

AI APPLICATIONS IN FINANCE, BANKING AND FINTECH

Abirami A¹, Aarthi Nivetha S V², Dr. P.S. Venkateswaran³

¹Student, Department of Management Studies, PSNA College of Engineering and Technology, Dindigul, Tamil Nadu, India. Email ID: abiramia25mba@psnacet.edu.in

²Student, Department of Management Studies, PSNA College of Engineering and Technology, Dindigul, Tamil Nadu, India. Email ID: aarthinivethas25mba@psnacet.edu.in

³Professor, Department of Management Studies, PSNA College of Engineering and Technology, Dindigul, Tamil Nadu, India. Email ID: venkatespsna07@psnacet.edu.in

Abstract—Artificial Intelligence is transforming financial services, delivering predictive analytics, real-time decision-making, fraud detection, credit assessment, algorithmic trading and financial advisory across banking, financial management and FinTech. Research on these applications is scattered across case studies and institutional silos, lacking a joined-up account of how AI capabilities create value across the financial ecosystem. This conceptual paper takes on the challenge by synthesizing the literature around machine learning, natural language processing, intelligent automation, generative AI and decision-support systems to propose a framework linking these capabilities to operational efficiency, risk management, customer experience, financial inclusion and strategic competitiveness, while arguing that the value of AI depends jointly on organizational readiness, data governance, regulatory environments, employee capability and customer trust, rather than just the technology.

The framework also incorporates algorithmic bias, privacy, cybersecurity and model opacity as governance concerns intertwined with technological capability. The paper provides a basis for empirical testing and practical guidance for institutions, FinTech firms and policymakers building resilient, customer-centric AI-powered financial ecosystems.

Keywords: Artificial Intelligence; Financial Services; Banking; FinTech; AI Governance; Financial Innovation.

1. Introduction

Finance has always been about information. A bank decides who to lend to by guessing who will pay back, an insurer decides a price for a policy by guessing who will claim, and a market moves because thousands of people guess differently about the same company. Whoever guesses better, and does it faster and cheaper, tends to win. It is not surprising, then, that artificial intelligence (AI) has found its way into almost every corner of the sector. What is surprising is how fast it happened. Ten years ago most branch managers would have described AI as a research topic; today the same managers rely on systems that score loan applications, flag suspicious transfers and answer customer queries at two in the morning.

Several things came together to make this possible. Payment and lending moved onto smartphones, which created a continuous stream of digital data. Cloud computing made it affordable to store and process that data. And the post-crisis regulatory climate opened space for new entrants, the so-called FinTech firms, that were willing to compete on technology rather than on branches (Arner, Barberis and Buckley, 2016). In India, the story has a particular flavour. Universal identity through Aadhaar, mass account opening under Jan Dhan and the Unified Payments Interface (UPI) built a digital base on which AI-driven services could sit (Reserve Bank of India, 2017). The Financial Stability Board (2017) noted early on that AI and machine learning were already being used for credit scoring, insurance pricing, trading, and regulatory compliance, and that the trend was likely to accelerate.

Despite this activity, the academic conversation is oddly fragmented. One group of papers studies credit models, another studies fraud, a third studies robo-advisors, and they seldom talk to each other.

A manager reading this literature would struggle to answer a simple question: what does AI actually change for a financial institution as a whole, and what must be in place for that change to be beneficial? This paper tries to answer that question.

1.1 Objectives and approach

The paper has three objectives: (i) to describe the main AI capabilities used in financial services in plain terms; (ii) to map those capabilities onto applications in banking, finance management and FinTech; and (iii) to propose a framework that explains when AI creates value and when it creates trouble. The study is conceptual. We reviewed peer-reviewed articles, regulatory and central bank reports, and industry publications, and then grouped the findings by capability, by application and by outcome. No primary data were collected, so the framework should be read as a set of propositions that future empirical work can test, not as proven results.

