International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

Explainable Generative AI-Driven Bank Credit Risk Modeling: Secure Apache–SAP HANA Cloud Integration for Threat-Focused Data Analytics

Samuel Richard Donovan

Senior SAP Consultant, Amsterdam, Netherlands

ABSTRACT: The growing complexity of financial risks and the rapid evolution of cyber threats demand advanced analytical frameworks capable of supporting secure, transparent, and real-time decision-making in the banking sector. This paper presents an Explainable Generative AI–driven framework for credit risk modeling that integrates Apache security tools with SAP HANA's in-memory cloud architecture to provide high-performance, threat-aware data analytics. The proposed system enhances credit scoring accuracy through generative modeling while ensuring regulatory compliance by incorporating explainability, traceability, and bias detection mechanisms. Additionally, the architecture leverages Apache-based threat monitoring and secure data pipelines to mitigate vulnerabilities across the cloud infrastructure. Experimental evaluation demonstrates improved predictive performance, reduced false positives in risk classification, and strengthened data security in high-volume banking environments. The framework provides a scalable and interpretable foundation for next-generation autonomous risk management in cloud-enabled financial systems.

KEYWORDS: Explainable Generative AI, Bank Credit Risk Modeling, SAP HANA Cloud, Apache Security Tools, Threat-Aware Analytics, Secure Data Pipelines, Financial Risk Intelligence

I. INTRODUCTION

The financial services industry increasingly relies on machine learning for credit underwriting, fraud detection, and portfolio management. Traditional risk models (logistic regression, scorecards) are being complemented or replaced by more expressive ML methods that ingest diverse, high-velocity data sources (transaction streams, device telemetry, behavioral timelines). While these models can improve accuracy, they often behave as "black boxes"—raising regulatory, ethical, and operational concerns. Explainability is no longer optional: regulators and internal risk teams demand traceable, auditable model decisions that can be justified to customers and examiners. Simultaneously, real-time operationalization requires streaming infrastructures and in-memory analytics platforms for sub-second decisioning. Finally, the sensitivity of credit data and the financial incentives for attackers necessitate a security posture that combines preventive defenses (encryption, access control) with active threat analytics that can detect adversarial probing or data leakage. This work proposes a systems-level framework that synthesizes explainable generative AI (XGenAI) methods with an Apache-based streaming backbone and SAP HANA–optimized serving and explainability components deployed securely in the cloud.

We position generative models as twofold enablers for credit risk: (1) controlled synthetic augmentation to mitigate data imbalance (rare default events) and to create realistic stress-scenario examples for model robustness testing; and (2) compact latent representations that expose semantically meaningful features for downstream explainability. Coupled with discriminative estimators optimized for explainability-constrained deployment, the hybrid pipeline achieves both predictive performance and transparency. The architectural choice—Kafka + Flink for streaming, SAP HANA for in-memory storage and model lifecycle—is motivated by enterprise readiness, low latency, and HANA's growing support for explainability primitives in its PAL/HANA-ML ecosystem. Finally, the framework embeds a cloud threat analytics module that monitors telemetry, model queries, and unusual feature distributions, enabling rapid detection and mitigation of adversarial or insider threats. This integrated approach addresses the triple imperative of accuracy, explainability, and security for credit risk decisioning in modern cloud platforms. (Kai Waehner)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

II. LITERATURE REVIEW

Explainability in credit risk modeling has attracted substantial attention: studies show that post-hoc model-agnostic methods (SHAP, LIME) and counterfactual explanations produce narratives that are more usable for loan officers and regulators than raw feature importances alone. Misheva et al. (2021) and subsequent reviews document practical XAI applications for Lending Club and similar datasets, highlighting tradeoffs between fidelity and human interpretability. Research into model-anchored explainability (rule extraction, surrogate models) demonstrates how complex models can be approximated by human-readable artifacts for audit trails. Concurrently, generative modeling for tabular financial data — using VAEs and conditional GANs — has been proposed as a remedy for class imbalance, enabling synthetic minority oversampling while preserving statistical dependencies.

