

ARTIFICIAL INTELLIGENCE ADOPTION IN VIETNAM'S BANKING SECTOR: CURRENT PRACTICES, CONSTRAINTS, AND SOLUTIONS

Nguyen Thi Quy¹

Email: quynt@hou.edu.vn, ORCID: 0009-0001-7629-3861

Received: 06/02/2026. Revised: 07/08/2026. Accepted: 24/08/2026

DOI: 10.59266/houjs.2026.1451

Abstract: *Artificial intelligence (AI) is moving from isolated experimentation toward more embedded use in Vietnam's banking sector, but the depth of adoption and the quality of public performance evidence remain uneven. This article examines current practices, implementation constraints, and governance priorities through a structured qualitative review updated to 24 August 2026. The review combines peer-reviewed banking research, international risk-governance guidance, Vietnamese legal and supervisory documents, and recent official disclosures from major commercial banks. The evidence indicates that digital identity and biometric authentication are operational and strongly regulation-driven, while machine-learning applications in underwriting, personalization, fraud monitoring, and decision support are increasingly embedded at leading banks. Generative AI has also moved beyond experimentation in several institutions, although sector-wide maturity remains uneven. Recent disclosures provide more concrete operational metrics than earlier sources, but most remain self-reported and cannot establish causal effects on profitability or risk. The main constraints concern fragmented data and legacy integration, model and conduct risk, cybersecurity and third-party dependence, uneven governance capacity, and shortages of interdisciplinary expertise. The article applies a technology–organization–environment framework with a cross-cutting AI-risk governance layer and proposes a sequenced implementation matrix for banks and supervisors. The central implication is that the value of AI in banking depends less on acquiring models than on integrating data, controls, accountable ownership, and human oversight throughout the AI lifecycle.*

Keywords: *artificial intelligence, banking, digital transformation, model risk, data governance, Vietnam*

I. Introduction

Artificial intelligence is becoming part of the operating architecture of modern banking rather than a stand-alone technology project. Machine-learning

models support credit assessment, fraud detection, customer segmentation, and risk monitoring; natural-language systems increasingly assist both customers and employees. The business case, however, is

¹ Hanoi Open University, Hanoi, Vietnam

conditional. Model performance has little value if data are fragmented, integration is weak, outputs cannot be challenged, or the institution lacks controls over cyber, conduct, and third-party risk (Königstorfer & Thalmann, 2020; Rodrigues et al., 2022).

Vietnam is a useful setting for examining these trade-offs because digital banking has expanded rapidly while the regulatory framework has become materially more demanding. By 2025, many credit institutions were reporting that more than 90% of transactions were conducted through digital channels, and the State Bank of Vietnam (SBV) continued to position digital transformation as a sector-wide priority (State Bank of Vietnam [SBV], 2025a). Biometric authentication requirements have tightened identity controls for online transactions, while rules on personal data protection and online banking security have increased the compliance burden associated with data-intensive services (National Assembly of Vietnam, 2025a; SBV, 2024a, 2024b, 2025b). The Law on Artificial Intelligence, effective from March 2026, adds a risk-based legal framework and expressly emphasizes human responsibility and oversight (National Assembly of Vietnam, 2025b).

The empirical picture is nevertheless less straightforward than the pace of announcements suggests. Earlier Vietnamese accounts often cataloged eKYC, chatbots, robotic process automation, credit scoring, or fraud-detection tools without linking any particular use case to a dated primary source. They also tended to blur the boundary between AI, machine learning, optical character recognition (OCR), robotic process automation (RPA), and biometric technologies.

This article addresses two research questions. First, in which banking functions is AI adoption most clearly documented in Vietnam as of August 2026? Second, which technological, organizational, and regulatory constraints are most relevant to effective and responsible implementation, and how should governance responses be prioritized? The contribution is analytical rather than causal. It updates institution-level evidence, separates deployment from measurable outcomes, distinguishes AI from adjacent automation technologies, and applies a technology–organization–environment (TOE) framework with an AI-risk governance overlay. This design is narrower than an impact evaluation, but it is better aligned with what public evidence can support.

