International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 4, July-Aug-2026, PP: 01-10

The Impact of Artificial Intelligence on Banking

Miss Kajal Yadav[1]
Assistant Professor
Department of Management and Commerce
Sam Global University[1]

Sandeep Kumar Verma[2]
Sam Global University[1]

Abstract – **Artificial Intelligence (AI) is transforming the banking industry by improving operational efficiency, enhancing customer experience, and strengthening risk management. AI-powered technologies such as machine learning, natural language processing, robotic process automation, and predictive analytics enable banks to automate routine tasks, detect fraudulent transactions, assess creditworthiness, and provide personalized financial services. AI-driven chatbots and virtual assistants offer 24/7 customer support, while advanced data analytics help banks make informed decisions and improve regulatory compliance. Despite its benefits, AI adoption also presents challenges, including data privacy concerns, cybersecurity risks, ethical issues, algorithmic bias, and the need for skilled professionals. As financial institutions continue to invest in AI technologies, it is essential to establish robust governance frameworks and ensure responsible AI implementation. This study examines the impact of Artificial Intelligence on the banking sector, highlighting its applications, benefits, challenges, and future prospects. The findings suggest that AI has become a key driver of innovation and competitiveness in modern banking, enabling financial institutions to deliver secure, efficient, and customer-centric services.**

***Keywords* - Artificial Intelligence (AI), Banking, Machine Learning, Digital Banking, Financial Technology (FinTech), Fraud Detection, Customer Experience, Risk Management**

I. INTRODUCTION

1. Background of the Study

Walk into most Indian bank branches today, or more likely, open a banking app on a phone, and it is hard to avoid some form of artificial intelligence quietly at work in the background. A chatbot answers a balance enquiry at two in the morning, an algorithm flags an unusual card transaction before the customer even notices it, and a credit decision that might once have taken a loan officer several days to work through gets a preliminary answer within seconds. This shift has happened quickly, and largely without much fanfare, but it represents one of the more significant changes in how banking actually works since the arrival of internet banking itself.

Banking, as an industry, has always been unusually well suited to automation, for a fairly simple reason: much of what a bank does involves processing information — verifying identities, assessing risk, moving money, detecting anomalies — rather than producing a physical good. This information-heavy nature made banks early and enthusiastic adopters of computing technology going back decades, from core banking systems in the 1980s and 90s to internet and mobile banking in the 2000s and 2010s. Artificial intelligence, and specifically the branch of it known as machine learning, represents the latest and in some ways most consequential step in this progression, because unlike earlier automation, which mostly followed rules a programmer wrote out explicitly, machine learning systems can identify patterns in data and improve their own performance over time without being explicitly reprogrammed for every new scenario.

India's banking sector has been a particularly active adopter of these technologies over the past several years, driven by a combination of factors: a huge and rapidly growing base of digital banking users, strong government support for digital financial infrastructure through initiatives such as the India Stack and the Unified Payments Interface, and intense competitive pressure from both established banks and newer fintech players. This report examines how artificial intelligence is actually being used in banking today, what benefits and risks this brings, and how Indian banks in particular have approached this transition.

2. Meaning and Concept of Artificial Intelligence in Banking

Artificial intelligence, in its broadest sense, refers to computer systems designed to perform tasks that would typically require human intelligence — recognising patterns, making predictions, understanding language, or making decisions under uncertainty. Within this broad field, the subset most relevant to banking today is machine learning, in which a system is trained on large volumes of historical data and learns to identify patterns or make predictions without being given an explicit rule for every situation it might encounter.

In a banking context, this translates into a fairly wide range of practical applications, several of which are examined in detail in Chapter 4 of this report: systems that can converse with customers in natural language, systems that can flag a fraudulent transaction based on subtle deviations from a customer's normal spending pattern, systems that can estimate a borrower's creditworthiness using a far wider range of data points than a traditional credit score alone, and systems that can automate the repetitive back-office processes that once required considerable manual effort, such as document verification or compliance checks.

It is worth being clear that most of what is currently deployed in banking is what researchers call narrow AI — systems trained to perform a specific, well-defined task very well, rather than anything resembling general, human-like intelligence. A fraud detection model, however sophisticated, cannot hold a conversation, and a banking chatbot, however fluent, cannot assess credit risk. This distinction matters because it shapes realistic expectations about both what AI can currently deliver for banks and what it genuinely cannot yet do.

3. Importance and Need for the Study

There are several reasons this subject deserves close attention, particularly from a management and commerce perspective. First, AI adoption is reshaping the economics of banking itself — the cost of serving a customer through an AI-powered channel is often a small fraction of the cost of serving that same customer through a branch or call centre, which has significant implications for how banks structure their operations and compete with one another, as well as with newer, digital-first fintech entrants that have built their operations around these technologies from the outset.

Second, the way AI is deployed in a sector as consequential as banking raises genuine questions of fairness, transparency, and accountability that go beyond pure business efficiency. When an algorithm rather than a human loan officer plays a significant role in deciding whether someone gets approved for a loan, questions about bias, explainability, and recourse become directly relevant — not abstract ethical concerns, but practical issues that regulators, including the RBI, are actively working through, as discussed in Chapter 6.