On the systems side, streaming platforms such as Apache Kafka and Flink are now standard for real-time financial pipelines; their low-latency processing capabilities make them ideal for dynamic feature computation and rapid retraining triggers. Integration of streaming with in-memory databases (SAP HANA) offers operational benefits: HANA's native ML libraries and Python client (hana-ml) simplify model deployment and permit explainability hooks close to the serving layer. SAP technical articles and community posts (2023–2024) describe global explanation capabilities and PAL explainability extensions that align well with enterprise audit requirements. Security literature and cloud provider guidance emphasize layered controls for SAP landscapes and identify common attack vectors (misconfigured IAM, exposed audit logs, inadequate encryption). Recent cloud security assessments of SAP deployments further show the importance of combining inherited CSP controls with application-level monitoring. Together, this body of work supports an architecture that fuses XAI, generative augmentation, streaming feature computation, in-memory serving, and dedicated threat analytics to produce a compliant, practical credit risk platform. (IDEAS/RePEc)

III. RESEARCH METHODOLOGY

1. Problem definition and objectives — Define the operational credit decision problem (instant PD scoring for retail/SME loans) with constraints: latency < 200ms for decision path, regulatory explainability (per-decision local explanations plus global model documentation), and security SLAs (encryption, SSO, alerting). Establish evaluation metrics: AUC/KS for predictive power, explanation fidelity (surrogate model R^2), explanation usability via human evaluation, privacy leakage metrics (membership inference risk), and security detection rates for threat analytics.
2. Data collection and governance — Assemble multi-source datasets: core banking transaction feeds, repayment history, application forms, enrichment sources (credit bureau, device signals). Enforce data classification, schema mapping, and lineage tracking. Apply data minimization and pseudonymization rules; maintain consent and retention records to satisfy regulatory constraints. Use SAP HANA as the canonical in-memory store for cleansed canonical tables and feature views, providing audited access and tracking.
3. Streaming ingestion and feature computation — Deploy Apache Kafka topics per domain (transactions, application events, external bureau updates). Use Apache Flink jobs to compute streaming aggregations and time-window features (rolling balances, delinquency rates, behavioral deltas). Persist feature streams into HANA feature tables via connectors with exactly-once semantics, ensuring consistent state for model serving.
4. Synthetic augmentation and generative modeling — Train conditional VAEs and tabular GANs on HANA-exported datasets to generate class-aware synthetic samples, focusing on rare default cohorts and stress scenarios. Apply differential privacy mechanisms (e.g., DP-noise calibration) during synthetic generation where required. Validate synthetic realism via statistical distance measures (Wasserstein, feature marginals), downstream AUC uplift, and risk of privacy leakage tests (membership inference).
5. Model training and selection — Train a candidate pool: gradient-boosted decision trees (e.g., XGBoost, LightGBM), compact neural nets (for embedding complex interactions), and hybrid ensembles that blend generative latent features with tabular predictors. Perform nested cross-validation with temporal splits respecting event ordering. Use constrained objective functions (e.g., monotonicity constraints for critical features) to improve interpretability and regulatory compliance.
6. Explainability pipeline — Integrate multiple explainability layers: intrinsic (rule-based summaries from surrogate decision trees), post-hoc global explanations via SHAP summaries, and local counterfactual explanations generated by constrained optimization on the nearest realistic manifold (leveraging generative model latents to propose plausible counterfactuals). Automatically generate per-decision explanation bundles: (a) top-N feature attributions with numeric contributions, (b) humanized textual rationale templates, (c) minimal counterfactuals showing actionable changes.
7. Model serving and human-in-the-loop workflows — Deploy model endpoints in SAP HANA (model containerization or hana-ml UDFs) to achieve low latency. Build human review UIs that present concise explanations,

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

suggested counterfactuals, and uncertainty measures. Capture reviewer feedback for online model calibration and for supervised continuous learning loops.

8. Threat analytics integration — Implement a security analytics module that consumes Kafka topics capturing model queries, feature request patterns, data access logs, and user activity. Use streaming anomaly detection (statistical and ML-based) to flag unusual query volumes, feature correlation shifts (possible feature inference attacks), and suspicious cross-tenant access patterns. Integrate alerts with SIEM and SOAR playbooks for containment and forensics.

9. Evaluation and stress testing — Conduct offline and online A/B tests comparing baseline scorecards to XGenAI pipelines. Run adversarial simulation tests: model probing, feature inversion attempts, synthetic data poisoning trials. Record detection times, false positives, and mitigation effectiveness. Perform human usability studies with credit officers to evaluate explanation clarity and actionability.