II. Literature review

2.1. AI applications, effects, and technological boundaries in banking

Research on AI in banking can be organized around three broad functions. The first is customer interaction and identity. Chatbots, virtual assistants, eKYC, and biometric authentication can reduce routine service workloads, increase availability, and make remote access more convenient. Yet automated service can also damage trust when responses are inaccurate, difficult to contest, or poorly connected to human support (Huang & Rust, 2018; Königstorfer & Thalmann, 2020).

The second function is risk analysis and decision support. Machine learning can capture nonlinear relationships and process high-dimensional, structured or unstructured data in credit scoring, transaction monitoring, fraud detection, customer care, money laundering, data analysis, customer behavior prediction,

online credit granting, and other risk applications (Nguyen & Nguyen, 2025). The literature documents substantial methodological progress, but predictive accuracy is only one part of a bank’s control problem. Explainability, bias, model drift, validation, override governance, and accountability remain central in regulated decision processes (Leo et al., 2019; Doumpos et al., 2023).

The third function is operational automation and knowledge work. RPA and OCR can accelerate repetitive back-office processes, while generative AI extends automation to document summarization, information retrieval, coding, call analysis, and conversational assistance. These applications can improve productivity, but they also create risks that are less prominent in conventional predictive models: hallucinated outputs, prompt injection, confidential data leakage, intellectual property uncertainty, and dependence on external foundation models or cloud providers (Crisanto et al.,

2024; Financial Stability Board [FSB], 2024). Bibliometric evidence also shows that AI and machine-learning research in financial services now spans fintech, risk management, anti-money laundering, customer service, and data analytics rather than a single application domain (Pattnaik et al., 2024).

2.2. Analytical framework

The analysis combines the TOE framework with an AI-risk governance layer. The technology dimension covers data quality and availability, interoperability, model performance, cybersecurity, and compatibility with legacy systems. The organizational dimension includes leadership commitment, process ownership, specialist skills, investment discipline, validation capacity, and internal control. The environmental dimension encompasses regulation, digital identity infrastructure, customer expectations, and dependence on vendors or technology partners (Tornatzky & Fleischer, 1990).




Figure 1. Operational analytical framework for AI adoption in Vietnamese banking

Note. TOE dimensions define the conditions for adoption; AI-risk governance is treated as a cross-cutting layer. The framework is applied at application-group level through deployment status, potential value, principal limitations, and governance response. Adapted from Tornatzky and Fleischer (1990) and NIST (2023).

Risk governance cuts across all three dimensions. The NIST AI Risk Management Framework emphasizes governance, mapping, measurement, and ongoing management rather than one-off technical testing (National Institute of Standards and Technology [NIST], 2023). In banking, this implies a documented inventory of material models and AI use cases, accountable ownership, independent challenge, human escalation, post-deployment monitoring, and incident response. International supervisory work reaches a similar conclusion: many AI risks are extensions of familiar model, data, conduct, cyber, and outsourcing risks, but generative AI and concentration in external providers can intensify them (Crisanto et al., 2024; FSB, 2024).

III. Research methods

3.1. Search strategy and source selection

The study uses a structured qualitative document review rather than a systematic meta-analysis. The purpose is to build a transparent, source-linked account of adoption and governance, not to estimate an average treatment effect. The final search and verification cycle was conducted from 18 to 24 August 2026, and materials published on or before 24 August 2026 were eligible.

Four source families were searched. First, peer-reviewed international studies were identified through Google Scholar and publisher or DOI records using combinations of “artificial intelligence” AND bank*, “machine learning” AND bank* AND (credit OR fraud OR risk), and (“generative AI” OR “large language model”) AND (bank* OR “financial sector”).

Second, international governance materials were searched directly on the BIS, FSB, and NIST websites. Third, Vietnamese legal and supervisory materials were reviewed on the Government and SBV portals using the terms “personal data”, “online banking security”, “biometric authentication”, and “artificial intelligence”. Fourth, institutional evidence was searched on the official investor, sustainability, and corporate-news pages of Techcombank, TPBank, Vietcombank, VPBank, and MSB using the bank name together with AI, machine learning, GenAI, chatbot, credit scoring, fraud, OCR, RPA, and biometric.