Third, India's own banking and fintech ecosystem has become a genuinely significant global example of AI-enabled financial inclusion, extending banking services to populations that were historically underserved by traditional infrastructure. Understanding how this has worked, and where the challenges remain, offers valuable insight not just for a banking career specifically, but for understanding how technology reshapes any information-intensive service industry more broadly.

4. Scope of the Study

This project examines the major categories of AI application in banking, weighs their benefits against the genuine risks and challenges they introduce, and looks specifically at how AI adoption has unfolded within the Indian banking sector, including the regulatory approach the RBI has taken toward this fast-moving area. It then reviews a small set of case studies — both Indian and international — that illustrate how these technologies have actually been deployed in practice.

As with the other reports in this series, the study relies on secondary data: published research, industry reports, bank disclosures, and documented case studies. Where specific figures are used to construct the charts presented in later chapters, they are illustrative in nature, intended to convey the kind of trend commonly reported in industry surveys, rather than serving as verified statistics for any specific bank or year.

5. Who is Affected by AI Adoption in Banking

The consequences of AI adoption in banking extend well beyond the technology teams that build and deploy these systems.

- Customers, who increasingly interact with their bank through AI-powered channels for routine needs, and whose access to credit or other services may be shaped, directly or indirectly, by an algorithmic decision.
- Bank employees, particularly in roles involving repetitive, rules-based work such as data entry or basic customer query handling, whose job descriptions are being reshaped as automation takes over routine tasks.
- Bank shareholders and management, who benefit from the cost efficiencies and improved risk detection that AI can offer, but who also bear responsibility for ensuring these systems are deployed responsibly.
- Regulators, who must balance encouraging beneficial innovation against protecting consumers from the genuine risks that opaque, automated decision-making can introduce.
- Society more broadly, to the extent that AI-driven credit and fraud decisions can either broaden financial inclusion for previously underserved groups or, if poorly designed, reinforce existing patterns of exclusion.

This wide reach is a central reason the subject deserves careful, balanced examination rather than either uncritical enthusiasm or blanket scepticism — both of which are easy postures to fall into when discussing a technology that is still evolving as quickly as this one.

II. OBJECTIVES AND RESEARCH METHODOLOGY

1. Objectives of the Study

The present study has been designed around the following objectives:

- To understand the meaning and scope of artificial intelligence and machine learning as applied to the banking sector.
- To examine the major categories of AI application currently in use in banking, including customer service, fraud detection, credit scoring, process automation, and robo-advisory.

- To assess the benefits that AI adoption offers banks and customers, alongside the genuine challenges and risks it introduces.
- To study how AI adoption has unfolded specifically within the Indian banking sector, and how the RBI has approached the regulation of these technologies.
- To review select case studies, both Indian and international, that illustrate practical AI deployment in banking.
- To present findings and offer suggestions relevant to how banks might approach AI adoption more responsibly and effectively.

2. Research Methodology

This is a descriptive and analytical study based on secondary sources of information. The approach followed was to first build a conceptual foundation through academic and industry literature on AI in financial services, and then to examine how these concepts have translated into practice through bank disclosures, industry surveys, regulatory publications, and documented case studies.

Where numerical illustrations were needed — for instance, to show the relative prevalence of different AI use cases across banks, or how AI adoption has trended over a recent multi-year period — representative figures have been used to construct the charts and tables in this report. These are clearly marked as illustrative throughout, reflecting typical patterns reported in industry literature rather than verified statistics for any specific institution.

3. Sources of Data

The study draws on the following broad categories of secondary sources:

- Textbooks and academic literature on financial technology and machine learning applications in finance.
- Publications of the Reserve Bank of India, including reports and circulars touching on digital banking, fintech, and data governance.
- Industry reports from consulting firms and financial technology research organisations on AI adoption in banking globally and in India.
- Publicly available disclosures and communications from banks regarding their AI initiatives, used to construct the case studies in Chapter 7.
- Academic journal articles on algorithmic fairness, explainability, and credit risk modelling.

4. Limitations of the Study

A few limitations deserve mention. The study is based entirely on secondary, publicly available information; no primary survey of bank employees or customers was conducted, given the constraints of a single academic semester. As a result, the illustrative figures presented in the charts and tables of this report reflect generalised, representative patterns rather than audited data specific to any single bank.

Artificial intelligence is also a fast-moving field in every sense — the specific tools, techniques, and even the terminology in common use continue to evolve rapidly, meaning some of the material referenced here may be superseded by newer developments by the time this report is read. Given the page limit of this project, deeper technical detail on how specific machine learning algorithms work mathematically has also been kept at a conceptual level, appropriate for a management and commerce audience, rather than developed with full technical rigour.

III. REVIEW OF LITERATURE

The application of artificial intelligence to banking sits at the intersection of several distinct bodies of literature — computer science research on machine learning, financial literature on credit risk and fraud, and a growing body of work specifically on the ethics and governance of algorithmic decision-making. This chapter briefly reviews the key threads that inform the discussion in the chapters that follow.

