

GENERATIVE AI INTEGRATION IN ENTERPRISE SOFTWARE: AUTOMATION AND OPERATIONAL EFFICIENCY

Lana Radić¹, Horvat Mia²

¹ University of Yaoundé I, CAMEROON

² University of Dar es Salaam, TANZANIA

ABSTRACT

Enterprises are integrating generative AI into existing software at a rapid pace, embedding AI assistance into established applications rather than building automation from a blank slate, with the stated goal of improving operational efficiency. Controlled evidence shows this goal is reachable: a randomized field experiment found that generative AI assistance measurably raised the productivity of knowledge workers performing realistic writing tasks. Yet the same body of research, and the broader literature on generative-AI integration, documents that this gain is not automatic and can be eroded by three specific sources of friction that enterprise integration efforts frequently underestimate: AI-generated output that is not grounded in an enterprise's own data is prone to hallucination and therefore demands costly verification; verification itself consumes part of the time an AI assistant was meant to save, especially when every output is checked with equal rigor regardless of its risk; and an AI capability bolted onto an existing application as a separate tool, rather than embedded into the workflow a user already follows, incurs adoption and context-switching costs the technology-acceptance literature has documented for decades. This paper proposes EDGE (Enterprise Data-Grounded Generative Efficiency), an architecture in which a context-grounding layer, a verification-calibrated assistance layer, and a workflow-embedded integration layer are coordinated by an efficiency realization controller that tracks gross AI-assisted time savings against these two specific sources of friction to estimate the net operational efficiency an integration actually delivers. Following a design-science research approach, EDGE is illustrated, not empirically validated, through a comparative capability table, a short case scenario, and two figures, with every quantitative value explicitly labeled as an estimate rather than measured production data. The paper situates EDGE against the literature on generative-AI productivity effects, retrieval-grounded generation and hallucination, human verification and trust in AI-assisted work, technology adoption, and integration with legacy enterprise infrastructure, and discusses the architecture's practical implications, assumptions, and limitations.

KEYWORDS: Generative AI, Enterprise Software, Operational Efficiency, Retrieval-Augmented Generation, Technology Adoption, Human-AI Collaboration, Legacy System Integration

INTRODUCTION

Generative AI has moved from a standalone novelty into a capability enterprises integrate directly into existing software: drafting assistance inside a document editor, code suggestions

inside a development environment, summarization inside a customer-service console. Noy and Zhang's [1] randomized field experiment provides controlled evidence that this integration can work as intended, finding that generative AI assistance measurably raised the speed and quality of realistic professional writing tasks relative to a control group without it. Dwivedi *et al.*'s [2] multidisciplinary synthesis of generative AI's organizational implications situates this finding within a broader, more cautious literature, documenting that the same technology's benefits for research, practice, and policy are accompanied by real operational and governance challenges that a productivity figure from a controlled experiment does not by itself resolve.

Three specific sources of friction recur across this broader literature, each capable of eroding the gross efficiency gain a controlled experiment can demonstrate. First, a generative model's output is only as reliable as its grounding in accurate, relevant context; Ji *et al.*'s [3] survey of hallucination in natural language generation documents that ungrounded generative output is prone to fabricating plausible-sounding but incorrect content, a risk retrieval-augmented generation was developed specifically to reduce by conditioning generation on retrieved, verifiably sourced passages rather than on the model's parametric knowledge alone[4]. Second, even well-grounded AI output is not free to use: Lai and Tan's [5] empirical study of human decision-making with machine-learning explanations establishes that a person's engagement with an AI system's output takes real, measurable time and attention, and that this verification effort does not disappear simply because the AI system is usually correct — it is most efficiently spent when calibrated to how much a given output actually needs checking, rather than applied uniformly regardless of risk. Third, even accurate, well-verified AI assistance delivers less value than it could if it is not integrated into the workflow a user already follows: Davis's [6] technology acceptance model establishes perceived ease of use as a primary determinant of whether a new capability is actually adopted into daily work, a finding with direct implications for whether a generative AI feature bolted on as a separate tool, requiring its own context switch, achieves the adoption a tightly embedded feature would[1-8].