The updated search log contained 30 candidate records. Eight were excluded after source-level review: six older institutional disclosures superseded by newer primary evidence covering the same use cases, one duplicate institutional item, and one earlier bibliographic record whose title and publication details could not be reconciled. The final analytical corpus therefore comprised 22 retained sources. This count describes the bounded corpus used in this article; it is not presented as an exhaustive estimate of the global literature.

3.2. Extraction, coding, and quality assessment

Each retained source was extracted into an evidence sheet containing source type, jurisdiction or institution, banking function, technology described, deployment status, implementation date, outcome measures, governance information, principal risk or limitation, and relevance to the TOE dimensions.

Institutional disclosures were used to establish that a use case had been implemented and, where numerical outcomes were reported, to record those outcomes as self-reported institutional evidence. They were not treated as independent proof that AI caused the reported financial or operational result.

Evidence was coded into six application groups: digital identity and authentication; customer interaction; credit and portfolio analytics; fraud and transaction monitoring; process automation; and generative AI. A primary-function rule resolved overlaps. For example, eKYC, combining OCR, facial recognition, and liveness detection, was coded under digital identity; an LLM-based assistant was coded under generative AI; and RPA, combined with OCR for document handling, was coded under process automation unless identity verification was the dominant function. This prevents a single deployment from being counted as several independent applications.

Source quality was assessed on provenance, bibliographic traceability, and evidential proximity. Peer-reviewed articles, official laws and regulations, and audited or formally published annual and sustainability reports were treated as high-provenance sources. Official bank releases were accepted for deployment and self-reported outcome evidence but not for causal inference. Screening and coding were conducted by the author, followed by a second-pass claim-to-source check for every institutional metric and legal statement. Because the review has one coder, no inter-rater reliability statistic is reported; ambiguous classifications were resolved conservatively using the primary-function rule.

IV. Results and discussion

4.1. Current practices of AI adoption in Vietnamese commercial banks

AI adoption is most mature where digital transformation and regulatory requirements reinforce one another. Identity and authentication are the clearest examples. SBV measures introduced biometric verification for specified online transactions and strengthened requirements for matching customer identity information. At the same time, subsequent online banking security rules expanded the control framework for authentication and application security (SBV, 2024a, 2024b, 2025b). In this domain, OCR, facial recognition, liveness checks, and database matching are better described as operational infrastructure than experimental AI projects.

More recent institutional evidence shows a second stage of adoption: AI is being embedded in data, risk, service, and knowledge-work platforms across several large banks. Techcombank's 2025 annual report describes an enterprise-wide AI and data monetization engine and reports VND 1.3 trillion in incremental revenue delivered through advanced analytics and machine learning. It also reports a data platform that manages 5 petabytes and processes 8 billion data points per day, along with AI-powered underwriting and model-monitoring infrastructure. These are substantial indicators of scale, but the revenue figure remains institution-reported and should not be interpreted as an independently estimated causal effect of AI (Techcombank, 2026).

TPBank reports a smart-bank architecture built around AI, data, and cloud infrastructure. Its 2025 annual report states that more than 80% of systems

were integrated with AI and that 99.5% of transactions were conducted through digital channels. A 2026 institutional disclosure further identifies AI credit scoring, real-time fraud detection, limit optimization, personalization, and internal management support as active use cases. The same disclosure emphasizes that human managers remain responsible for risk oversight and final accountability, an important governance qualification to the bank's AI-native positioning (TPBank, 2026a, 2026b).

Customer service and generative AI are also becoming more concrete. Vietcombank's 2025 sustainability report states that its customer support center processed approximately 3 million requests across channels and that the bank had implemented next-generation generative-AI solutions to improve automation and task processing. The same report provides customer-satisfaction and service metrics, but it does not isolate the marginal contribution of AI to those outcomes (Vietcombank, 2026). VPBank

reports a broader ezGenAI ecosystem across customer support, collections, policy, coding, and data tasks. Its 2025 disclosure states that response time fell from about 15 minutes to 6–10 seconds, self-service exceeded 95%, and employee productivity increased by 20–30%; these figures are useful operational evidence, though they remain self-reported (VPBank, 2025).