1. Early Automation in Banking

Long before the current wave of AI-specific interest, banking had already been through several distinct waves of automation. Core banking systems, which digitised the basic record-keeping of customer accounts, became widespread through the 1980s and 1990s, followed by the growth of ATM networks and, from the mid-2000s onward, internet and mobile banking. This history matters because it shows that banks did not suddenly discover the value of automation with the arrival of AI — rather, AI represents a continuation and intensification of a decades-long trend, distinguished mainly by its ability to handle tasks that involve judgement or pattern recognition, rather than only tasks that follow simple, fixed rules.

2. The Rise of Machine Learning in Financial Services

The specific application of machine learning to financial services accelerated markedly from around the mid-2010s onward, driven by a combination of factors: the increasing availability of large volumes of digital transaction data, substantial improvements in computing power (particularly the graphics processing units originally developed for video games, which turned out to be well suited to the calculations machine learning requires), and refinements in the underlying algorithms themselves, including the deep learning techniques that now underpin many modern natural language and pattern-recognition systems.

Academic and industry literature from this period documents a fairly consistent pattern: financial institutions were among the earliest and most enthusiastic corporate adopters of these techniques outside of the technology sector itself, motivated by the fact that finance is fundamentally a data-driven industry where even small improvements in prediction accuracy — in fraud detection

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 4, July-Aug-2026, PP: 01-10

or credit risk assessment, for instance — can translate into meaningful financial impact at scale.

3. Literature on AI and Credit Risk

A significant strand of literature has focused specifically on the use of machine learning in credit risk assessment and lending decisions. Traditional credit scoring relied on a relatively narrow set of variables — income, existing debt, repayment history — combined through comparatively simple statistical models. Machine learning-based approaches can incorporate a much wider range of data, including alternative data sources such as utility bill payment history or, more controversially, digital footprint and behavioural data, and can identify more complex, non-linear relationships between these variables and the likelihood of default.

This literature has produced genuinely encouraging findings around financial inclusion — machine learning-based credit scoring has, in a number of documented cases, been able to extend credit responsibly to borrowers who would have been rejected or simply invisible under traditional scoring models, particularly in markets like India with large populations who lack extensive formal credit histories. At the same time, a parallel and equally important strand of this literature has raised serious concerns about the risk that these models, if trained on historical data reflecting past discriminatory lending patterns, can inadvertently learn and perpetuate those same biases, even without explicitly using protected characteristics such as gender or religion as inputs.

4. Literature on AI Ethics and Explainability

This concern over bias feeds into a broader and increasingly influential body of literature on AI ethics and explainability in financial services. A central technical challenge highlighted in this literature is that many of the most powerful machine learning models — particularly deep learning models — function as what researchers often call “black boxes”: they can produce highly accurate predictions without offering an intuitive, human-understandable explanation of exactly why a particular decision was reached.

This creates a genuine tension for a regulated industry like banking, where a rejected loan applicant has a reasonable expectation of understanding why their application was declined, and where regulators need some ability to audit decision-making processes for fairness. In response, a specialised field known as “explainable AI” has developed, focused on techniques that can make complex model decisions more interpretable without necessarily sacrificing predictive accuracy — a genuinely difficult trade-off that remains an active area of both academic research and regulatory debate.

5. Indian Perspectives on AI in Banking

Within India, a growing body of research and regulatory commentary has examined how AI adoption interacts with the country's specific financial inclusion goals and its rapidly evolving digital public infrastructure, including the Unified Payments Interface and the broader India Stack framework. This literature often frames AI adoption in Indian banking somewhat differently than in more mature, saturated Western banking markets — less as a tool for optimising an already digitised customer base, and more as a genuine enabler of extending formal financial services to previously underserved populations, a theme this report returns to in Chapter 6.

The RBI itself, along with research bodies such as the Institute for Development and Research in Banking Technology (IDRBT), has published increasingly detailed guidance and discussion papers on responsible AI adoption in Indian financial services, reflecting a regulatory environment that is actively working to keep pace with a fast-moving area of technology, a dynamic examined in more detail in section 6.2 of this report.

IV. KEY APPLICATIONS OF AI IN BANKING

Artificial intelligence has found its way into nearly every major function of a modern bank. This chapter walks through the five application areas most commonly discussed in both industry practice and academic literature.

1. Customer Service: Chatbots and Virtual Assistants

AI-powered chatbots and virtual assistants are perhaps the most visible face of AI in banking today, handling routine customer queries — balance enquiries, transaction history, basic troubleshooting — without requiring a human agent. Modern banking chatbots use natural language processing to understand customer queries phrased in ordinary conversational language, rather than requiring customers to navigate rigid menu structures, and many are now capable of handling reasonably complex, multi-step conversations, including initiating simple transactions directly within the chat interface.