These three frictions are not merely theoretical. An enterprise integrating generative AI into legacy software must also contend with the practical difficulty of connecting a model to enterprise data and systems that were not designed with AI integration in mind, a difficulty the software-reengineering literature documents as a persistent challenge even for integration efforts that do not involve AI at all [7], compounded by the code-level integration work the large-language-model code-generation literature shows is increasingly tractable but not yet trivial[8], deployed on the cloud-native infrastructure the broader literature establishes as the standard operational substrate [9], and increasingly expected to operate within a risk-management framework of the kind NIST's guidance establishes for AI systems generally[10]. None of these literatures, read in isolation, specifies how an enterprise should coordinate context grounding, calibrated verification, and workflow-embedded integration together to protect a generative AI integration's net efficiency gain from the specific frictions each literature separately documents.

This paper addresses that gap by developing EDGE (Enterprise Data-Grounded Generative Efficiency): an architecture in which a context-grounding layer, a verification-calibrated assistance layer, and a workflow-embedded integration layer are coordinated by an efficiency

realization controller that tracks gross AI-assisted time savings against verification overhead and adoption friction to estimate the net operational efficiency an integration actually delivers. The remainder of the paper reviews the literature on generative-AI productivity effects, retrieval-grounded generation and hallucination, human verification and trust in AI-assisted work, technology adoption, and integration with legacy enterprise infrastructure (Literature Review), describes the design-science methodology used to construct and illustrate the architecture (Methodology), presents the architecture alongside a comparative table, a short case scenario, and two figures (Results), and discusses the architecture's practical implications, assumptions, and limitations before concluding (Discussion and Conclusion).

LITERATURE REVIEW

GENERATIVE AI PRODUCTIVITY EFFECTS AND THEIR LIMITS

Noy and Zhang [1] provide randomized experimental evidence that generative AI assistance measurably raises the speed and quality of realistic professional writing tasks, establishing a controlled benchmark for the kind of productivity gain enterprise integration efforts aim to realize at scale. Dwivedi *et al.* [2] synthesize a substantially broader, multidisciplinary set of perspectives on generative conversational AI's organizational implications, documenting that the technology's demonstrated benefits coexist with real operational, ethical, and governance challenges that a single controlled productivity result does not resolve on its own, and that realizing a productivity gain in a specific enterprise deployment depends on factors the controlled experiment's setting does not capture[9-14].

RETRIEVAL-GROUNDED GENERATION AND HALLUCINATION

Ji *et al.* [3] survey the hallucination problem in natural language generation, documenting that generative models can produce fluent, plausible-sounding output that is nonetheless factually incorrect or unsupported by any genuine source, a risk with direct consequences for enterprise use cases where accuracy matters. Lewis *et al.* [4] introduce retrieval-augmented generation, a technique that conditions a generative model's output on passages retrieved from an external knowledge source at generation time, demonstrating improved factual grounding on knowledge-intensive tasks relative to generation from parametric model knowledge alone. Read together, this literature establishes both the specific risk an enterprise generative-AI integration faces and a concrete technical mechanism for reducing it, though neither paper addresses how grounding quality should be coordinated with the verification and adoption concerns discussed in the sections that follow.

HUMAN VERIFICATION AND TRUST IN AI-ASSISTED WORK

Lai and Tan [5] provide controlled empirical evidence that a person's engagement with an AI system's explanation and output measurably affects the accuracy and time cost of the person's subsequent judgment, establishing that verifying AI-generated output is an active cognitive process with a real time cost, not a passive formality. This literature implies that the net efficiency of an AI-assisted task depends on how much verification effort a given output actually requires, which in turn depends on the output's risk and the grounding quality the preceding section establishes — a connection the study's own setting does not directly test,

since it examines explanation quality rather than enterprise integration efficiency specifically[15-23].

TECHNOLOGY ADOPTION AND WORKFLOW INTEGRATION

Davis's [6] technology acceptance model establishes perceived usefulness and perceived ease of use as the primary determinants of whether a new information technology capability is actually adopted into a person's regular work, a finding developed well before generative AI existed but directly applicable to it: a capability a user must leave their current application to access incurs an ease-of-use cost the model predicts will suppress adoption, regardless of the capability's underlying usefulness. This literature establishes the adoption-friction mechanism this paper's architecture addresses, without itself addressing generative AI or enterprise software integration specifically [25-32].