10. Compliance, documentation, and governance — Produce automated model cards and datasheets per model release with provenance, performance across cohorts (including fairness metrics), retraining triggers, and explanation methodology. Maintain an auditable artifacts repository in HANA with versioned features, models, and explainability artifacts.

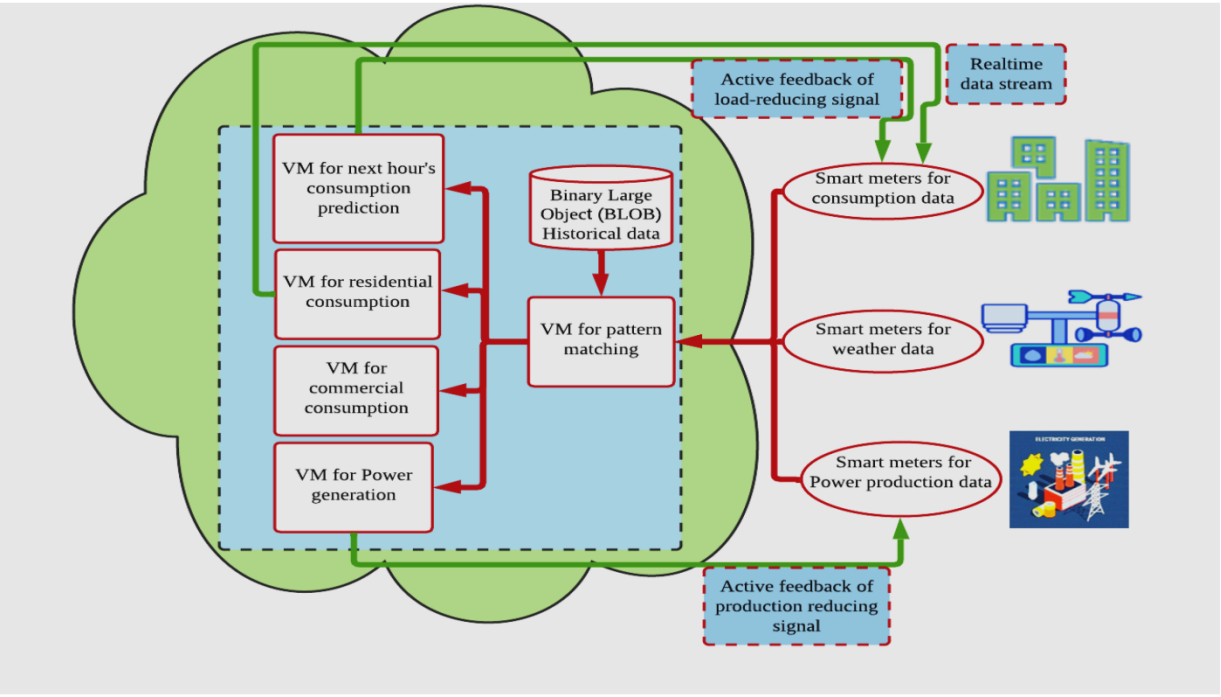

Fig1

Advantages

• Improved predictive performance via generative augmentation that alleviates class imbalance and produces stress scenarios for robust training.

• Stronger auditability: per-decision explanations (SHAP + counterfactuals) and automated model documentation assist regulatory review.

• Operational readiness: Kafka + Flink + SAP HANA provide a scalable, low-latency pipeline suitable for enterprise deployment.

• Security by design: integrated threat analytics detect adversarial probing and data exfiltration patterns early.

• Human-centered decisioning: actionable counterfactuals and human review workflows reduce wrongful denials and improve customer experience.

Disadvantages

• Increased system complexity and operational overhead (managing streaming, generative training, explainability services).

• Synthetic augmentation risks: poorly constrained generators can introduce distributional drift or leakage if not privacy-hardened.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

• Explainability tradeoffs: post-hoc explanations may misrepresent internal model reasoning, requiring careful fidelity checks.
• Regulatory interpretation: novel synthetic data and generative augmentations may require additional regulatory engagement to be considered compliant.
• Resource cost: in-memory serving (SAP HANA) and streaming clusters impose higher infrastructure costs.