2. Review of Literature

2.1 Credit assessment and lending

Credit scoring is the oldest and best-documented use of statistical prediction in finance, and it is where machine learning (ML) has been studied most carefully. Khandani, Kim and Lo (2010) trained ML models on consumer data from a large commercial bank and showed that they could forecast delinquencies noticeably better than conventional scoring, which translated into meaningful savings on losses. Later work asked a more uncomfortable question: better for whom? Fuster et al. (2022), studying U.S. mortgages, found that ML raised predictive accuracy overall but that the gains were not shared equally, with some minority borrowers ending up relatively worse off. On the other side of the ledger, Jagtiani and Lemieux (2019) reported that alternative data used by a FinTech lender helped some borrowers with weak traditional scores to obtain credit at lower cost, and Berg et al. (2020) showed that simple digital footprints, such as the type of device and email provider a customer uses, can predict default about as well as a credit bureau score.

2.2 Fraud, markets and text

Fraud detection was an early success for data mining. Ngai et al. (2011) surveyed the literature and found that classification, clustering and outlier-detection techniques had been applied to bank, insurance and securities fraud, although they also noted that the field lacked common data sets, which makes comparison difficult. In asset management, Gu, Kelly and Xiu (2020) compared a wide range of ML methods for predicting stock returns and concluded that non-linear methods such as trees and neural networks outperform linear models, mainly because they capture interactions between predictors. Heaton, Polson and Witte (2017) went a step further and built "deep portfolios" using neural networks. On the text side, Loughran and McDonald (2011) showed that general-purpose dictionaries misread financial language (the word "liability" is not negative in a balance sheet), which is the reason that finance-specific language models are now standard.

2.3 Advice, fairness and explainability

Robo-advice has drawn attention because it puts a machine directly in front of a retail investor. D'Acunto, Prabhala and Rossi (2019) found that automated advice improved diversification for investors who were poorly diversified to begin with, and helped reduce some behavioural biases, although the benefit was small for those who already held balanced portfolios. Bartlett et al. (2022) examined mortgage pricing and found that minority borrowers paid more under both face-to-face and algorithmic lending, with the algorithmic discrimination being smaller but not absent. This is an important reminder that algorithms inherit the world they learn from. To respond, researchers have developed tools that explain individual predictions, notably LIME (Ribeiro, Singh and Guestrin, 2016) and SHAP (Lundberg and Lee, 2017). Cao (2022) surveys the wider field.

3. Core AI Technologies in Financial Services

"AI" is used loosely in boardrooms, so it helps to separate the actual capabilities. Table 1 lists the five that matter most for this paper, what each does, and where it typically appears in a financial institution.

Table 1. Core AI capabilities and their use in financial services

Capability	What it does	Typical use in finance
Machine learning	Learns patterns from historical data to predict outcomes such as default, fraud or price movement.	Credit scoring, fraud scoring, churn prediction, return forecasting
Natural language processing	Reads and interprets text and speech.	Chatbots, sentiment analysis of news, contract and KYC document review
Intelligent automation	Combines robotic process automation with ML to handle rule-heavy, repetitive work.	Account opening, reconciliation, loan processing, reporting
Generative AI	Produces new text, summaries or code from prompts.	Drafting customer replies, summarising research, assisting relationship managers
Decision-support systems	Brings model output together with rules and human judgement.	Credit committees, treasury dashboards, portfolio rebalancing


Two points deserve emphasis. First, these capabilities are rarely used alone; a loan journey may use NLP to read documents, ML to score the applicant and automation to update records. Second, generative AI is the newest and least settled member of the group. A system that writes convincingly is not the same as one that is correct, and that matters in a regulated industry.

4. AI in Banking

4.1 Fraud detection and anti-money laundering

Banks were among the first to adopt AI at scale, and fraud control is the clearest example. With digital payments running into millions of transactions a day, no team of humans can check each one. Instead, models score every transaction against the customer's usual behaviour: amount, time, device, location and merchant type. Anything that looks out of pattern is stopped or challenged before the money leaves. Figure 1 shows the general flow. The interesting part is the feedback loop at the bottom: analysts confirm which alerts were real, and those labels are used to retrain the model, so that it adapts as fraudsters change tactics.

Figure 1. Typical real-time fraud detection workflow in a bank (author's schematic)



The same logic supports anti-money laundering (AML) monitoring. Traditional rule-based systems generate large numbers of false alerts, and compliance staff spend much of their day clearing them. ML-based systems can rank alerts by risk and pick up unusual networks of accounts that rules would miss. The trade-off is that regulators expect the bank to explain why an account was flagged, which is harder with a complex model.