MSB provides another indication that generative AI and intelligent document processing are moving into routine operations. In June 2026, the bank reported that its Gamma AI agent handled about 20,000 questions per month and that its intelligent document-processing system had processed more than 300,000 document sets at 600–1,800 images per second. The system was estimated to save more than 2,000 workdays annually. MSB also disclosed an enterprise AI platform and hybrid AI-cloud architecture intended to standardize and govern AI applications at scale (MSB, 2026).

Table 1. Source-linked institutional evidence of AI and AI-enabled banking applications in Vietnam

Institution	Use case	Technology	Evidence type	Status / date	Reported outcomes	Governance information	Citation
Techcombank	Personalization, underwriting, lead generation, fraud-related analytics; GenAI/agent workflows	Advanced analytics; ML; GenAI	2025 annual report	Operational enterprise-wide; GenAI scaling / 2025	VND 1.3 trillion incremental revenue attributed to advanced analytics and ML; 5 PB data platform; 8 billion data points/day. Revenue is not an independent causal estimate.	Enterprise feature store; AI/ML Ops; automated retraining, monitoring, lineage, and governance controls.	Techcombank (2026)
TPBank	Credit scoring, fraud detection, personalization, customer interaction, internal support	AI/ML; GenAI; automation	2025 annual report + 2026 bank disclosure	Operational; transition toward AI-native model / 2025–2026	More than 80% of systems reported as AI-integrated; 99.5% of transactions digital. Metrics do not isolate AI's causal contribution.	Human responsibility for risk oversight and final accountability explicitly emphasized.	TPBank (2026a, 2026b)

Institution	Use case	Technology	Evidence type	Status / date	Reported outcomes	Governance information	Citation
Vietcombank	Customer support and knowledge assistance	GenAI; service automation	2025 sustainability report	Operational/ selective / 2025	About 3 million customer-support requests processed across channels; next-generation GenAI deployed. Satisfaction metrics are reported but not attributed solely to AI.	Real-time service monitoring, information-security controls, and structured customer-quality monitoring reported.	Vietcombank (2026)
VPBank	Customer support, collections, policy, coding, data and agent-assisted tasks	GenAI; LLM-based applications	Official 2025 institutional disclosure	Operational across multiple functions / 2024–2025	Response time reported to fall from about 15 minutes to 6–10 seconds; self-service above 95%; employee productivity up 20–30%.	Common AI platform supporting multiple business functions; public model-validation detail remains limited.	VPBank (2025)
MSB	Internal AI agent, intelligent document processing, HR assistance	AI agents; GenAI; IDP/OCR	Official 2026 institutional disclosure	Operational and scaling / 2026	Gamma handles about 20,000 questions/month; more than 300,000 document sets processed; estimated annual saving above 2,000 workdays.	Hybrid AI cloud and enterprise AI platform; data control, compliance, and scalable governance emphasized.	MSB (2026)

Note. Outcome measures are reproduced only where the institution itself reported them. They are treated as self-reported evidence of deployment or operational scale, not independently validated estimates of AI's causal effect. "Status" applies to the cited use case, not to the bank as a whole.

Table 2. Application of the TOE + AI-risk governance framework by application group

Application group	Technology condition	Organizational condition	Environment / regulation	Principal governance test	Analytical maturity
Digital identity and authentication	High readiness where OCR, biometrics, and database matching are integrated.	Requires coordination across onboarding, operations, compliance, cyber, and customer support.	Strongly shaped by biometric and online-banking requirements.	Privacy, false matching, spoofing, availability, and auditability.	Operational and regulation-driven
Customer interaction	Chatbots and assistants are technically mature; answer quality and knowledge integration vary.	Requires content ownership, escalation routes, and service-quality monitoring.	Digital-service expectations interact with conduct and data-protection duties.	Accuracy, escalation, complaints, automated-interaction disclosure, and data protection.	Operational and scaling
Credit and portfolio analytics	Data quality, feature governance, explainability, and model stability are critical.	Requires business ownership, independent challenge, and override rules.	Credit and AI duties increase accountability for adverse decisions.	Bias, validation, drift, override frequency, explainability, and accountability.	Selective to advanced at leading banks; uneven sector-wide
Fraud and transaction monitoring	Real-time data integration and anomaly detection are increasingly feasible; false positives remain costly.	Requires coordination among fraud, operations, cyber, risk, and investigation teams.	Online-banking security rules and threat evolution create strong external pressure.	False positives/negatives, adversarial manipulation, incident response, and data sharing.	Selective but increasingly embedded
Process automation	RPA/OCR are mature for repetitive processes; exception handling and legacy integration remain constraints.	Value depends on process redesign and clear ownership, not automation alone.	Indirectly shaped by auditability, outsourcing, and operational-resilience requirements.	Access control, change management, exception handling, continuity, and vendor risk.	Operational in defined processes
Generative AI	Capability is advancing quickly, but hallucination, prompt injection, and data leakage remain material.	Requires approved knowledge sources, controlled access, human review, and usage monitoring.	AI law, data-protection rules, and provider dependence raise governance demands.	Grounding, output verification, leakage, IP provenance, provider concentration, and human oversight.	Selectively operational at leading banks; early and uneven sector-wide

