring, cash-flow analytics	Faster underwriting, thin-file inclusion	Bias, opacity, adverse decisions
Payments/Fraud	Behavioural analytics, network models	Real-time fraud detection	False positives, adversarial attacks
AML/Compliance	Alert prioritisation, NLP, RegTech	Lower alert burden, better surveillance	Explainability, missed suspicious cases
Operations	Process intelligence, GenAI copilots	Lower turnaround time, staff productivity	Automation bias, data leakage

7. CUSTOMER EXPERIENCE, TRUST AND FINANCIAL INCLUSION

AI can improve customer experience through speed, availability and relevance. Routine service becomes available twenty-four hours a day, waiting time can fall, and customers can receive guidance that reflects their transaction history rather than generic product categories. Shaikh et al. (2024) provide evidence from Indian customers that AI-based banking can improve accessibility and time efficiency. The same study, however, indicates that human interaction retains greater satisfaction value in important circumstances.

This finding has a direct strategic implication: the optimal model is not “digital only” but “AI first, human when needed.” Customers should move seamlessly from automated service to a skilled employee without repeating the entire problem. AI can

prepare the employee by summarising the customer’s journey, while the employee provides judgement, empathy and accountability. This hybrid design is particularly important in disputed transactions, bereavement, loan distress, fraud victimisation and complex financial advice.

For financial inclusion, AI offers two major possibilities. The first is credit inclusion through alternative-data scoring and cash-flow analysis. The second is interface inclusion through local-language, speech and simplified conversational banking. But inclusion can reverse into exclusion if models work poorly for customers who generate less digital data, speak non-standard dialects, share devices or live in areas with weak connectivity. Therefore, inclusive AI requires deliberate testing across demographic, linguistic and geographic segments. A bank should not assume that a model is fair merely because protected



characteristics are not directly used; proxy variables can reproduce similar effects.

Trust is the central outcome connecting customer experience and inclusion. A customer may accept automation for a balance enquiry but demand explanation when a loan is rejected or an account is frozen. As decision consequence increases, the need for explanation, appeal and human accountability also increases. AI maturity should therefore be measured not only by automation rates but by whether customers understand decisions and can obtain effective redress.

8. CREDIT UNDERWRITING AND RISK MANAGEMENT

Credit is where AI can create both the greatest economic value and the most serious fairness concerns. Conventional credit systems rely heavily on documented income, bureau history, collateral and predefined scorecards. These methods work well for established borrowers but can under-serve microenterprises, first-time borrowers and individuals with irregular income. AI can potentially analyse bank-account cash flows, transaction consistency, digital invoices and other permitted information to create a more complete risk picture.

For banks, improved risk differentiation can reduce adverse selection and enable more precise pricing, limits and monitoring. AI can also support early-warning systems by identifying deterioration in account behaviour before default becomes visible in conventional delinquency measures. This enables proactive contact, restructuring or risk mitigation. In corporate and MSME lending, document intelligence can reduce the time spent reading financial statements, invoices and legal documents, although final judgement for material exposures should remain subject to qualified human review.

The analytical danger is model overconfidence. A credit model can learn from historical approvals and repayments, but those records reflect earlier policies and structural inequalities. If certain groups historically received smaller loans or fewer approvals, the model may interpret the resulting data as evidence of higher risk. Explainability is therefore not a cosmetic feature. Banks need reason codes, challenger models, fairness testing and periodic validation. RBI's responsible-AI position emphasises that AI may inform decisions but accountability must remain with the institution and human actors (Sankar, 2025).

A prudent Indian approach is risk-tiered automation. Low-value, reversible decisions can be highly automated. High-impact decisions such as significant loan denial, pricing changes or account restrictions should require stronger explainability and defined human oversight. This preserves the efficiency advantage of AI without surrendering institutional responsibility.

9. Fraud Detection, AML and Cybersecurity

The rapid growth of digital payments has increased both convenience and the attack surface available to fraudsters. Traditional fraud controls depend heavily on rules: unusual amount, unusual location, excessive attempts or blacklisted beneficiary. These controls are necessary but increasingly insufficient because fraud adapts. AI can learn complex

behavioural patterns and detect combinations of signals that individually look normal.

India's experience with mule accounts illustrates the value of network intelligence. Mule accounts are used to receive, layer or transfer fraudulent proceeds. The suspicious feature may not be the absolute transaction value but the relationship among multiple accounts, rapid fund movement, beneficiary changes and abnormal behavioural sequences. RBI's MuleHunter.ai initiative reflects the use of machine learning for this problem; as reported by the RBI Deputy Governor in October 2025, the model had been deployed in about twenty commercial banks and was designed to improve precision while reducing false positives (Sankar, 2025).