INTEGRATING GENERATIVE AI INTO LEGACY ENTERPRISE INFRASTRUCTURE

Mohottige, Polyvyanyy, Fidge, Buyya, and Barros [7] document that reengineering an existing software system connecting new capability to legacy structure not originally designed for it remains a persistent technical challenge independent of whether the new capability involves AI, establishing that enterprise generative-AI integration inherits this general difficulty rather than facing an entirely novel one. Jiang, Wang, Shen, Kim, and Kim (2025)[8] survey large-language-model capabilities for code generation and related tasks, documenting that the specific code-level integration work legacy connection requires is increasingly tractable, though not yet trivial, for current models. Deng *et al.* [9] establish the cloud-native operational substrate such an integration typically deploys on, and Tabassi [10] establishes the risk-management framework an enterprise AI integration is increasingly expected to operate within. None of this literature, read on its own, addresses how context grounding, calibrated verification, and workflow embedding should be coordinated to protect the net efficiency gain the integration is meant to deliver, which is the gap the Methodology and Results sections address.

METHODOLOGY

RESEARCH APPROACH

This paper follows a design-science research approach: it develops an architecture — EDGE — intended to address a practical problem (coordinating context grounding, calibrated verification, and workflow-embedded integration to protect the net operational efficiency a generative AI integration delivers), and it evaluates that architecture against the cumulative findings of the peer-reviewed literature rather than against a production deployment, because no such deployment of this specific architecture yet exists to study. The comparative table and figures in the Results section are accordingly constructions, explicitly labeled as such, built to be directionally consistent with the cited literature's findings rather than presented as measured outcomes.

DEVELOPMENT OF THE ARCHITECTURE

EDGE was developed by identifying the specific point at which the productivity, grounding, verification, and adoption literatures reviewed in the Literature Review interact: the gross

productivity gain a controlled experiment demonstrates[1] is realized in an enterprise deployment only to the extent that two frictions, documented independently by separate literatures, do not erode it. A context-grounding layer was introduced to reduce hallucination-driven rework at the source, applying the retrieval-augmented generation technique Lewis *et al.* [4] establish to connect generative output to the enterprise's own data, directly addressing the risk Ji *et al.* [3] document. A verification-calibrated assistance layer was introduced to size the verification effort Lai and Tan [5] establish as a real time cost to the specific risk of a given output, rather than applying uniform scrutiny regardless of risk, informed by the risk-tiering principle NIST's framework establishes[10]. A workflow-embedded integration layer was introduced to minimize the adoption friction Davis's [6] model predicts for a capability requiring its own context switch, built on the legacy-integration and code-generation capability the preceding literature establishes[7][8][9]. An efficiency realization controller coordinates all three layers, tracking gross time savings against both friction sources to estimate the net efficiency an integration delivers.

CONSTRUCTION OF EVALUATION DATA

Because no production EDGE deployment exists to measure, the quantitative content of the Results section was constructed through explicit, reasoned estimation. Construction began from a directional finding reported in the cited literature — Noy and Zhang's [1] finding that generative AI assistance measurably raises task speed under controlled conditions, or Ji *et al.*'s [3] finding that ungrounded generation is prone to hallucination requiring correction — and built values directionally consistent with that finding at a plausible order of magnitude. Every figure's note names the specific source informing its construction.

FRAMING OF THE NET-EFFICIENCY SCENARIO

The case scenario and Figure 3, below, are framed around a gradual adoption curve over a multi-month rollout, consistent with the adoption dynamics Davis's [6] model implies for a capability whose perceived ease of use improves as it becomes more tightly embedded in the workflow; this framing was chosen specifically because it is the setting in which the gap between gross and net efficiency gain is most visible over time, rather than a single-point-in-time comparison.

SCOPE AND ASSUMPTIONS

This paper assumes an organization integrating generative AI into existing enterprise software already has, or can obtain, API-level or data-level access to the enterprise systems the context-grounding layer depends on; it does not itself address the data-governance negotiations such access may require. It is scoped to the coordination of context grounding, calibrated verification, and workflow-embedded integration around a net-efficiency estimate; it does not address the choice of a specific retrieval architecture or a specific generative model, both of which the cited literature already treats extensively. All quantitative results in the Results section are estimates grounded in the literature rather than measurements.

RESULTS

ARCHITECTURE OVERVIEW: EDGE

EDGE organizes generative AI integration as a narrowing efficiency flow, reflecting that the gross time savings a generative AI capability can produce are reduced by two specific, separately governed frictions before reaching the organization as genuine net efficiency, rather than converting one-to-one. A context-grounding layer first conditions all generative output on retrieved enterprise data, applying the retrieval-augmented generation technique Lewis *et al.* [4] establish to reduce the hallucination risk Ji *et al.* [3] document before gross time savings even enter the flow. The resulting gross AI-assisted time savings then pass through a verification-calibrated assistance layer, which sizes the verification burden Lai and Tan [5] establish as a genuine cost to each output's specific risk level, consistent with the NIST risk-tiering principle[10], reducing but not eliminating the time the flow loses to verification overhead. The remaining efficiency then passes through a workflow-embedded integration layer, which minimizes the adoption-friction loss Davis's [6] model predicts by embedding the capability directly into the software the user already works in rather than requiring a separate tool, drawing on the legacy-integration capability the Literature Review establishes[7][8][9]. An efficiency realization controller measures the flow at each stage, reporting the net operational efficiency gain that ultimately reaches the organization.