Note. Maturity is an analytical synthesis of public evidence and does not rank individual banks. It distinguishes application readiness from governance adequacy.

4.2. *Principal constraints*

The first constraint is technological. AI performance depends on the quality, lineage, and accessibility of data, yet banks typically operate multiple product, channel, and legacy platforms built at different times. Inconsistent definitions, missing values, limited interoperability, and changing pipelines weaken reproducibility. Recent disclosures from Techcombank and MSB are informative precisely because both institutions describe substantial investment in common data and AI platforms. Their experience suggests that reliable AI at scale is an infrastructure and operating-model problem, not merely a model-development problem (Techcombank, 2026; MSB, 2026).

A second constraint is organizational. AI projects require joint ownership by business units, risk management, compliance, information security, data teams, legal functions, and internal audit. If responsibility remains concentrated within technology teams, a use case can be technically successful yet lack a clear process owner, validation standard, or measure of business value. The shortage is therefore not only one of data scientists.

The third constraint concerns model and conduct risk. Credit and fraud models can reproduce historical bias, become unstable when customer behavior shifts, or produce outputs that staff cannot adequately explain. Generative AI introduces a distinct set of concerns: plausible but incorrect responses, prompt injection, leakage of confidential information, uncertain source provenance, and overreliance on model outputs. These risks matter more in banking because decisions may affect access to credit, customer assets, complaint handling, or regulatory reporting (NIST, 2023; Crisanto et al., 2024).

Cybersecurity and third-party concentration form a fourth constraint. Advanced models are often built on external cloud, software, or foundation-model ecosystems. This can accelerate deployment but also creates dependencies that are difficult to exit, especially after applications become embedded in core workflows. The FSB identifies third-party dependence, cyber risk, model risk, and correlated use of common technologies as channels through which AI may amplify financial-sector vulnerabilities (FSB, 2024).

Finally, the regulatory environment is becoming denser. The Law on Personal Data Protection No. 91/2025/QH15, effective from January 2026, provides the current statutory framework for personal-data protection; Circular No. 50/2024/TT-NHNN and its 2025 amendment strengthen online-banking security requirements; and the Law on Artificial Intelligence No. 134/2025/QH15, effective from March 2026, introduces a risk-based framework with explicit human-centered and oversight principles (National Assembly of Vietnam, 2025a, 2025b; SBV, 2024b, 2025b).

Most of these barriers are not uniquely Vietnamese. Data quality, explainability, model drift, cyber exposure, specialist-skill shortages, and dependence on external providers recur internationally (Rodrigues et al., 2022; Crisanto et al., 2024; FSB, 2024). What is specific to Vietnam is its configuration and timing. Rapid growth in digital banking is occurring alongside heterogeneous legacy architectures, mandatory biometric controls, uneven public disclosure of model outcomes, and a newly effective AI law.

4.3. Solutions for more effective and responsible adoption

The recommendations below are analytical implications derived from the reviewed evidence and regulatory framework; they should not be read as descriptions of current practice at every bank. The priority is use-case discipline. Banks should rank proposed applications by expected value, data readiness, operational criticality, and potential harm. Low-risk internal tools can be tested under controlled access, while credit decisions, fraud controls, biometric systems, and customer-facing generative AI require stricter validation and approval. A proof of concept should proceed to scale only when it meets predefined operational, risk, and financial thresholds.