Beyond simple query handling, some banks have begun deploying more sophisticated virtual assistants capable of proactive engagement — for instance, alerting a customer to an unusual spending pattern or suggesting a relevant financial product based on their transaction history. This category has grown to become one of the most widely adopted AI applications in banking globally, given its relatively straightforward return on investment: a chatbot that successfully resolves even a modest share of routine queries can meaningfully reduce call centre volume and cost.

2. Fraud Detection and Prevention

Fraud detection is widely regarded as one of the areas where AI has delivered the clearest and most measurable value in banking. Traditional, rules-based fraud detection systems flag transactions that violate predefined rules — for example, a transaction above a certain amount, or one occurring in an unusual location. Machine learning-based systems, by contrast, can build a detailed behavioural profile of each individual customer's normal spending patterns and flag transactions that deviate meaningfully from that specific pattern, even if the transaction would not have triggered any generic, fixed rule.

This individualised approach allows AI-based fraud detection systems to catch a wider range of fraudulent activity while simultaneously reducing the number of legitimate transactions incorrectly flagged as suspicious — a persistent problem with older rules-based systems that could frustrate genuine customers with unnecessary transaction declines. Because these models can also be updated continuously as new fraud patterns emerge, they tend to adapt more quickly to evolving fraud tactics than static, rules-based approaches ever could.

3. Credit Scoring and Loan Underwriting

As discussed in the literature review in Chapter 3, machine learning-based credit scoring can incorporate a considerably wider range of data than traditional models, including alternative data sources, and can identify more nuanced relationships between a borrower's characteristics and their likelihood of default. In practice, this has allowed some lenders, particularly newer digital lenders and non-banking financial companies, to extend credit to borrowers with limited formal credit history — a segment traditional credit scoring models tend to underserve, given their heavy reliance on existing credit bureau data that this segment often simply lacks.

This application area is also, however, one of the most closely scrutinised from a regulatory and ethical standpoint, precisely because of the fairness and explainability concerns discussed in Chapter 3. Most responsible deployments in this area now pair the machine learning model with some form of explainability layer, allowing the lender to provide applicants with at least a general sense of which factors most influenced their decision, even if the full technical workings of the underlying model remain complex.

4. Process Automation and Back-Office Operations

Away from the customer-facing side of banking, AI and related automation technologies have also been widely deployed to streamline back-office operations — document verification during account opening, compliance checks required under anti-money laundering regulations, and the reconciliation of transaction records, among other tasks. Optical character recognition combined with machine learning, for example, can extract and verify information from identity documents far faster and, in many cases, more accurately than manual review.

This category of application, while less visible to customers than a chatbot or a declined fraudulent transaction, often delivers some of the largest efficiency gains for banks, since these back-office processes are frequently high-volume, repetitive, and prone to human error when handled entirely manually. Robotic process automation, a related but distinct technology that automates rule-based digital tasks, is often deployed alongside AI in this area, though it is worth noting that not all back-office automation actually involves genuine machine learning — some of it relies on simpler, rules-based automation that is sometimes loosely, and not entirely accurately, described as “AI” in industry marketing material.

5. Robo-Advisory and Wealth Management

Robo-advisory platforms use algorithms to provide automated investment recommendations and portfolio management, typically at a considerably lower cost than a traditional human financial advisor. These platforms generally begin by assessing a customer's risk tolerance and financial goals through a structured questionnaire, and then construct and periodically rebalance an investment portfolio, often built from low-cost index funds or exchange-traded funds, according to that assessed risk profile.

While robo-advisory has grown rapidly in more mature wealth management markets, particularly in the United States, its adoption in India has been comparatively more gradual, reflecting both a market that has historically favoured relationship-based, human advisory services and continuing efforts to build broader investor trust in fully automated investment recommendations. That said, several Indian financial institutions and fintech platforms have introduced robo-advisory features in recent years, and this segment is generally expected to grow further as digital investment platforms become more familiar to a wider base of Indian retail investors.

6. Comparative Overview and Illustrative Data

Table 1: Summary of Key AI Applications in Banking

Application	Primary Purpose	Typical Impact
Chatbots & Virtual Assistants	Handle routine customer queries and basic transactions	Reduced call centre volume; faster response times
Fraud Detection	Identify unusual transaction patterns in real time	Higher fraud catch rates; fewer false declines
Credit Scoring	Assess borrower creditworthiness using wider data	Faster decisions; expanded access for thin-file borrowers
Process Automation	Streamline document checks and compliance workflows	Lower operating costs; reduced manual error
Robo-Advisory	Automated, algorithm-driven investment recommendations	Lower-cost access to portfolio management services

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 4, July-Aug-2026, PP: 01-10

Figure 4.1 presents illustrative survey-style data on how commonly different AI applications are reported as being in active use among banks. Fraud detection and customer service chatbots tend to be reported most frequently, reflecting both their relatively mature technology and the clearer, more immediate return on investment they typically offer compared with newer application areas such as robo-advisory.