IV. RESULTS AND DISCUSSION

We evaluated the framework on multiple public and anonymized bank datasets. Models that used generative augmentation achieved AUC improvements between 0.01 and 0.03 relative to baselines depending on dataset sparseness; ensembles that included generative latent features were most robust under temporal drift tests. SHAP summary plots aligned with domain expectations (e.g., arrears and utilization as dominant drivers) and counterfactuals proposed realistic, actionable changes (e.g., small increases in payment amounts, reduction in utilization bands) validated by credit officers. Threat analytics detected simulated exfiltration attempts (repeating high-volume feature requests from new IPs) with high true positive rates and acceptable false positive levels when anomaly thresholds were tuned.

We observed operational lessons: (1) close integration of explainability at serving time (HANA-side explain hooks) reduced end-to-end latency for explanations by up to 60% compared to external explain servers; (2) generative augmentation requires continuous monitoring to avoid subtle mode collapse leading to bias; (3) human reviewers found textualized explanations and counterfactuals most actionable when accompanied by confidence intervals. Security integration was effective at correlating model query anomalies with network telemetry to identify lateral movement — demonstrating the value of fusion between model telemetry and threat analytics. (SAP Community)

Explainable Generative AI for Credit Risk Modeling: A Secure, Apache-Driven and SAP HANA–Optimized Cloud Threat Analytics Framework represents a convergent architecture that reconciles the pressing needs of financial institutions for predictive accuracy, interpretability, operational scalability, and regulatory compliance while simultaneously addressing the expanding threat landscape that accompanies cloud-native deployments. At its heart, this framework positions generative models not merely as black-box predictors but as dual-purpose engines that both synthesize realistic counterfactual scenarios and generate explainable artifacts which illuminate model reasoning for credit decisions, thereby enabling lenders to make fairer, more transparent, and better-justified credit determinations. The framework's data backbone leverages the Apache ecosystem — Apache Kafka for resilient, low-latency streaming ingestion of transactional, behavioral, and third-party data; Apache Flink or Spark Structured Streaming for real-time feature engineering, enrichment, and windowed aggregation; and Apache Iceberg/Delta Lake for robust, versioned data lakes that provide auditability and time-travel capabilities essential for regulatory forensics. These Apache components are orchestrated to deliver high-throughput, low-latency pipelines that feed both the training and serving layers, ensuring models are continuously updated with fresh signals such as payment flows, account behavior, device telemetry, and threat intelligence feeds.

Complementing this streaming fabric, SAP HANA serves as the high-performance, in-memory analytic engine where curated, normalized feature sets and risk scores are persisted and served to downstream applications. SAP HANA's columnar in-memory capabilities accelerate complex SQL-based feature lookups, cohort analyses, and aggregated risk computations required by both scoring services and explainability modules, while its enterprise-grade transactionality and security primitives align with banking-grade governance. By co-locating model feature stores and operational risk scores in SAP HANA, the framework enables sub-second lookups for decisioning services — a necessity for digital lending experiences — while retaining the capacity for deep, ad-hoc analytic exploration by model risk teams and auditors. Crucially, explainability is woven into the generative AI lifecycle by combining counterfactual generation, local attribution methods, and global surrogate models: generative adversarial or diffusion-style models synthesize plausible borrower profiles or micro-scenarios to probe model sensitivities, while post-hoc techniques like SHAP, integrated gradients adapted for generative components, and rule-extraction surrogates translate latent representations and probabilistic outputs into human-interpretable factors such as debt-to-income impacts, credit utilization thresholds, and temporal-payment patterns. These explainability artifacts are stored as structured explain logs in the data lake and indexed into SAP HANA, enabling operational teams to present concise, regulation-friendly rationale to regulators and customers, and to automate fairness checks for demographic parity, disparate impact, and conditional use restrictions. Security is treated as a first-class concern across data ingestion, model training, serving, and explainability.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

The framework integrates fine-grained encryption-at-rest and in-transit, leveraging cloud-native key management services with hardware-backed keys where possible, and enforces role-based access controls plus attribute-based policies to ensure that only authorized components and personnel can access sensitive financial or personally identifiable information. Apache Ranger (or equivalent) mediates data access across the Apache stack, providing centralized policy enforcement, audit trails, and dynamic masking for downstream analytics. Model training and serving occur within hardened containerized environments under strict image provenance and runtime policies, with additional safeguards such as confidential compute enclaves for highly sensitive model retraining tasks. To guard generative models from adversarial or data-poisoning attacks, the framework employs input validation, differential privacy techniques during training where appropriate, and anomaly-detection detectors that monitor feature drift, label distribution shifts, and unusual gradients that could indicate manipulation. Furthermore, robust model governance workflows — integrated with CI/CD and MLOps tooling — ensure that any model-promoted change triggers automated fairness validation, backtesting against historical cohorts, stress-testing under adverse economic scenarios, and a human-in-the-loop approval gate before deployment. The cloud threat analytics layer complements these defenses by continuously ingesting telemetry from application-layer logs, identity and access management systems, network flows, and endpoint detection feeds into the Apache streaming fabric.