AI, however, strengthens both defence and attack. Generative AI lowers the cost of creating convincing phishing messages, synthetic identities, deepfake audio or fake documents. The FSB (2024) identifies cyber risk, model risk and third-party dependencies among the important vulnerabilities associated with AI in finance. This means banks need AI-enhanced security but must also secure the AI systems themselves against data poisoning, adversarial manipulation, prompt injection, model theft and unauthorised access.

The correct objective is not simply "maximum fraud blocking." Excessive false positives can freeze legitimate payments, harm customer trust and push users away from digital channels. Banks should optimise detection using both loss prevention and customer-friction metrics. High-risk alerts need fast investigation, transparent communication and a reliable appeal process. Effective fraud AI is therefore a socio-technical system combining models, investigators, customer education, strong authentication, transaction controls and inter-institutional intelligence sharing.

10. OPERATIONAL EFFICIENCY AND THE BANKING WORKFORCE

A common narrative suggests that AI transforms banking mainly by reducing headcount. This is incomplete. The more significant change is task reallocation. Banks contain large volumes of repetitive work: data entry, document classification, reconciliations, routine queries, exception sorting and report preparation. AI can reduce the time required for these activities, but savings materialise only if processes are redesigned. Automating a poorly designed workflow may simply accelerate errors.

Generative AI introduces a new layer of employee assistance. It can summarise policies, draft routine communication, search internal knowledge, convert unstructured documents into structured information and help analysts explore scenarios. The most productive use is often "copilot" rather than autonomous decision maker. Employees remain responsible for checking the output, especially in legal, regulatory, credit and customer-facing contexts.

Employee readiness is therefore a core implementation variable. Kumar et al. (2025) found that perceived usefulness, ease of use, top-management support and competitive pressure were significant drivers of AI-adoption intention among Indian



banking employees. The implication is practical: training should be role-specific and linked to real workflow benefits. Employees need to know what the model can do, what it cannot do, when to override it and how to report problems.

AI will reduce demand for some repetitive tasks while increasing demand for model risk specialists, data engineers, cybersecurity professionals, product managers, compliance technologists and relationship roles requiring judgement. Banks that treat AI only as a cost-cutting programme may face resistance and weak adoption. Banks that treat it as a capability-building programme can combine productivity with better service quality.

11. STRATEGIC AND COMPETITIVE IMPLICATIONS

AI changes competition in banking because it increases the return to data, digital distribution and organisational learning. Institutions with richer data and better feedback loops can improve models faster, while better models can attract more usage and generate more data. This creates scale advantages. At the same time, India's interoperable public digital infrastructure can reduce entry barriers by allowing smaller institutions and fintech partners to build services without owning every layer of infrastructure.

The strategic risk is concentration in a different part of the value chain. Even if customer-facing banking remains competitive, banks may become dependent on a small number of cloud, model or specialised hardware providers. The FSB has highlighted third-party concentration as a potential financial-stability vulnerability (FSB, 2024). Vendor dependence also creates issues of data location, audit access, model transparency, business continuity and switching costs.

For Indian banks, durable competitive advantage will therefore come less from possessing a chatbot and more from combining four capabilities: high-quality governed data, strong integration with core banking workflows, disciplined model-risk management and a culture that learns from customer and operational feedback. AI features are easy to imitate; trusted execution is harder to imitate.

12. RISKS, ETHICAL CHALLENGES AND LIMITS OF AI

The first major risk is bias. AI models learn patterns from historical data, and banking history may contain unequal access, regional differences and policy effects. Even when sensitive variables are excluded, other variables can act as proxies. A fairness programme should therefore test outcomes across meaningful groups and investigate unexplained disparities. The objective is not to force identical outcomes regardless of risk, but to ensure that differences are economically relevant, legally permissible and not the product of hidden discrimination.

The second risk is opacity. Complex models may provide accurate predictions without a clear human-readable explanation. This is problematic when a customer needs to know why credit was denied, a transaction was blocked or an account was classified as high risk. Lakshmi et al. (2026) show that perceived

transparency and competence are important predictors of trust in AI within Indian financial institutions. Explainability therefore has both compliance value and operational value. A model that cannot be challenged is difficult to improve.

The third risk is privacy. AI rewards data richness, creating incentives to collect or retain more information. The Digital Personal Data Protection Act, 2023 and the Digital Personal Data Protection Rules, 2025 strengthen the need for lawful processing, clear notices, security safeguards and responsible handling of digital personal data (Government of India, 2023; MeitY, 2025). Banks should apply data minimisation even when additional variables could marginally improve model accuracy. A small accuracy gain may not justify a large privacy cost.