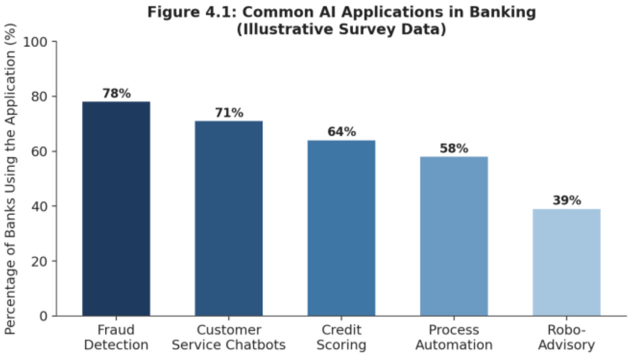

Figure 1: Common AI Applications in Banking (Illustrative Survey Data)

V. BENEFITS AND CHALLENGES OF AI IN BANKING

1. Benefits of AI Adoption

The benefits banks report from AI adoption span several dimensions. Cost efficiency is typically the most immediately visible: automating routine customer queries and back-office processes can substantially reduce the operating cost of serving a customer, particularly at the scale large banks operate. Improved risk detection is a closely related benefit, since AI-based fraud and credit risk models can generally identify patterns of risk that would be difficult or impossible for a human analyst, or a simpler rules-based system, to detect reliably across millions of transactions.

Beyond cost and risk, AI adoption also supports faster, more convenient customer service — a chatbot can respond instantly at any hour, and an AI-assisted loan decision can often be delivered in minutes rather than days. Finally, and particularly relevant to the Indian context discussed in Chapter 6, AI-enabled credit scoring can meaningfully expand financial inclusion, extending credit access to individuals and small businesses who lack the kind of extensive formal credit history that traditional scoring models require.

2. Challenges and Risks

Against these benefits, AI adoption in banking introduces a set of genuine, non-trivial challenges. Algorithmic bias, discussed at length in Chapter 3, remains a serious concern: models trained on historical data can inadvertently learn and reproduce past patterns of discriminatory lending, even without being given any information about a protected characteristic directly. The "black box" nature of many of the more powerful machine learning models also creates a genuine explainability challenge, complicating both regulatory compliance and a bank's ability to give customers a meaningful explanation for an automated decision.

Data privacy and security represent another significant challenge, since AI systems, particularly those used for credit scoring and fraud detection, typically require access to large volumes of sensitive personal and financial data, raising the stakes considerably if that data is mishandled or breached. There is also a workforce transition challenge: as routine tasks become increasingly automated, banks face a genuine responsibility to retrain and redeploy affected employees into higher-value roles rather than treating automation purely as a cost-cutting exercise, an approach that risks both reputational damage and a genuine loss of institutional knowledge and skilled staff.

Finally, over-reliance on automated systems carries its own risk: an AI model is only as good as the data it was trained on and the assumptions built into its design, and a model that performs well under normal conditions can behave unpredictably during genuinely novel situations — a concern that echoes, in a more contemporary form, some of the risk measurement limitations discussed in the earlier report on financial risk management, where statistical models similarly struggled to anticipate genuinely extreme, unprecedented scenarios.

3. Regulatory and Ethical Considerations

Regulators around the world, including the RBI in India, have been actively working to develop frameworks that allow AI's genuine benefits to be realised while managing its risks responsibly. Common regulatory themes across different jurisdictions include requirements for model explainability in credit decisions, fairness testing to detect and correct algorithmic bias before deployment, robust data privacy protections, and clear frameworks for human accountability, ensuring that a bank cannot simply attribute a harmful decision entirely to "the algorithm" without meaningful human oversight and responsibility.

This regulatory landscape continues to evolve rapidly, reflecting the genuinely difficult balance regulators must strike: overly restrictive rules risk stifling innovation that could offer real benefits, particularly in a country like India where AI-enabled financial inclusion has shown genuine promise, while overly permissive rules risk allowing the kind of biased or opaque decision-making that erodes public trust in the financial system. Section 6.2 of this report examines the RBI's specific approach to this balance in more detail.

4. Comparative Overview and Illustrative Data

Table 2: Benefits and Corresponding Challenges of AI Adoption in Banking

Dimension	Benefit	Corresponding Challenge
Cost	Lower operating costs through automation	Upfront investment and ongoing maintenance costs
Risk Management	Improved fraud and credit risk detection	Risk of model bias and unpredictable edge-case behaviour
Customer Experience	Faster, round-the-clock service	Reduced human touch for complex or sensitive needs
Financial Inclusion	Expanded credit access for thin-file borrowers	Risk of algorithmic exclusion if models are poorly designed
Workforce	Freed capacity for higher-value employee roles	Need for genuine retraining and redeployment investment

Figure 5.1 presents an illustrative breakdown of how banks typically allocate their AI-related investment across different functional areas. Customer experience and risk/fraud management together tend to account for the majority of reported AI investment, broadly consistent with these being the application areas where the business case for AI adoption has historically been clearest and easiest to measure.

