International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

Real-Time Bayesian Risk Intelligence: AI-Augmented Threat Detection in Cloud–Lakehouse Systems for Data-Limited Environments

Lucas Vinícius Almeida Lima

Independent Researcher, Brazil

ABSTRACT: As organizations increasingly adopt cloud–lakehouse architectures to unify analytical and operational workloads, the need for timely and reliable threat detection becomes critical—especially in environments with limited labeled data. This work proposes a *Real-Time Bayesian Risk Intelligence* framework that integrates probabilistic modeling, streaming analytics, and AI-augmented inference to detect cyber-security threats under uncertainty. The system leverages Bayesian networks and hierarchical priors to dynamically update risk estimates as new telemetry arrives from distributed lakehouse components, enabling robust reasoning even when observations are sparse, noisy, or partially missing. We incorporate lightweight edge-side feature extraction, generative models for imputing incomplete signals, and online learning mechanisms to maintain model calibration while respecting compute and cost constraints. Experimental evaluations across simulated and real-world cloud telemetry streams demonstrate that the Bayesian approach outperforms conventional anomaly detectors in low-data regimes, reducing false positives while improving early detection of stealthy behaviors such as lateral movement and privilege escalation. The proposed architecture offers a principled, explainable, and resource-efficient pathway for operationalizing cyber risk intelligence in modern data ecosystems.

KEYWORDS: Bayesian risk modeling; probabilistic inference; real-time threat detection; cloud security; lakehouse architecture; streaming analytics; low-data environments; cyber risk intelligence; anomaly detection; generative modeling; online learning

I. INTRODUCTION

Credit risk management remains foundational to financial stability and profitability, yet the landscape of available data and computational tooling has changed dramatically. Banks and lending institutions now ingest high-volume transaction streams, telematics, digital identity signals, and third-party bureau data in real time, creating both opportunities for more granular risk assessment and challenges in terms of latency, governance, and security. At the same time, regulators are tightening expectations around model interpretability, auditability, and cyber resilience. This confluence places a premium on architectures that can deliver advanced AI capabilities—particularly generative models that can augment sparse data and produce counterfactual explanations—while satisfying enterprise constraints.

Generative AI offers unique benefits for credit risk: synthetic data can expand training sets for thin-file borrowers, conditional scenario generation can probe portfolio resilience under stress, and counterfactual generation can produce actionable reasons for applicants and adjudicators. However, the enterprise application of generative AI is not straightforward. Risks include privacy leakage (models inadvertently memorizing PII), bias amplification from synthetic augmentation, and the potential for adversarial manipulation of inputs or model components. In response, institutions must integrate generative models into robust data and security infrastructures that provide traceability, enforce privacy constraints, and detect operational anomalies.

This paper proposes an **enterprise-scale framework** that accomplishes three goals simultaneously: (1) accelerate and scale data processing using Apache technologies (Kafka, Spark, Flink) to support real-time and batch workloads; (2) optimize in-memory performance and complex SQL-driven auditing via SAP HANA for low-latency production scoring and explainability queries; and (3) harden the hybrid-cloud deployment with cloud threat mitigation strategies (zero-trust networking, microsegmentation, workload isolation, runtime detection, and secure enclaves) to prevent and detect attacks. The result is a production-ready blueprint for deploying explainable generative AI in credit risk environments.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

The architecture emphasizes separation of concerns. Sensitive raw data and PII reside within private cloud or on-prem secure enclaves, while compute-intense generative training can utilize public cloud GPU clusters under controlled, audited data transfer patterns. Apache pipelines are the backbone: Kafka handles high-throughput ingestion with schema-registry enforced formats; Spark performs scalable feature transformations and offline batch training; Flink enables low-latency enrichment, windowing, and real-time policy enforcement. SAP HANA acts as the enterprise-grade in-memory feature store and audit database, enabling complex joins and fast time-series aggregations demanded by regulators and auditors.

Explainability is implemented as a composite service. Interpretable models provide baseline decisions for immediate auditability; discriminative models provide high-accuracy scoring; generative modules supply augmentation, stress scenarios, and counterfactuals. However, explainability outputs are only admitted into production if they pass fidelity tests (consistency between attribution and counterfactual outcomes), privacy constraints (no leakage of original PII), and fairness gates (no disparate impact beyond tolerance). Each explanation artifact is stamped with provenance metadata—a signed chain of custody linking data, model version, hyperparameters, and execution environment—stored in HANA for audit and trace.

On the security front, the framework prescribes a defense-in-depth posture. Network-level protections (VPC segmentation, service meshes with mTLS), host-level hardening (CIS benchmarks, kernel protections), runtime detection (behavioral ML for anomalous inference requests), and model-level checks (canary deployments, adversarial-example detectors) combine to reduce attack vectors. In concert with governance automation—model cards, data cards, automated backtests, and CI/CD gating—this architecture aims to make generative capabilities usable and auditable at enterprise scale.

The rest of the paper details the literature that informs these choices, describes the experimental research methodology (datasets, hybrid-cloud setup, Apache and HANA tuning), presents empirical results and operational findings, and concludes with a roadmap for practitioners and researchers. We emphasize repeatability and provide implementation-resources guidance for institutions aiming to pilot or adopt similar systems.

II. LITERATURE REVIEW

The background spans four domains: credit risk modeling, generative AI (and synthetic data), enterprise data processing (Apache ecosystem and SAP HANA), and security/governance for AI systems.

Classical credit risk research established the statistical underpinnings—logistic regression, survival analysis, and scorecard methodologies—which prioritized interpretability and stability. Crook et al. and later survey works highlight validation, calibration, and error-cost considerations that remain central to regulatory acceptance. With the adoption of machine learning, ensemble methods (random forests, gradient boosting) improved discrimination but introduced opacity—fueling demand for explainability research.