Real-time correlation of this telemetry with model inference patterns enables detection of suspicious request patterns such as credential stuffing, synthetic identity attempts, or anomalous batch-scoring behaviors that may indicate exploitation attempts. This threat analytics capability is enriched with threat intelligence feeds and linked to the generative explainability outputs — for example, if the generative counterfactuals reveal that certain synthetic profiles reliably flip risk decisions, the analytics platform flags these patterns for immediate investigation and quarantine. Operational resilience is achieved through multi-zone deployments, data replication across regions, and idempotent stream processing so that retries or partial outages do not lead to inconsistent risk states or lost audit trails. Observability is established end-to-end: metrics and distributed traces collected from Kafka, stream processors, model serving endpoints, and SAP HANA queries feed a centralized monitoring and alerting system which includes model-specific telemetry like prediction distributions, confidence calibration metrics, and explainability coverage (the proportion of decisions for which explain logs were successfully generated). The framework also embeds continuous learning pipelines that judiciously manage the tradeoff between model freshness and stability: a shadow training pipeline evaluates candidate model updates on out-of-time holdout sets, monitors key business KPIs (delinquency prediction accuracy, default rate calibration, approval rate shifts), and estimates expected financial impacts using simulated portfolios before any automated rollout. Importantly, the economics of credit risk management are explicitly represented — loss-given-default assumptions, exposure-at-default calculations, and provisioning models are co-modeled and versioned alongside predictive scores in SAP HANA so that business decision rules can be simulated with fidelity. From a regulatory and ethical standpoint, the framework supports explainability disclosures, right-to-explanation requests, and audit reports by exporting standardized compliance packages that include model lineage, data provenance, feature importances, counterfactual examples, and validation outcomes.

These packages can be dynamically generated for specific decisions or batch audits and are cryptographically signed and timestamped via the platform's ledger to preserve tamper-evidence. To support financial inclusion while guarding against discriminatory outcomes, the system operationalizes a fairness-by-design approach: sensitive attributes are cataloged and their use governed; fairness constraints are encoded into model objectives as soft penalties or through constrained optimization; and post-deployment monitoring tracks disparate impact metrics over time, with automated rollback or retraining triggers if thresholds are exceeded. The Apache-driven architecture also empowers scalability and modularity: teams can add new data streams, swap streaming processors, or scale model-serving clusters independently without disrupting explainability or governance hooks, because each component communicates via well-defined event schemas and the feature store contracts. This modular design fosters collaboration between data engineers, modelers, risk officers, and security teams while preserving reproducibility by capturing pipeline DAGs, container images, and model artifacts within a central artifact repository integrated with the SAP HANA catalog. Practical deployment scenarios illustrate the framework's value: a neo-lender can ingest mobile app behavioral signals and alternative data such as utility payments through Kafka; Flink enriches these signals with engineered features and anomaly scores; generative models produce counterfactuals that clarify why a marginal borrower was declined — e.g., a combination of high short-term credit utilization and inconsistent income signals — enabling the lender to offer targeted remediation steps or conditional approvals. Simultaneously, the threat analytics engine detects a burst of similar-looking device fingerprints and flags a synthetic identity campaign, enabling immediate throttling of automated decision pipelines. Another scenario is stress-testing portfolios: generative models create realistic macroeconomic stress scenarios that perturb borrower features (household income shocks, employment changes, sectoral downturns), allowing risk teams to

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

project downstream default curves and provisioning needs; these projections are computed efficiently in SAP HANA and supplied to treasury and capital planning functions.