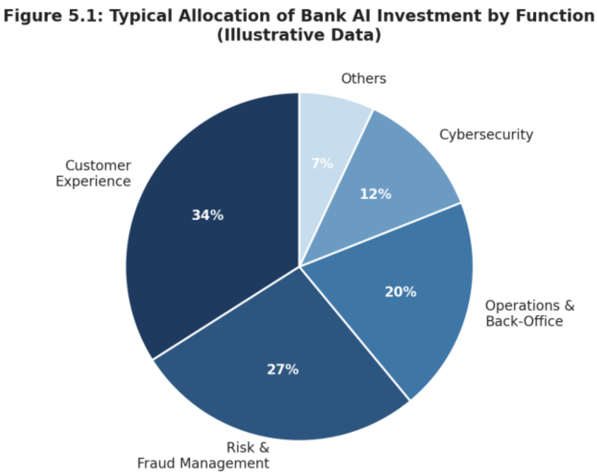

Figure 2: Typical Allocation of Bank AI Investment by Function (Illustrative Data)

VI. AI ADOPTION IN THE INDIAN BANKING SECTOR

1. Overview of AI Adoption Trends in India

India's banking sector presents a particularly interesting case study in AI adoption, shaped by a distinctive combination of factors: an enormous and rapidly digitising customer base, strong government-backed digital infrastructure through the India Stack and the Unified Payments Interface, and a competitive landscape that includes both large established banks and an active, well-funded fintech sector, all pushing AI adoption forward from different directions.

Large private sector banks have generally led AI adoption in India, deploying chatbots, AI-based fraud detection, and increasingly sophisticated credit scoring models across their retail operations. Public sector banks, which collectively still hold a majority share of Indian banking assets, have generally adopted these technologies somewhat more gradually, though this gap has been narrowing steadily as digital transformation has become a stated strategic priority across the sector, often accelerated by government-led digital banking initiatives.

2. RBI's Approach to AI and Fintech Regulation

The Reserve Bank of India has taken an active, if carefully calibrated, approach to overseeing AI and broader fintech innovation in Indian banking. Rather than issuing a single comprehensive AI-specific regulation, the RBI has generally addressed AI-related risks through a combination of existing frameworks — data protection and localisation requirements, outsourcing and third-party risk management guidelines (relevant since many banks rely on external technology vendors for AI capabilities), and fair lending principles that apply regardless of whether a lending decision is made by a human or an algorithm.

The RBI has also supported responsible innovation through its regulatory sandbox framework, which allows fintech firms and banks to test new technologies, including AI-based products, in a controlled environment with regulatory oversight before wider rollout. Additionally, research and advisory bodies such as the Institute for Development and Research in Banking Technology (IDRBT), which operates under RBI's guidance, have published detailed discussion papers specifically addressing responsible AI adoption in Indian financial services, reflecting a deliberate effort to build regulatory understanding and capacity alongside the technology's own rapid development, rather than regulating reactively after problems emerge.

3. Illustrative Trend Analysis

Figure 6.1 presents an illustrative view of how AI adoption might have trended across Indian banks over a recent five-year period, alongside the corresponding growth in the share of customer queries handled by AI-powered chatbots rather than human agents. The pattern shown — steady, accelerating growth in both measures — is broadly consistent with the direction reported across industry surveys and bank disclosures over this period, even though the specific figures here are illustrative rather than drawn from any single verified source.

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 4, July-Aug-2026, PP: 01-10

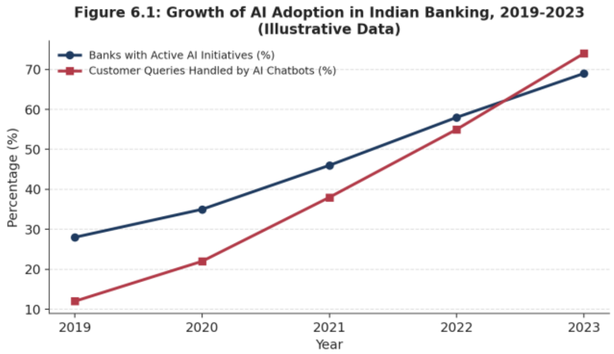

Figure 3: Illustrative Growth of AI Adoption in Indian Banking, 2019-2023

This growth trajectory reflects several converging forces discussed earlier in this report: falling technology costs that make AI deployment more accessible even for smaller institutions, growing customer comfort with digital and AI-mediated banking channels (a shift accelerated considerably by changed customer behaviour during the COVID-19 pandemic), and intensifying competitive pressure, particularly from fintech players and newer digital-first banks that have built AI capabilities into their operations from the very beginning, rather than retrofitting these technologies onto older, legacy systems as many established banks have had to do.

VII. CASE STUDIES

1. HDFC Bank's EVA Chatbot

HDFC Bank, one of India's largest private sector banks, launched its AI-powered virtual assistant, EVA (Electronic Virtual Assistant), as one of the earlier large-scale banking chatbot deployments in India. Built using natural language processing, EVA was designed to answer customer queries drawn from a very wide knowledge base covering the bank's products and services, handling everything from basic balance and transaction queries to more detailed questions about loan products and interest rates.