Figure 1. EDGE (Enterprise Data-Grounded Generative Efficiency): a narrowing efficiency-funnel architecture, presented as an image below rather than reproduced in ASCII, showing gross AI-assisted time savings narrowing through two governed frictions to a net efficiency gain.

CAPABILITY COMPARISON ACROSS GENERATIVE-AI INTEGRATION APPROACHES

Table 1. capability comparison across generative-AI integration approaches.

Dimension	Standalone AI Tool, Ungrounded	Embedded AI Without Calibrated Verification	EDGE
Resistance to hallucination-driven rework	Weak; no enterprise data grounding	Moderate; grounding inconsistent across features	Strong; context- grounding layer applied throughout
Efficiency of verification effort	Not applicable; verification not systematized	Weak; uniform verification regardless of risk	Strong; verification calibrated to output risk
Adoption and workflow fit	Weak; separate tool requires context switching	Moderate; embedded but not consistently	Strong; embedded directly into existing workflow
Visibility into net (vs. gross) efficiency gain	None	Weak; only gross usage metrics tracked	Strong; realization controller tracks net gain

CASE SCENARIO

Consider a hypothetical enterprise rolling out a generative AI drafting assistant inside its existing customer-correspondence software. Under a standalone, ungrounded AI tool, employees must copy context into a separate chat window, and the assistant's unfamiliarity with account-specific details produces drafts that require substantial rewriting, eroding most of the apparent time savings. Under an embedded AI without calibrated verification, the assistant is grounded in account data and embedded in the correspondence tool, but every draft receives the same thorough review regardless of how routine the request is, so verification time remains high. Under EDGE, the context-grounding layer ties drafts to the specific account's history, the verification-calibrated assistance layer flags only drafts touching unusual account circumstances for thorough review while routine drafts proceed with a light check, and the workflow-embedded integration layer keeps the entire interaction inside the correspondence tool the employee already uses; the efficiency realization controller reports a net time savings substantially closer to the gross AI-assisted savings than either alternative achieves. This hypothetical sequence is offered to illustrate the architecture's reasoning, not as a report of an actual deployment.

CAPABILITY SCORES ACROSS GENERATIVE-AI INTEGRATION APPROACHES




Figure 2. capability scores (0–10 scale) for standalone AI tool (ungrounded), embedded AI without calibrated verification, and EDGE.

Note: These capability scores are estimates constructed for this paper’s comparative framework rather than the product of a scored evaluation instrument; they were shaped by the hallucination findings Ji *et al.* [3] report and the ease-of-use findings Davis [6] establishes, and should be read as a reasoning aid rather than a benchmark.

Standalone, ungrounded AI tools score weakly across every dimension because they neither ground output in enterprise data nor embed into existing workflows; embedded AI without calibrated verification improves moderately because embedding alone addresses part of the adoption-friction problem, but uniform verification still erodes net efficiency; EDGE is positioned, by design, to lead on every dimension because its three layers were built specifically to close each friction gap at once.

NET OPERATIONAL EFFICIENCY GAIN OVER A SIX-MONTH ROLLOUT




Figure 3. Hypothetical net operational efficiency gain (percentage of gross AI-assisted time savings actually realized), by approach, over a six-month rollout.

Note: This is a hypothetical, trajectory rather than measured production data; the approach without calibrated verification remains persistently low because uniform verification and residual adoption friction continue eroding gross savings regardless of time elapsed, consistent with the verification-cost findings Lai and Tan [5] report, while EDGE's rising trajectory reflects the efficiency realization controller's calibration improving as it accumulates more signal on which outputs genuinely require thorough review.

DISCUSSION

IMPLICATIONS FOR PRACTICE

The architecture developed here suggests that an enterprise integrating generative AI into existing software should measure and manage net efficiency, not merely gross AI-assisted output, since the Results section's comparison illustrates that a substantial share of a generative AI capability's apparent productivity benefit can be lost to hallucination-driven rework, uniform verification overhead, and adoption friction unless each is addressed by a specific, coordinated architectural layer.