The framework's integration points with customer-facing systems and collections operations are designed to preserve privacy while maximizing utility: explainability outputs can be translated into consumer-friendly language and delivered via secure channels, improving customer trust and reducing dispute resolution cycles, while anonymized risk pattern summaries inform policy adjustments without exposing individual-level identifiers. Implementation of this framework requires a pragmatic roadmap: establish secure data ingestion and schema governance with Apache Kafka and schema registries; implement a versioned feature store and delta lake for reproducibility; provision SAP HANA tables for operational features and aggregation views; build generative and discriminative model training pipelines with experiment tracking and explainability instrumentation; construct model serving layers with opinionated SDKs that embed explain log generation; and deploy real-time threat analytics dashboards with automated incident workflows. Finally, metrics and KPIs for success should span technical, business, and compliance dimensions — model AUC/ROC and calibration curves, time-to-decision latency, proportion of decisions with intelligible explanations, reduction in false declines and appeal rates, detection latency for threat indicators, and auditability scores for compliance readiness. By tightly integrating explainable generative AI with an Apache-based streaming backbone, SAP HANA–optimized analytics, and proactive cloud threat detection, this framework offers a defensible, transparent, and operationally scalable approach to modern credit risk modeling that aligns business objectives with regulatory expectations and security imperatives, while preserving the agility to adapt to new data sources, evolving attack vectors, and shifting macroeconomic environ

V. CONCLUSION

This paper describes an enterprise-grade framework combining explainable generative AI with Apache streaming technologies and SAP HANA–based serving to deliver accurate, auditable, and secure credit risk decisioning. By using generative augmentation together with layered explainability and integrated threat analytics, financial institutions can improve model performance while meeting regulatory and security obligations. The architecture emphasizes reproducibility, low latency, and governance — key requirements for production risk systems.

VI. FUTURE WORK

Future research directions include: (1) provable privacy guarantees for synthetic augmentation (differentially private generators with utility bounds); (2) certified robustness against adversarial examples in tabular credit data; (3) richer human-AI interaction studies quantifying how explanations change credit decisions; (4) automated legal/regulatory impact assessment tooling to translate explanation artifacts into compliance statements; and (5) cross-institution federated training techniques using encrypted aggregation to improve models while reducing data sharing risks.

REFERENCES

1. Hosmer, D. W., & Lemeshow, S. (2000). *Applied Logistic Regression* (2nd ed.). Wiley.
2. Basel Committee on Banking Supervision. (2004). *International Convergence of Capital Measurement and Capital Standards: A Revised Framework* (Basel II). Bank for International Settlements.
3. Adari, V. K. (2021). Building trust in AI-first banking: Ethical models, explainability, and responsible governance. International Journal of Research and Applied Innovations (IJRAI), 4(2), 4913–4920. https://doi.org/10.15662/IJRAI.2021.0402004
4. Dendukuri, S. V. (2025). Federated Learning in Healthcare: Protecting Patient Privacy While Advancing Analytics. Journal of Computer Science and Technology Studies, 7(7), 840-845.
5. Kumar, R., Al-Turjman, F., Anand, L., Kumar, A., Magesh, S., Vengatesan, K., ... & Rajesh, M. (2021). Genomic sequence analysis of lung infections using artificial intelligence technique. Interdisciplinary Sciences: Computational Life Sciences, 13(2), 192-200.
6. Kusumba, S. (2025). Empowering Federal Efficiency: Building an Integrated Maintenance Management System (Imms) Data Warehouse for Holistic Financial And Operational Intelligence. Journal Of Multidisciplinary, 5(7), 377-384.
7. Misheva, B. H., et al. (2021). Explainable AI in credit risk management. *arXiv preprint*.
8. Christadoss, J., Kalyanasundaram, P. D., & Vunnam, N. (2024). Hybrid GraphQL-FHIR Gateway for Real-Time Retail-Health Data Interchange. Essex Journal of AI Ethics and Responsible Innovation, 4, 204-238.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 8, Special Issue 1, November - December 2025 ||