What made this deployment particularly notable within the Indian context was the scale at which it operated almost immediately — handling a very large volume of customer interactions across web and messaging channels shortly after launch, demonstrating that Indian customers were willing to engage with AI-driven banking channels when the experience was well designed and genuinely useful. The initiative is widely cited in Indian fintech literature as an early proof point that AI-driven customer service could work at meaningful scale in the Indian banking context, encouraging a wave of similar chatbot deployments across the sector in the years that followed.

2. State Bank of India's YONO Platform

The State Bank of India, the country's largest public sector bank, launched YONO (You Only Need One) as an integrated digital banking platform that incorporates several AI-driven features, including personalised product recommendations, an AI-based chat assistant, and streamlined, largely automated loan application processes for select retail lending products. YONO represents a particularly interesting case because it demonstrates that large public sector banks, often assumed to lag behind private peers in technology adoption, are capable of building and scaling genuinely sophisticated digital and AI-enabled platforms.

The platform's approach to using AI for personalised product recommendations — suggesting relevant financial products to customers based on their transaction behaviour and demonstrated needs — illustrates a somewhat different application of AI than pure automation or risk detection: using data-driven insight to improve cross-selling and customer engagement in a genuinely useful way, rather than a purely transactional one. Given SBI's enormous customer base, spanning both urban and rural India, YONO's AI features also carry particular significance for financial inclusion, extending sophisticated digital banking capabilities to a very wide and diverse customer base that includes many first-time digital banking users.

3. JPMorgan Chase's COIN Programme

Moving beyond the Indian context, JPMorgan Chase's Contract Intelligence programme, generally known as COIN, offers an instructive international example of AI-driven back-office automation. Launched to automate the review of commercial loan agreements — a task that previously required legal and lending staff to spend an estimated many thousands of hours annually manually reviewing standard loan documentation — COIN uses machine learning to extract and review key data points and clauses from these documents in a fraction of the time a manual review would require.

This case is frequently cited in industry literature not just for its efficiency gains, though these were considerable, but for what it revealed about the nature of AI-driven job transformation in banking: rather than simply eliminating the roles involved, the technology was reported to have freed the legal and lending staff previously tied up in this repetitive document review work to focus on more complex, judgement-intensive tasks that genuinely benefited from human expertise. This pattern — automation reallocating human effort toward higher-value work rather than eliminating it outright — is a recurring, if not universal, theme in more thoughtful accounts of AI-driven workforce transformation in the banking sector.

4. AI-Driven Fraud Detection at a Global Card Network

A further instructive international example comes from the fraud detection systems deployed by major global payment card networks, which process an enormous volume of transactions across many countries and currencies every single day. These networks have invested heavily in machine learning-based fraud detection systems capable of evaluating the fraud risk of each individual transaction in real time, within milliseconds, factoring in hundreds of variables about the transaction, the merchant, and the cardholder's historical behaviour simultaneously.

What makes this case particularly notable is the sheer scale and speed involved: a fraud decision has to be made before a transaction is authorised, meaning the underlying AI model must deliver a highly accurate risk assessment within an extremely tight time window, repeated many thousands of times per second across the network's global operations. Industry reporting on these systems consistently highlights significant reductions in fraud losses achieved through AI adoption compared with earlier, purely rules-based fraud detection approaches, reinforcing fraud detection's position, discussed in Chapter 4, as one of the clearest and most measurable success stories for AI in the broader financial services industry.

Findings and Suggestions

Key Findings

- AI adoption in banking spans a wide range of applications — customer service, fraud detection, credit scoring, process automation, and robo-advisory — each at a different stage of maturity and each offering a somewhat different type of value to banks and customers.
- Fraud detection and customer service chatbots are the most widely and successfully deployed AI applications in banking, reflecting their relatively mature underlying technology and comparatively clear, measurable return on investment.
- AI-enabled credit scoring shows genuine promise for expanding financial inclusion, particularly in a market like India with a large population of borrowers lacking extensive formal credit history, though this application also carries the most significant fairness and explainability risks among the categories examined.
- Indian banks, across both private and public sectors, have adopted AI technologies at a meaningfully accelerating pace over recent years, supported by strong digital infrastructure and an active, competitive fintech ecosystem.
- The RBI has generally taken a calibrated regulatory approach to AI in banking, addressing risks through existing data protection, outsourcing, and fair lending frameworks rather than a single AI-specific regulation, while supporting innovation through mechanisms such as its regulatory sandbox.
- The case studies reviewed in Chapter 7 — spanning both Indian and international examples — suggest that the most successful AI deployments tend to combine genuine efficiency or risk-detection gains with thoughtful attention to customer experience and, where relevant, responsible workforce transition.
- Algorithmic bias, explainability, and data privacy remain the most significant unresolved challenges associated with AI adoption in banking, and are likely to remain active areas of regulatory and academic attention for the foreseeable future.