Second, data and model governance should be integrated rather than managed in parallel. Common definitions, data lineage, quality checks, approved training and reference data, accountable ownership, independent validation, version control, and drift monitoring are all necessary. Human-in-the-loop controls are especially important where outputs affect customer rights or significant financial exposure. The governance logic should be proportional: a rule-based OCR workflow does not require the same controls as an adaptive credit model or a generative-AI assistant that produces customer-facing content.

Third, cybersecurity and third-party risk require lifecycle controls. Banks need strong identity and access management, secure development, adversarial testing, incident response, and clear rules for sensitive data. Vendor due diligence should cover data location, subcontractors, model updates, audit rights, business continuity, concentration, and exit arrangements. A low subscription cost is not evidence of low total risk; switching costs and operational dependencies often appear only after the model becomes embedded in processes.

Fourth, human capital should be developed by role. Data scientists need banking and regulatory knowledge; risk and compliance staff need enough model literacy to challenge assumptions; managers need to understand when AI can support judgment and when it cannot replace it. Cross-functional product teams, combined with independent validation and audit, are more likely to convert technical capability into controlled business value than isolated technology teams.

For the SBV and related authorities, the priority is proportionate supervisory guidance. The AI Law establishes broad principles, but banks would benefit from more specific expectations regarding model risk classification, explainability, generative AI controls, outsourcing, incident reporting, and documentation. Shared fraud-intelligence mechanisms and supervisory capacity-building can also reduce duplication.

Table 3. Priority matrix for implementation

Priority action	Responsible actor	Implementation horizon	Expected indicator	Principal risk addressed
Establish an AI/model inventory and risk-tiering rule for all material use cases.	Commercial banks: Board/Risk, Model Risk, IT/Data	0–6 months	Share of material models/use cases inventoried and risk-classified; named accountable owner.	Uncontrolled model use; accountability gaps
Set minimum data-quality, lineage, validation, override, and monitoring standards before scale-up.	Commercial banks: Data, Risk, Business, Internal Audit	0–12 months	Data-quality pass rate; validation completion; drift alerts; override and exception rates.	Bias; model failure; weak reproducibility

Priority action	Responsible actor	Implementation horizon	Expected indicator	Principal risk addressed
Apply stricter approval and human-oversight rules to credit, fraud, biometric, and customer-facing GenAI.	Commercial banks: Risk, Compliance, Business owners	0–12 months	Human-review coverage; escalation rate; false-positive/false-negative trends; incident count.	Customer harm; conduct risk; hallucination
Strengthen AI-specific cybersecurity and vendor due diligence, including exit and concentration testing.	Commercial banks: CISO, Procurement, Legal, Third-party Risk	6–18 months	Share of critical vendors assessed; red-team findings closed; tested exit/continuity plans.	Data leakage; cyberattack; provider concentration
Issue proportionate supervisory guidance on model governance, GenAI, outsourcing, and incident reporting.	SBV and related authorities	6–18 months	Guidance issued; supervisory templates; reporting taxonomy; staff trained.	Regulatory ambiguity; uneven supervisory treatment
Develop shared fraud-intelligence and technical-standard mechanisms while protecting lawful data use.	SBV, banks, payment infrastructure, relevant agencies	12–24 months	Participation coverage; time to share verified fraud indicators; reduced duplicate controls.	Fragmented fraud data; systemic operational risk

Note. Horizons are sequencing priorities derived from the review, not regulatory deadlines. Indicators are illustrative management measures and should be calibrated to the materiality of each use case.

V. Conclusion

AI adoption in Vietnam’s banking sector is now more clearly documented than it was only a few years ago, but it remains uneven across functions and institutions. Digital identity and biometric authentication are operational and regulation-driven. Machine learning is increasingly embedded in personalization, underwriting, fraud monitoring, and data-driven decision support. Most importantly, evidence available by August 2026 shows that generative AI has moved from pilot status to selective operational use at several large banks. Techcombank, TPBank, Vietcombank, VPBank, and MSB all provide primary evidence of AI-enabled platforms or applications, and several disclose measurable operational outcomes.