Fourth is model drift. Customer behaviour, fraud strategies, economic conditions and product design change over time. A model that was accurate during one period can deteriorate silently. Monitoring must therefore include stability, error rates, fairness metrics, false positives, override rates and business outcomes. Material model changes should trigger revalidation.

Fifth is automation bias. Employees may defer to AI recommendations because the system appears objective. This can be dangerous when models are wrong in unusual situations. Human oversight must be meaningful, not symbolic. An employee should have authority, information and training to challenge the system, and the organisation should study override patterns rather than penalise them automatically.

Sixth is generative-AI reliability. Large language models can generate plausible but false content. In banking, hallucination can misstate a policy, product condition or compliance requirement. For this reason, generative AI should be grounded in approved knowledge sources, constrained by role-based access and prevented from taking high-impact actions without verification. Sensitive customer data should not be exposed to uncontrolled external models.

Seventh is cybersecurity and adversarial risk. Attackers can attempt to manipulate model inputs, poison training data, extract confidential information or exploit integrated tools. AI expands the cyber threat model beyond software vulnerabilities to model behaviour and data pipelines. Security testing must therefore include red teaming of AI systems.

Finally, there is systemic risk. If many banks rely on similar models, datasets or vendors, they may respond similarly to market signals, creating correlated behaviour. The FSB (2024) identifies market correlations and service-provider concentration as potential vulnerabilities. This is why responsible AI cannot be addressed solely by individual model accuracy; supervisors must also consider common dependencies and system-wide effects.

13. REGULATORY AND RESPONSIBLE-AI GOVERNANCE IN INDIA

The direction of Indian policy is increasingly clear: innovation is encouraged, but accountability remains with the regulated institution. RBI leadership has argued for a balance between innovation and financial stability and has emphasised trust, people-first design, fairness and accountability,



understandability, safety and resilience as core responsible-AI principles (Sankar, 2025). This approach is consistent with NITI Aayog's earlier Responsible AI framework and with international financial-sector guidance.


A bank-level governance architecture should begin with an inventory of every material AI system, including vendor models and embedded AI features. Each use case should be classified by impact. A chatbot that answers general FAQs does not create the same risk as a model that rejects credit or blocks an account. Higher-impact systems should face stronger requirements for data documentation, independent validation, explainability, fairness testing, human review and board-level reporting.

Data governance is the foundation. Banks need clear ownership of datasets, data-quality controls, lineage records, retention rules, consent management and access restrictions. Model governance should document the objective, training data, assumptions, limitations, performance metrics and known failure modes. Validation should be independent of model development where practical. Once deployed, models should be continuously monitored for drift, unusual errors and changes in customer impact.

Third-party governance is equally important because many institutions will use external cloud services, APIs or foundation models. Contracts should provide audit rights, incident-notification obligations, data-protection safeguards, service-continuity arrangements and clarity on whether bank data are used to train external models. Outsourcing technology does not outsource regulatory responsibility.

The RBI's exploration of AI-enabled fraud tools and payment-intelligence infrastructure suggests that supervisors may increasingly use advanced analytics as well. This creates an opportunity for SupTech: regulators can identify anomalies, concentration risks and emerging misconduct more quickly. But supervisory AI must be governed with the same principles of data quality, proportionality and explainability expected from banks.

The broader legal environment reinforces this direction. India's data-protection framework places stronger emphasis on responsible processing of digital personal data, while RBI's KYC, fraud-risk, cyber-security and customer-protection requirements continue to apply regardless of whether a decision is made by a human or a model. AI governance should therefore be integrated into existing risk, compliance, audit and customer-redress structures instead of being managed as a separate ethical checklist.



Responsible AI is a lifecycle discipline, not a one-time model approval.

Figure 3. Responsible-AI governance cycle for banks

Table 2. Traditional banking controls versus AI-enabled banking controls

Area	Traditional model	AI-enabled model	Required control
Credit assessment	Periodic scorecard and manual review	Continuous and granular risk estimation	Explainability + human review for adverse decisions
Fraud control	Static thresholds and blacklists	Behavioural and network anomaly detection	False-positive monitoring + investigator escalation
Customer service	Branch/call-centre queues	Conversational 24x7 support	Grounded knowledge + human handoff
Compliance	Rule-based alert generation	Risk-based alert prioritisation and NLP	Audit trail + independent validation
Operations	Manual document processing	Intelligent extraction and workflow automation	Exception management + quality sampling



14. FUTURE DIRECTION OF AI IN INDIAN BANKING

The next phase of AI in Indian banking is likely to move from isolated predictive models toward integrated and increasingly agentic systems. Predictive AI answers questions such as “What is the probability of default?” Generative AI can explain or summarise information. Agentic AI may be able to plan a sequence of tasks, call tools and execute parts of a workflow. This creates large efficiency opportunities but also raises the importance of permission controls and human approval.