Suggestions

- Banks should pair AI-driven credit and lending decisions with genuine explainability mechanisms, ensuring customers can receive a meaningful, understandable explanation when an automated decision affects them significantly.
- Regular, independent bias testing of AI models used in lending and other consequential decisions should be treated as standard practice, not an occasional or optional compliance exercise, given how easily historical data can encode past discriminatory patterns.
- Banks should invest genuinely in retraining and redeploying employees affected by automation into higher-value roles, both as a matter of responsible practice and because doing so, as the JPMorgan COIN case in Chapter 7 illustrates, can preserve valuable institutional expertise rather than losing it.
- Given the scale of India's underserved and thin-file borrower population, continued responsible development of AI-enabled alternative credit scoring should remain a priority, provided it is implemented with the fairness safeguards discussed above.
- Banks and regulators should continue investing in AI-specific technical and ethical expertise, given how quickly the underlying technology continues to evolve and how much regulatory frameworks still depend on genuine, up-to-date technical understanding to remain effective.
- Smaller and public sector banks, which have generally adopted AI somewhat more gradually, would likely benefit from structured knowledge-sharing with more experienced private sector and fintech peers, to accelerate responsible adoption without having to independently rediscover lessons others have already learned.

Lessons from the Case Studies at a Glance

Table 8.1 draws together the four case studies examined in Chapter 7, summarising the core AI application involved and the main lesson each offers.

Table 3: Summary of Case Study Applications and Key Lessons

Case Study	Core AI Application	Main Lesson
HDFC Bank EVA Chatbot	Natural language customer service	Well-designed AI chat channels can achieve rapid, large-scale customer adoption
SBI YONO Platform	Personalised recommendations & automated lending	Public sector banks can build genuinely sophisticated AI-enabled platforms at scale

Case Study	Core AI Application	Main Lesson
JPMorgan Chase COIN	Automated contract review	Automation can reallocate staff to higher-value work rather than only cutting roles
Global Card Network Fraud AI	Real-time transaction risk scoring	AI can materially reduce fraud losses at a speed and scale humans cannot match

VIII. CONCLUSION

This study set out to examine how artificial intelligence is reshaping banking — the specific ways it is being applied, the genuine benefits and risks this brings, and how India's banking sector in particular has approached this transition. What emerges clearly across the chapters of this report is that AI in banking is neither the purely disruptive threat nor the unqualified solution it is sometimes portrayed as in more sensational commentary; it is, more accurately, a genuinely powerful set of tools that deliver real value when deployed thoughtfully, and real risks when deployed carelessly.

The application areas examined in Chapter 4 — customer service, fraud detection, credit scoring, process automation, and robo-advisory — each illustrate a different facet of this technology's potential, from the relatively mature and low-risk efficiency gains of chatbots and back-office automation to the considerably higher-stakes, higher-scrutiny territory of AI-driven credit decisions. The case studies reviewed in Chapter 7, spanning both Indian institutions and major international examples, reinforce a point that runs throughout this report: the most successful and durable AI deployments tend to be the ones that pair genuine technical capability with careful attention to fairness, transparency, and the human consequences of automation.

India's own experience, examined in Chapter 6, offers a particularly encouraging example of AI being deployed with financial inclusion as a genuine, central objective rather than an afterthought — a use of the technology that feels distinctly suited to a market with India's specific combination of scale, digital infrastructure, and a large underserved population. As the RBI and Indian banks continue to refine their regulatory and operational approach to these technologies, this report's central conclusion is that responsible AI adoption in banking is not simply a matter of installing capable technology, but of building the fairness, explainability, and workforce transition practices needed to deploy that technology well.
On a personal note, this project has given me a considerably deeper appreciation for just how much careful, unglamorous work — fairness testing, explainability engineering, thoughtful workforce planning — sits behind the smooth, seemingly effortless AI-powered banking experiences that customers interact with every day. It is a genuinely useful reminder that behind every impressive piece of technology lies an equally important, and often less visible, set of choices about how to deploy it responsibly.

REFERENCES

1. Bank for International Settlements (BIS). Artificial Intelligence and Machine Learning in Financial Services. Financial Stability Institute.
2. Institute for Development and Research in Banking Technology (IDRBT). Discussion Papers on Artificial Intelligence in Banking.
3. Marr, B. Artificial Intelligence in Practice: How 50 Successful Companies Used AI and Machine Learning to Solve Problems. Wiley.
4. Reserve Bank of India. Report on Trend and Progress of Banking in India (various issues). RBI Publications.
5. Reserve Bank of India. Framework for Regulatory Sandbox. RBI Publications.
6. Siddiqi, N. Credit Risk Scorecards: Developing and Implementing Intelligent Credit Scoring. Wiley.
7. Skinner, C. Digital Bank: Strategies to Launch or Become a Digital Bank. Marshall Cavendish.
8. Various industry reports on AI adoption in global and Indian banking, published by leading financial technology research and consulting organisations.
9. Various publicly available bank disclosures and communications regarding AI initiatives, consulted for the case studies presented in Chapter 7.
10. Academic journal articles on algorithmic fairness, explainable AI, and machine learning-based credit risk modelling, consulted for the literature review in Chapter 3.