4.2 Credit underwriting

In lending, AI shortens the time between application and decision from days to minutes. Models draw on bureau history, bank statements, GST or income records and, increasingly, alternative data. For thin-file customers, such as first-time borrowers or small shopkeepers, this can be the difference between being scored and being ignored. However, as the literature in Section 2 showed, a model that is accurate on average can still be unfair to specific groups, so accuracy alone cannot be the acceptance test.

4.3 Customer service and personalisation

Virtual assistants now handle balance enquiries, card blocking, statement requests and basic product questions in several languages. The benefit is round-the-clock availability and lower cost per query. The limit is trust: customers are happy to ask a bot about a balance but often want a person when money is disputed. Good banks therefore design a clean hand-off to a human agent instead of trapping customers in the bot. Behind the scenes, recommendation models suggest products (a savings plan, an insurance cover) based on life stage and spending pattern, which, done well, feels helpful and, done badly, feels intrusive.

5. AI in Financial Management and Investment

5.1 Algorithmic trading and portfolio management

In capital markets, algorithms have executed trades for decades, but ML changed what they can learn. Instead of following fixed rules, modern models search large sets of signals, such as prices, volumes, company fundamentals, news and even satellite or web-traffic data, for patterns that predict returns. The evidence from academic work is encouraging but sobering. Gu, Kelly and Xiu (2020) show real predictive gains, yet financial markets are competitive, so any pattern that is easy to find tends to fade once many participants exploit it. Models also fail when the environment shifts, as happened in the sudden market moves of March 2020. For these reasons, most institutions treat AI as an aid to portfolio managers rather than a replacement.

5.2 Robo-advisory and wealth management

Robo-advisors ask a customer a few questions about goals, horizon and risk tolerance, and then build and rebalance a diversified portfolio, usually from low-cost funds. Their real contribution is access: advice that once needed a minimum balance and a personal relationship can now be offered on a phone to a young salaried person with a small monthly investment. D'Acunto et al. (2019) found benefits for under-diversified investors, which supports this view. The weakness is that a questionnaire cannot capture everything about a person. Hybrid models, in which an algorithm builds the plan and a human advisor talks the client through it during a market fall, seem to be the most sensible direction.

5.3 Risk management and forecasting

Risk teams use AI to estimate credit losses, stress-test portfolios, detect early-warning signals in loan books and monitor liquidity. Text analytics adds another layer: reading earnings calls, filings and news to detect changes in tone that numbers alone might not show (Loughran and McDonald, 2011). For corporate finance and treasury departments, forecasting cash flows, receivables and foreign-exchange exposure with ML has become an appealing use case because the payoff, better working-capital management, is easy to measure.

5.4 Compliance, audit and insurance

So-called RegTech tools scan transactions and communications for policy breaches, help prepare regulatory reports and check customer documents during onboarding (KYC). Auditors use anomaly detection to test whole populations of entries rather than small samples. In insurance, ML supports underwriting and claims triage, and image analysis is used to assess vehicle damage. These are unglamorous applications, but they are often where the earliest and most dependable savings are found, because the work is high in volume, rule-heavy and expensive to do manually.

6. AI in FinTech and Financial Inclusion

FinTech firms are, in a sense, AI companies with a financial licence. They have no legacy branches to protect, their data pipelines are built for machine learning from day one, and their whole pitch is speed and convenience. Payment apps use ML to block fraud in the background so that the customer sees only a smooth checkout. Digital lenders decide on small-ticket loans in minutes, often using data such as cash-flow patterns on a merchant's payment account. Insurtech firms use telematics and app data to price motor cover by how a person actually drives. Personal finance apps categorise spending and nudge users toward saving.

The strongest argument for AI in this space is financial inclusion. Roughly speaking, traditional credit systems only work for people who already have a credit history, which excludes exactly those who need a first loan. Alternative data can help. Berg et al. (2020) and Jagtiani and Lemieux (2019) both suggest that digital information can extend credit to people the old scorecards would miss. Language technology adds a second channel: voice bots in Tamil, Hindi and other Indian languages can serve customers who are uncomfortable with English-only apps. The Reserve Bank of India's working group on FinTech and digital banking (2017) explicitly linked such technologies with wider access to formal finance.

There is, however, a flip side, and it is worth stating plainly. Digital lending has also produced aggressive recovery practices, opaque pricing and over-borrowing among people who found it too easy to get a loan. Inclusion that leads a first-time borrower into a debt trap is not inclusion. AI makes it cheaper to reach customers, but it does not by itself make the reaching responsible.