One likely development is the growth of multilingual conversational banking. India’s linguistic diversity creates a strong use case for voice and text interfaces that translate complex financial processes into local-language conversation. If designed well, this can reduce literacy barriers. If designed poorly, translation errors and overconfident advice can mislead vulnerable users. Banks should evaluate local-language systems for accuracy in financial terminology, not merely general language fluency.

A second direction is embedded intelligence in payments. Real-time systems can move from post-transaction fraud analysis toward pre-transaction risk scoring and adaptive authentication. RBI’s discussion of a Digital Payments Intelligence Platform indicates the policy relevance of real-time intelligence (Sankar, 2025). The challenge will be coordinating data and signals without creating excessive surveillance or friction.

Third, AI is likely to deepen cash-flow-based lending. As consent-based data infrastructure becomes more mature, lenders can evaluate small businesses and individuals using transaction evidence rather than relying only on collateral. The potential gain is faster and more inclusive credit. The governance requirement is transparent permissioning, clear adverse-action reasons and protection against excessive data extraction.

Fourth, generative AI will become a standard internal productivity layer. Bank employees may use secure assistants for policy search, code support, report drafting, complaint summarisation and document review. The main managerial question will shift from “Should employees use AI?” to “Which tasks can be safely delegated, under what controls, and how is output verified?”

Finally, responsible-AI capability itself may become a competitive advantage. The FSB’s 2026 consultation on sound practices reflects a global move toward organisation-wide lifecycle governance (FSB, 2026). Institutions that can demonstrate reliable, explainable and well-controlled AI may gain greater customer and regulatory trust than institutions that pursue rapid automation without safeguards.

A further change will be the use of AI in supervisory dialogue and internal management information. Instead of relying only on monthly or quarterly reports, banks may develop near-real-time dashboards that combine operational, conduct, cyber and model-risk indicators. This can improve early intervention, but it also creates a danger of excessive confidence in dashboards. Senior managers should retain the ability to question data definitions, model assumptions and missing information. The future bank will therefore need both technological sophistication

and stronger critical judgement. The institutions that benefit most from AI will not be those that automate the largest number of decisions, but those that know which decisions should remain contestable, reversible and accountable to people.

15. RECOMMENDATIONS FOR INDIAN BANKS

First, banks should establish an enterprise AI policy approved by senior management and linked to the existing risk-management framework. Every material model should have a named business owner, technical owner and risk owner. Accountability must remain identifiable even when a model is purchased from a vendor.

Second, use cases should be classified by customer and financial impact. High-impact applications such as credit decisions, fraud blocking, collections and account restrictions should have stricter explainability, validation and human-review requirements than low-risk productivity tools.

Third, banks should build a governed data foundation before expanding the number of AI pilots. Data quality, lineage, consent, access control and retention are more important than model novelty. Poor data combined with a sophisticated model produces sophisticated error.

Fourth, customer-facing AI should be designed around escalation. Customers must be able to reach a human when the issue is complex, disputed or consequential. The success metric should include first-contact resolution, complaint recurrence, fairness and customer trust, not only chatbot containment rates.

Fifth, banks should create independent model-validation and AI-audit capabilities. Validation should test performance, drift, fairness, explainability, cyber resilience and vendor dependence. Red-team exercises should include attempts to manipulate prompts, inputs and data pipelines.

Sixth, employee capability building should be continuous. Training should distinguish between awareness training for all staff and advanced technical training for data, risk, audit and technology teams. Front-line employees should understand when model recommendations can be challenged and how exceptions are escalated.

Seventh, financial inclusion should be a design objective. Banks should test models and interfaces across rural users, local languages, older customers, thin-file borrowers and customers with limited digital footprints. AI should reduce the cost of serving these groups rather than creating a new digital barrier.

Eighth, banks and regulators should collaborate on standards, shared fraud intelligence and safe experimentation. Regulatory sandboxes and common evaluation methods can reduce duplication while ensuring that smaller institutions are not excluded from responsible innovation.