Those disclosures should still be interpreted cautiously. They differ in scope, metric definitions, and verification, and they rarely provide comparable model accuracy, risk-adjusted profitability, or counterfactual performance measures. The evidence is therefore strong enough to establish deployment and, in some

cases, scale; it is not strong enough to rank banks or estimate the causal effect of AI on financial performance.

The study remains limited by its reliance on public secondary evidence and single-author coding. Future research should use bank-level data, expert interviews, or carefully designed case studies to test whether specific AI applications improve efficiency, risk outcomes, and customer welfare, and to distinguish self-reported deployment benefits from independently verified performance effects.

References

- Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. (2024). *Regulating AI in the financial sector: Recent developments and main challenges (FSI Insights No. 63)*. Bank for International Settlements. <https://www.bis.org/fsi/publ/insights63.htm>
- Doumpos, M., Zopounidis, C., Gounopoulos, D., Platanakis, E., & Zhang, W. (2023). Operational research and artificial intelligence methods in banking. *European Journal of Operational Research*, 306(1), 1–16. <https://doi.org/10.1016/j.ejor.2022.04.027>

- Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. <https://www.fsb.org/uploads/P14112024.pdf>
- Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172. <https://doi.org/10.1177/1094670517752459>
- Königstorfer, F., & Thalmann, S. (2020). Applications of artificial intelligence in commercial banks: A research agenda for behavioral finance. *Journal of Behavioral and Experimental Finance*, 27, Article 100352. <https://doi.org/10.1016/j.jbef.2020.100352>
- Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Risks*, 7(1), Article 29. <https://doi.org/10.3390/risks7010029>
- MSB. (2026, June 17). *GreenNode and MSB expand strategic partnership to accelerate the transition from digital banking to AI banking*. <https://www.msb.com.vn/ve-chung-toi/tin-tuc-va-su-kien/ve-msb/greennode-va-msb-mo-rong-hop-tac-chien-luoc-thuc-day-hanh-trinh-tu-ngan-hang-so-den-ngan-hang-ai/>
- National Assembly of Vietnam. (2025a). Law on Personal Data Protection No. 91/2025/QH15. <https://vanban.chinhphu.vn/?classid=1&docid=214590&pageid=27160&typegroupid=3>
- National Assembly of Vietnam. (2025b). Law on Artificial Intelligence No. 134/2025/QH15. <https://vanban.chinhphu.vn/?docid=216334&pageid=27160&typegroupid=3>
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. <https://doi.org/10.6028/NIST.AI.100-1>
- Nguyen, T. H. L., & Nguyen, T. T. H. (2025). The impact of artificial intelligence on Vietnam's banking sector. *Journal of Science, Hanoi Open University*, 220-235. <https://doi.org/10.59266/houjs.2025.697>
- Pattnaik, D., Ray, S., & Raman, R. (2024). Applications of artificial intelligence and machine learning in the financial services industry: A bibliometric review. *Heliyon*, 10(1), Article e23492. <https://doi.org/10.1016/j.heliyon.2023.e23492>
- Rodrigues, A. R. D., Ferreira, F. A. F., Teixeira, F. J. C. S. N., & Zopounidis, C. (2022). Artificial intelligence, digital transformation and cybersecurity in the banking sector: A multi-stakeholder cognition-driven framework. *Research in International Business and Finance*, 60, Article 101616. <https://doi.org/10.1016/j.ribaf.2022.101616>
- State Bank of Vietnam. (2024a). *Banking technology highlights in 2024: Security and safety in the provision and use of banking services in cyberspace*. https://www.sbv.gov.vn/vi/web/sbv_portal/w/sbv622854
- State Bank of Vietnam. (2024b). Circular No. 50/2024/TT-NHNN on security and confidentiality in the provision of online banking services. <https://chinhphu.vn/?classid=1&docid=211772&orggroupid=4&pageid=27160>
- State Bank of Vietnam. (2025a, May 26). *Banking sector digital transformation event 2025*. <https://www.sbv.gov.vn/w/sbv637421>
- State Bank of Vietnam. (2025b). Circular No. 77/2025/TT-NHNN amending Circular No. 50/2024/TT-NHNN on security and confidentiality in the provision of online banking services. <https://chinhphu.vn/?docid=216580&pageid=27160>
- Techcombank. (2026). *Annual report 2025*. <https://techcombank.com/content/dam/techcombank/public-site/documents/techcombank-2025-annual-report-eng.pdf>