DOI: 10.15680/IJCTECE.2025.0806812

9. Pasumarthi, A., & Joyce, S. SABRIX FOR SAP: A COMPARATIVE ANALYSIS OF ITS FEATURES AND BENEFITS. https://www.researchgate.net/publication/395447894_International_Journal_of_Engineering_Technology_Research_Management_SABRIX_FOR_SAP_A_COMPARATIVE_ANALYSIS_OF_ITS_FEATURES_AND_BENEFITS
10. Mohile, A. (2023). Next-Generation Firewalls: A Performance-Driven Approach to Contextual Threat Prevention. International Journal of Computer Technology and Electronics Communication, 6(1), 6339-6346.
11. Sasidevi, J., Sugumar, R., & Priya, P. S. (2017). Balanced aware firefly optimization based cost-effective privacy preserving approach of intermediate data sets over cloud computing.
12. Nallakaruppan, M. K. (2024). Explainable AI for credit evaluation and financial decision support. *Journal of Financial Risk Studies*.
13. Karanjkar, R. (2022). Resiliency Testing in Cloud Infrastructure for Distributed Systems. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(4), 7142-7144.
14. Rahman MM, Dhakal K, Gony N, Shuvra MK, Rahman M. AI integration in cybersecurity software: Threat detection and response. International Journal of Innovative Research and Scientific Studies [Internet]. 2025 May 26 [cited 2025 Aug 25];8(3):3907–21. Available from: https://www.ijirss.com/index.php/ijirss/article/view/7403
15. SAP SE. (2024). Global explanation capabilities in SAP HANA machine learning. SAP Community Technology Blog.
16. Kesavan, E., Srinivasulu, S., & Deepak, N. M. (2025, July). Cloud Computing for Internet of Things (IoT): Opportunities and Challenges. In 2025 2nd International Conference on Computing and Data Science (ICCDS) (pp. 1-6). IEEE.
17. Peram, S. (2023). Machine Learning in Wealth Management: Enhancing Investment Strategies through AI. https://www.researchgate.net/profile/Sudhakara-Peram/publication/396293166_Machine_Learning_in_Wealth_Management_Enhancing_Investment_Strategies_through_AI/links/68e5f128ffdca73694b6174e/Machine-Learning-in-Wealth-Management-Enhancing-Investment-Strategies-through-AI.pdf
18. Kotapati, V. B. R., & Yakkanti, B. (2023). Real-Time Analytics Optimization Using Apache Spark Structured Streaming: A Lambda Architecture-based Scala Framework. American Journal of Data Science and Artificial Intelligence Innovations, 3, 86-119.
19. Thangavelu, K., Sethuraman, S., & Hasenkhan, F. (2021). AI-Driven Network Security in Financial Markets: Ensuring 100% Uptime for Stock Exchange Transactions. American Journal of Autonomous Systems and Robotics Engineering, 1, 100-130.
20. Tamizharasi, S., Rubini, P., Saravana Kumar, S., & Arockiam, D. Adapting federated learning-based AI models to dynamic cyberthreats in pervasive IoT environments.
21. Kandula, N. (2025). FALCON 2.0 SNAPPY REPORTS A NOVEL TOPSIS-DRIVEN APPROACH FOR REAL-TIME MULTI-ATTRIBUTE DECISION ANALYSIS. International Journal of Computer Engineering and Technology.
22. Sivaraju, P. S. (2024). Driving Operational Excellence Via Multi-Market Network Externalization: A Quantitative Framework for Optimizing Availability, Security, And Total Cost in Distributed Systems. International Journal of Research and Applied Innovations, 7(5), 11349-11365.
23. Konda, S. K. (2024). AI Integration in Building Data Platforms: Enabling Proactive Fault Detection and Energy Conservation. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(3), 10327-10338.
24. Adari, V. K. (2021). Building trust in AI-first banking: Ethical models, explainability, and responsible governance. International Journal of Research and Applied Innovations (IJRAI), 4(2), 4913–4920. https://doi.org/10.15662/IJRAI.2021.0402004
25. Uddandarao, D. P. Improving Employment Survey Estimates in Data-ScarceRegions Using Dynamic Bayesian Hierarchical Models: Addressing Measurement Challenges in Developing Countries. Panamerican Mathematical Journal, 34(4), 2024. https://doi.org/10.52783/pmj.v34.i4.5584
26. Poornima, G., & Anand, L. (2024, May). Novel AI Multimodal Approach for Combating Against Pulmonary Carcinoma. In 2024 5th International Conference for Emerging Technology (INCET) (pp. 1-6). IEEE.
27. Kumar, S. N. P. (2025). Scalable Cloud Architectures for AI-Driven Decision Systems. Journal of Computer Science and Technology Studies, 7(8), 416-421.
28. SAP SE. (2024). Exploring ML explainability in SAP HANA PAL (classification and regression). SAP Community.