Finally, boards should receive regular AI-risk reporting. Useful indicators include the number of high-impact models, validation status, drift incidents, customer complaints related to automated decisions, model overrides, fairness exceptions, vendor concentration, cyber events and unresolved audit findings. What gets governed at board level is more likely to be governed consistently across the institution.



16. CONCLUSION

Artificial intelligence is transforming the Indian banking sector by changing how information is converted into action. Its most important contribution is not the automation of isolated tasks but the creation of a continuous intelligence layer across customer service, onboarding, lending, payments, fraud control, compliance and operations. India's digital public infrastructure gives banks an exceptional foundation for this transformation, while the scale and diversity of the market create equally exceptional responsibilities.

The evidence reviewed in this paper indicates that AI can improve speed, accessibility, risk detection and personalisation, but it does not remove the need for human judgement. Indian customer research shows continuing value for human service, employee research highlights the importance of usefulness, ease of use and management support, and emerging institutional research demonstrates the relationship between transparency, trust and effective outcomes. These findings point toward a hybrid banking model in which AI performs high-volume analysis and employees remain responsible for judgement, empathy, escalation and accountability.

The future of AI in Indian banking will therefore be determined less by model sophistication alone and more by governance quality. Banks that combine data discipline, explainability, human oversight, privacy protection, cyber resilience and continuous validation will be able to use AI as a source of trust and inclusion. Institutions that chase automation without these safeguards may gain short-term efficiency while creating long-term operational, legal and reputational risk. The strategic objective should be responsible intelligence: using AI to make banking faster and more precise without making it less fair, less understandable or less human.

REFERENCES

1. Das, S. (2024, August 26). *Digital public infrastructure and emerging technologies* [Inaugural address at RBI@90 Global Conference]. Bank for International Settlements.
2. Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. Financial Stability Board.
3. Financial Stability Board. (2026). *Sound practices for responsible adoption of artificial intelligence (AI): Consultation report*. Financial Stability Board.
4. Government of India. (2023). *The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023)*. India Code.
5. Kalyani, S., & Gupta, N. (2023). *Is artificial intelligence and machine learning changing the ways of banking: A systematic literature review and meta analysis*. *Discover Artificial Intelligence*, 3, 41. <https://doi.org/10.1007/s44163-023-00094-0>
6. Kumar, J., Rani, M., Rani, G., & Rani, V. (2025). *Unveiling employee readiness for artificial intelligence adoption: Empirical insight from the perspective of the banking sector*. *Journal of Science and Technology Policy Management*. <https://doi.org/10.1108/JSTPM-08-2023-0130>
7. Lakshmi, K., Benny, V., & Sonny, M. (2026). *Artificial intelligence in financial institutions: Modelling transparency, competence, trust and conflict resolution in the Indian context*. *International Journal of Conflict Management*, 1-31. <https://doi.org/10.1108/IJCM-12-2025-0410>
8. Ministry of Electronics and Information Technology. (2025). *Digital Personal Data Protection Rules, 2025*. Government of India.
9. National Payments Corporation of India. (2026). *Unified Payments Interface (UPI): Product statistics*. NPCI.
10. NITI Aayog. (2021a). *Responsible AI #AIForAll: Approach document for India, Part 1 – Principles for responsible AI*. Government of India.
11. NITI Aayog. (2021b). *Responsible AI #AIForAll: Approach document for India, Part 2 – Operationalizing principles for responsible AI*. Government of India.
12. Reserve Bank of India. (2025). *Committee report on Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) in the financial sector*. Reserve Bank of India.
13. Sankar, T. R. (2025, October 7). *Responsible Artificial Intelligence (AI): Balancing innovation with financial stability* [Keynote address at the Global Fintech Festival, Mumbai]. Bank for International Settlements.
14. Shaikh, A. A., Kumar, A., Mishra, A., & Elahi, Y. A. (2024). *A study of customer satisfaction in using banking services through Artificial Intelligence (AI) in India*. *Public Administration and Policy*, 27(2), 167-181. <https://doi.org/10.1108/PAP-05-2023-0060>
15. Basel Committee on Banking Supervision. (2022). *Newsletter on artificial intelligence and machine learning*. Bank for International Settlements.
16. Bank for International Settlements. (2024). *Artificial intelligence and the economy: Implications for central banks*. *BIS Annual Economic Report 2024*.
17. Reserve Bank of India. (2025). *FAQs on Master Directions on Fraud Risk Management in regulated entities*. Reserve Bank of India.
18. Reserve Bank of India. (2016, as updated). *Master Direction – Know Your Customer (KYC) Direction, 2016*. Reserve Bank of India.
19. Approximate main-text word count (including abstract, excluding references, tables and figure captions): 6,019 words.