- Tornatzky, L. G., & Fleischer, M. (1990). *The processes of technological innovation*. Lexington Books.
- TPBank. (2026a). *Annual report 2025*. <https://tpb.vn/wps/wcm/connect/44f827ae-2699-4b05-8dd8-f7807530515d/AR%2B2026%2BENG%2B22.4%2BVIEW%2B%281%29.pdf>
- TPBank. (2026b, May 27). *Inside Vietnam's pioneering AI-native bank*. <https://tpb.vn/eng/tin-tuc/tin-tpbank/ben-trong-ngan-hang-viet-tien-phong-mo-hinh-ai-native>
- Vietcombank. (2026). *Sustainability report 2025*. https://www.vietcombank.com.vn/-/media/Project/VCB-Sites/VCB/Nha-Dau-tu/Bao-cao-PTBV/ENG_VCB_ESG-Report-2025.pdf
- VPBank. (2025). *VPBank recognized as an outstanding digital transformation enterprise at the Vietnam Digital Awards 2025*. <https://www.vpbank.com.vn/ve-chung-toi/giai-thuong/2025/vpbank-duoc-vinh-danh-doanh-nghiep-chuyen-doi-so-xuat-sac-tai-vnda-2025>

ỨNG DỤNG TRÍ TUỆ NHÂN TẠO TRONG NGÀNH NGÂN HÀNG VIỆT NAM: THỰC TIỄN TRIỂN KHAI, RÀO CẢN VÀ GIẢI PHÁP

Nguyễn Thị Quý¹

Tóm tắt: Trí tuệ nhân tạo (AI) đang chuyển từ các thử nghiệm đơn lẻ sang những ứng dụng tích hợp sâu hơn trong ngành ngân hàng Việt Nam, song mức độ triển khai và chất lượng bằng chứng công khai về hiệu quả vẫn chưa đồng đều. Bài viết thực hiện tổng quan tài liệu định tính có cấu trúc, cập nhật đến ngày 24/8/2026, kết hợp các nghiên cứu học thuật, tài liệu quản trị rủi ro quốc tế, văn bản pháp lý và giám sát của Việt Nam, cùng các công bố chính thức gần đây của các ngân hàng thương mại. Kết quả cho thấy định danh số và xác thực sinh trắc học đã ở giai đoạn vận hành và chịu tác động mạnh từ yêu cầu pháp lý; các ứng dụng học máy trong thẩm định tín dụng, cá nhân hóa, giám sát gian lận và hỗ trợ ra quyết định đang được tích hợp sâu hơn tại một số ngân hàng lớn. AI tạo sinh cũng đã chuyển từ thử nghiệm sang triển khai có chọn lọc trong thực tế, mặc dù mức độ trưởng thành trên toàn ngành còn không đồng đều. Các rào cản chính gồm dữ liệu phân tán và khó tích hợp hệ thống cũ, rủi ro mô hình và hành vi, an ninh mạng và phụ thuộc bên thứ ba, năng lực quản trị chưa đồng đều và thiếu nhân lực liên ngành. Bài viết vận dụng khung công nghệ–tổ chức–môi trường kết hợp lớp quản trị rủi ro AI xuyên suốt, đồng thời đề xuất ma trận ưu tiên triển khai cho ngân hàng và cơ quan giám sát. Kết luận quan trọng là giá trị của AI trong ngân hàng phụ thuộc không chỉ vào năng lực mô hình mà còn vào dữ liệu, cơ chế kiểm soát, trách nhiệm giải trình và giám sát của con người trong toàn bộ vòng đời AI.

Từ khóa: trí tuệ nhân tạo, ngân hàng, chuyển đổi số, rủi ro mô hình, quản trị dữ liệu, Việt Nam

¹ Trường Đại học Mở Hà Nội, Hà Nội, Việt Nam