Figure 2 summarises our reading of how strongly each AI capability currently supports each application area. It is a qualitative synthesis drawn from the literature reviewed, not a measured index, and it is intended to show where activity is concentrated (for example, ML in credit and fraud; NLP and generative AI in service and compliance) and where it is still emerging.

Figure 2. Qualitative map of AI capabilities against application areas (author's synthesis of the reviewed literature)

	Credit decisions	Fraud & AML	Customer service	Trading & investing	Compliance	Financial inclusion
Machine learning	Strong	Strong	Emerging	Strong	Moderate	Strong
NLP	Emerging	Moderate	Strong	Moderate	Strong	Moderate
Intelligent automation	Moderate	Moderate	Moderate	Emerging	Strong	Moderate
Generative AI	Emerging	Emerging	Strong	Moderate	Moderate	Moderate
Decision support	Strong	Moderate	Moderate	Strong	Moderate	Moderate

7. An Integrative Framework

Bringing the previous sections together, Figure 3 proposes a framework with three layers. On the left are the AI capabilities. In the middle are the three application areas where they are used. On the right are the outcomes that institutions and society care about: operational efficiency, better risk management, improved customer experience, financial inclusion and strategic competitiveness.




Figure 3. Integrative framework linking AI capabilities, applications and value outcomes

The key claim sits in the grey band at the bottom. Technology on its own does not deliver the outcomes on the right. Value depends jointly on five conditioning factors, and a weakness in any one of them can cancel out strength in the others.

Organizational readiness. Institutions with streamlined processes, clear project ownership and senior support benefit more from a given model than those that view AI as an IT experiment. **Data governance.** Models are only as good as the data they are trained on; poor quality inputs, missing consent documentation or walled-off data silos limit performance. **Regulatory environment.** Clear rules encourage investment, while regulatory uncertainty can slow deployment or push it into the shadows. **Employee capability.** Staff need to ask questions of a model's output, not just accept it, and need to feel that AI is a tool for them, not a threat. **Customer trust.** A technically perfect customer-facing service can be undermined if customers fear their data is being misused or if they believe the algorithm treats them unfairly.

From this framework we derive three propositions for future research. P1: AI adoption has a stronger impact on efficiency and risk outcomes for institutions with higher levels of data-governance maturity.

P2: The positive effect of AI on customer experience is moderated by levels of customer trust. P3: The contribution of AI to financial inclusion depends on the regulatory protections that constrain harmful lending practices. Survey research or case-studies of banks and FinTech firms could test each of these.

8. Challenges and Risks

8.1 Algorithmic bias

A model trained on historical lending decisions learns the patterns in them, including any prejudice. Even if protected variables like religion or caste are stripped away, other variables such as a pin code, a phone model or a spending pattern may act as proxies. The results of Fuster et al. (2022) and Bartlett et al. (2022) suggest this is not a mere hypothetical concern. Institutions should therefore test outcomes across customer groups before and after deployment, not just measure overall accuracy.

8.2 Privacy and data protection

AI in finance operates on personal data, from transaction histories to location and contacts. Customers often do not realise the extent of data collection or its uses. India's Digital Personal Data Protection Act, 2023 imposes obligations on

organisations that process personal data including purpose limitation and consent, and financial firms need to build these into their data pipelines.

8.3 Cybersecurity and adversarial risk

A greater digital presence equals a bigger attack surface. AI is a tool for both defenders and attackers: the latter can use it to devise convincing phishing emails, impersonate a voice, or create forged documents. Models themselves can be attacked: for example, by feeding inputs that fool a fraud-detection model. Security testing of AI systems must become as routine as that of applications.

8.4 Model opacity and explainability

Many of the most accurate models are also the least explainable. This causes real problems when a customer is denied a loan and asks why, or when a regulator demands an explanation from a bank.

Post-hoc tools such as LIME and SHAP (Ribeiro et al., 2016; Lundberg and Lee, 2017) are helpful but should not be mistaken for transparency: they provide approximations, not complete model explanation. A pragmatic approach is to match model complexity to the stakes: simpler, interpretable models for decisions that affect people's access to credit, and more complex ones where model error is cheap.