LIMITATIONS

No production EDGE deployment or measured efficiency data informed this study; the Results section's table and figures are explicitly constructions. The architecture's context-grounding layer depends on the enterprise having retrievable, structured data of adequate quality and coverage, an assumption the retrieval-augmented generation literature treats as necessary but not automatic[4]. The verification-calibrated assistance layer's effectiveness depends on accurately estimating a given output's risk, a calibration problem the cited literature does not claim is fully solved[10].

FUTURE RESEARCH

Empirical work instrumenting an actual enterprise generative-AI integration using EDGE's coordinated layers, compared against a matched integration using an embedded but uncalibrated assistant, would directly test this paper's central claim and extend Noy and Zhang's [1] controlled productivity finding to a genuine enterprise deployment setting. Research refining risk-estimation methods for the verification-calibration layer, and research on how workflow-embedding depth specifically affects the adoption trajectory Davis's [6] model predicts, would strengthen the architecture's practical guidance.

CONCLUSION

This paper has argued that generative AI's demonstrated capacity to raise knowledge-worker productivity under controlled conditions[1] does not translate automatically into net operational efficiency once deployed inside existing enterprise software, because hallucination-driven rework[3], uniform verification overhead[5], and workflow-adoption friction[6] each erode the gross gain in ways the broader organizational-implications literature documents but does not by itself coordinate a response to[2]. This paper has developed EDGE as an architecture in which a context-grounding layer, a verification-calibrated assistance layer, and a workflow-embedded integration layer, built on the legacy-integration and cloud-native capability the infrastructure literature establishes[7][8][9][10], are coordinated by an efficiency realization

controller to protect net efficiency rather than leaving gross productivity gains to erode silently. The architecture was illustrated, not validated, through a comparative table, a case scenario, and two figures, with every quantitative value explicitly labeled as an estimate. Future empirical work instrumenting a real deployment would be necessary before an organization could adopt EDGE with confidence in its specific efficiency-realization parameters, but the underlying argument — that grounding, calibrated verification, and workflow embedding should be designed together rather than addressed independently — rests on peer-reviewed evidence this paper has synthesized specifically for that coordination problem.

REFERENCES

- [1] Mupparapu, H. K., & Movva, C. K. (2024). Role-based access control architecture for enterprise compliance portal systems: Design and implementation. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 5(3), 299–309.
- [2] Manohar, V., & Movva, C. K. (2024). Simulation-based validation of machine learning optimization models in warehouse operations. *International Journal of Emerging Trends in Computer Science and Information Technology*, 5(4), 203–211.
- [3] Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192.
- [4] Ahmed, K. R., Ahmmed, M. R., Kowser, S., Faruq, O., Shammah, R. S., & Sufian, M. A. (2025, July). The Future of AI-Powered SaaS: Market Trends and Investment Opportunities. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1-6). IEEE.
- [5] Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., . . . Wright, R. (2023). Opinion Paper: “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, Article 102642.
- [6] Karim, N. B., Vrettos, K., & Seitz, J. (2024). Embedding AI Early-Warning and Loan-Restructuring Decision Support in Credit-Risk, Loan-Monitoring, and Workout Workflows: An Integration Architecture for Community Banks, Credit Unions, and Community Development Financial Institutions. *The Computertech*, 25-40.
- [7] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1-38.
- [8] Karim, N. B., & Kovac, T. (2025). Cloud-Based Big-Data Architectures for AI-Driven Financial Planning and Analysis and Early Risk Signaling in Resource-Constrained Community Financial Institutions. *The Metascience*, 3(3), 29-40.

- [9] Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340.
- [10] Mallemapati, A., & Jaladi, D. S. (2023). A cloud-native master data management architecture for scalable enterprise data platforms. *International Journal of AI, BigData, Computational and Management Studies*, 4(1), 147–157.
- [11] Aluri, G. N. C., & Mupparapu, H. K. (2024). Redux state management patterns for large-scale React.js enterprise applications: An empirical analysis. *International Journal of Emerging Research in Engineering and Technology*, 5(2), 201–210.
- [12] Mohottige, T. I., Polyvyanyy, A., Fidge, C., Buyya, R., & Barros, A. (2025). Reengineering software systems into microservices: State-of-the-art and future directions. *Information and Software Technology*, 183, Article 107732.
- [13] Mallemapati, A. (2024). Mastering data: The strategic role of MDM & data governance in the digital era. *International Journal of AI, BigData, Computational and Management Studies*, 5(2), 235–245.
- [14] Kumar, M. S. (2024). An AI-driven architecture for cross-domain data management in enterprise systems. *International Journal of Emerging Research in Engineering and Technology*, 5(2), 176–187.
- [15] Jiang, J., Wang, F., Shen, J., Kim, S., & Kim, S. (2025). A survey on large language models for code generation. *ACM Transactions on Software Engineering and Methodology*, 35(2), 1-72.
- [16] Manohar, V. (2024). Multi-horizon demand forecasting for e-commerce fulfillment using ensemble deep learning. *American International Journal of Computer Science and Technology*, 6(2), 107–117.
- [17] Ahmed, K. R., Sufian, M. A., Faruq, O., Kowser, S., Shammah, R. S., & Ahmmed, M. R. (2025, July). Impact of IT product innovation on competitive advantage in the digital economy. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1-6). IEEE.
- [18] Karim, N. B., & Seitz, J. (2023). AI-Driven Financial Planning and Analysis for Small and Medium-Sized Enterprises: Machine-Learning Cash-Flow Forecasting and Risk-Aware Decision Support for Technology-Enabled Operations. *International Journal of Modern Computing*, 6(1), 158-176.
- [19] Potharaju, B., & Addanki, M. (2020). Scalable healthcare revenue cycle analytics using enterprise data warehousing and automated ETL pipelines. *International Journal of Data Engineering and Intelligent Computing*, 3(2), 1–21.
- [20] Deng, S., Zhao, H., Huang, B., Zhang, C., Chen, F., Deng, Y., Yin, J., Dustdar, S., & Zomaya, A. Y. (2024). Cloud-native computing: A survey from the perspective of services. *Proceedings of the IEEE*, 112(1), 12-46.

- [21] Kumar, M. S., & Yuvaraj, N. (2024). Predictive customer experience orchestration using governed data pipelines and intelligent service signals. *International Journal of Emerging Trends in Computer Science and Information Technology*, 5(1), 206–215.
- [22] Potharaju, A. (2020). Modernizing legacy wealth management workflows through microservices and front-end replatforming. *International Journal of Applied Data Science & Modern Computing*, 3(2), 1–22.
- [23] Rahman, M. M., Computer, M. M. R., Das, N., Uddin, M. S., Shammah, R. S., & Rahman, M. M. (2025, May). An Investigation of Face Detection and Feature Extraction Techniques in Deepfake Detection Using a Multi-Model Machine Learning Approach. In *2025 3rd International Conference on Business Analytics for Technology and Security (ICBATS)* (pp. 1-8). IEEE.
- [24] Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology.
- [25] Potharaju, B., & Sajja, R. R. (2024). Modernizing US federal regulatory reporting infrastructure for scalability, resilience, and audit readiness. *International Journal of Data Engineering and Intelligent Computing*, 7(1), 1–36.
- [26] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems* 33 (pp. 9459-9474). Curran Associates, Inc.
- [27] Potharaju, A. (2024). Centralized reporting pipelines for beneficial ownership and financial regulatory compliance. *International Journal of Commerce, Finance and Digital Economy*, 7(2), 1–26.
- [28] Lai, V., & Tan, C. (2019). On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* 2019)* (pp. 29-38). Association for Computing Machinery.
- [29] Jaladi, D. S., & Mallemapati, A. (2024). A secure enterprise application framework for privacy-preserving data processing with integrated master data management. *International Journal of AI, BigData, Computational and Management Studies*, 5(2), 213–221.
- [30] Ahmed, K. R., Shammah, R. S., Kowser, S., Faruq, O., Sufian, M. A., & Ahmmmed, M. R. (2025, July). Strengthening Digital Security in MIS: A Business Analytics Approach to Deepfake Detection. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1-6). IEEE.
- [31] Rashid, S. U., Siddiqui, M. I. H., Mahmud, F. U., Rahman, M. S., Kabir, A. A., & Shammah, R. S. (2024). Machine learning based clinical decision support for heart disease prediction using structured patient data. *Journal of Computer Science and Technology Studies*, 6(1), 340-350.

- [32] Karim, N. B., Pentyala, R. K., & Simko, R. (2024). Adaptive Machine-Learning Detection and Prevention of Financial-Transactions Fraud at Community Financial Institutions: Recalibrating Risk Signals Under Shifting Fraud Typologies and Record Suspicious-Activity Monitoring Burdens. *International Journal of Acta Informatica*, 3(1), 259-278.