8.5 Workforce and concentration risk

Automation will shrink some roles, notably in back-office processing, while creating others in data, model risk and customer support. Banks that retrain staff early will outperform those that view the shift as pure cost-cutting. A separate concern, raised by the Financial Stability Board (2017), concerns concentration: if multiple institutions rely on the same few cloud or model providers, a single failure or shared blind spot could cascade across the system.

9. Governance and Regulation

The challenges above require more than better algorithms; they demand governance. Figure 4 illustrates a simple lifecycle that we believe institutions should follow for every material AI model.

Human accountability is placed at the centre: a named person or committee must own the model and be answerable for its performance.




Figure 4. Model governance lifecycle with human oversight at the centre Regulators are taking this approach, albeit unevenly. The Reserve Bank of India has promoted responsible digital innovation while tightening guidance on digital lending, and supervisors in other jurisdictions have issued principles on the fair and accountable use of AI. A reasonable

approach for regulators is to adopt proportional, risk-based oversight calibrated to the potential harms of a model, and to use regulatory sandboxes to test new approaches.

10. Future Directions

Three trends appear likely to gain traction within the next few years. Generative AI will transition from experiments to mainstream use in tasks such as drafting, summarising, and assisting staff, provided banks can control error rates and data leakage. Explainable and privacy-preserving techniques like federated learning will become more important as regulation tightens. And embedded finance, where credit and insurance are offered within non-financial apps, will introduce new considerations around responsibility when an algorithm goes wrong. For researchers, the most valuable next steps are empirical tests of the propositions in Section 7, ideally using Indian data.

11. Conclusion

AI has moved from the fringes to the heart of financial services. In banking it protects transactions and accelerates lending; in finance management it aids investing and risk control and supports compliance; and in FinTech it opens doors to customers that older systems missed. But the central thesis of this paper is that technology alone is not enough: whether AI delivers efficient, fair and trusted financial services depends on organizational readiness, data governance, regulation, employee capabilities and customer confidence.

For managers, the lesson is to invest in people, data quality and governance as well as models. For FinTech firms, it means embedding fairness and transparency into the product, not just the compliance function. For policymakers, it suggests encouraging innovation while demanding accountability. The framework proposed here is conceptual, and its propositions await empirical testing, but we hope it provides a common language for institutions and researchers to build customer-centric, resilient financial ecosystems enabled by AI.

References

- [1] Arner, D. W., Barberis, J. N., and Buckley, R. P. (2016). The evolution of Fintech: A new post-crisis paradigm? *Georgetown Journal of International Law*, 47(4), 1271-1319.
- [2] Bartlett, R., Morse, A., Stanton, R., and Wallace, N. (2022). Consumer-lending discrimination in the FinTech era. *Journal of Financial Economics*, 143(1), 30-56.
- [3] Berg, T., Burg, V., Gombovic, A., and Puri, M. (2020). On the rise of FinTechs: Credit scoring using digital footprints. *The Review of Financial Studies*, 33(7), 2845-2897.
- [4] Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys*, 55(3), Article 64.
- [5] D'Acunto, F., Prabhala, N., and Rossi, A. G. (2019). The promises and pitfalls of robo-advising. *The Review of Financial Studies*, 32(5), 1983-2020.
- [6] Financial Stability Board (2017). *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications*. Basel: FSB.
- [7] Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5-47.
- [8] Government of India (2023). *The Digital Personal Data Protection Act, 2023*. New Delhi: Ministry of Law and Justice.
- [9] Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223-2273.
- [10] Heaton, J. B., Polson, N. G., and Witte, J. H. (2017). Deep learning for finance: Deep portfolios. *Applied Stochastic Models in Business and Industry*, 33(1), 3-12.
- [11] Jagtiani, J., and Lemieux, C. (2019). The roles of alternative data and machine learning in fintech lending: Evidence from the LendingClub consumer platform. *Financial Management*, 48(4), 1009-1029.
- [12] Khandani, A. E., Kim, A. J., and Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking and Finance*, 34(11), 2767-2787.

- [13] Loughran, T., and McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1), 35-65.
- [14] Lundberg, S. M., and Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
- [15] Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., and Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559-569.
- [16] Reserve Bank of India (2017). Report of the Working Group on FinTech and Digital Banking. Mumbai: RBI.
- [17] Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144.