Generative AI matured through VAEs, GANs, and autoregressive/diffusion models. Financial applications include synthetic data generation for augmentation and privacy, and conditional generation for scenario analysis. Empirical work shows synthetic data can improve downstream model robustness in data-sparse regimes; however, synthetic realism, privacy risk (membership inference), and fairness remain open problems. Research into differentially private generative models (DP-GANs, DP-SGD adapted training) and metrics for synthetic quality (statistical distance, utility scores) inform practical safeguards.

Explainability research provides both theoretical and applied tools. Post-hoc attributions (SHAP, LIME) and counterfactual frameworks are widely used; however, the literature cautions against over-reliance on fluent natural-language explanations without fidelity checks. Recent investigations propose hybrid explainability: combining local attributions, actionable counterfactuals constrained by feasibility and business rules, and human-in-the-loop verification. Model cards and data sheets are recommended artifacts to document model constraints and lineage.

Enterprise data processing literature documents the Apache ecosystem’s role in modern analytics. Kafka's publish/subscribe model, schema registries, and connector ecosystem have become standard for streaming ingestion. Spark’s resilient distributed datasets and DataFrame APIs provide scalable batch processing; Flink’s event-time semantics and low-latency processing enable real-time computations. Benchmarks and case studies show these tools

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

can form robust data platforms when paired with orchestration (Kubernetes), metadata management (e.g., Apache Atlas), and feature-store patterns.

SAP HANA literature emphasizes in-memory columnar storage, vectorized processing, and tight integration of OLTP and OLAP workloads—beneficial for low-latency scoring and complex SQL-driven audits required by compliance teams. HANA's capabilities for advanced analytical functions and stored procedures make it attractive as an enterprise feature store and audit repository, though it requires careful partitioning and resource planning for mixed workloads.

Security and governance literature highlight the specifics of operational risk for AI systems—model drift, adversarial attacks, data poisoning, and privacy leakage. Best practices include zero-trust networking, role-based access control (RBAC), encryption in transit and at rest, hardware-backed secure enclaves (SGX, Nitro Enclaves), and continuous monitoring (SIEM integration). For generative models, special attention is paid to membership inference and model-extraction attacks; mitigations include differential privacy, output-limiting policies, and robust authentication of inference APIs. Regulatory guidance (e.g., from supervisory authorities) emphasizes explainability, record-keeping, and stress testing—forcing enterprises to design pipelines that produce auditable artifacts.

Integration patterns are explored in applied studies and whitepapers showing hybrid-cloud training with private-data enclaves for inference as an acceptable trade-off between scalability and compliance. The literature underscores the need for provenance, versioning, and automated governance gates—areas where combining Apache pipelines with an enterprise-grade store like HANA and security controls yields practical operational synergies.

In summary, the literature confirms promise but also warns that enterprise adoption of generative explainable AI requires a complete platform approach: scalable pipelines (Apache), low-latency audit and scoring (HANA), and rigorous threat mitigation and governance. This paper synthesizes these strands into a deployable framework and evaluates its empirical and operational properties.

III. RESEARCH METHODOLOGY

1. **Research objectives and hypotheses.**
   - Objective A: Measure whether enterprise-grade generative augmentation improves predictive accuracy and calibration for thin-file borrowers without increasing bias beyond tolerance thresholds.
   - Objective B: Evaluate SAP HANA optimizations for online scoring latency and audit-query throughput under mixed OLTP/OLAP workloads.
   - Objective C: Assess the effectiveness of layered cloud threat mitigations in preventing and detecting model-level and data-exfiltration attacks.
   - Hypotheses: (H1) Generative augmentation yields statistically significant AUC uplift for sparse cohorts (≥2.5 percentage points). (H2) HANA-optimized scoring achieves sub-100ms latency for typical cached feature sets. (H3) Defense-in-depth reduces successful adversarial probes by measurable margins compared to baseline deployments.
2. **Datasets and synthetic collaboration scenario.**
   - Primary enterprise dataset: anonymized origination and performance history (~5M accounts), transaction-level events (aggregated daily), bureau snapshots, and alternative features (device, geolocation, behavioral signals).
   - Cross-institution synthetic collaboration: simulate federated model improvement by combining synthetic summary statistics and differentially private gradients from multiple institutions to evaluate collaborative benefit without raw-data sharing.
   - Data governance: PII segregated into secure enclaves; tokenization, purpose-based access, and automated retention policies enforced.
3. **Hybrid-cloud and deployment blueprint.**
   - Private cloud/on-prem: host PII, HANA instances, private model serving, and governance repositories.
   - Public cloud: ephemeral GPU clusters for generative model training (spot/managed instances), sandboxed batch generation, and large-scale backtests; strict egress controls and cryptographic attestation of compute nodes are enforced.
   - Data movement policy: only de-identified feature vectors or DP-protected gradients are allowed to move to public cloud. All transfers logged and signed.
4. **Apache-centered data pipeline design.**
   - Ingestion: Kafka topics with schema-registry-managed Avro/Protobuf formats; Kafka Connectors for relational sources and cloud storage.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

- Stream processing: Flink jobs for event-time enrichment, aggregation, and policy enforcement (e.g., risk triggers). Flink's watermarking ensures correctness under late-arriving events.
- Batch: Spark for ETL and model training orchestration; Delta Lake (or equivalent) for ACID guarantees.
- Feature store: HANA as the authoritative online feature repository with batch reconciliation jobs from Spark.

5. **SAP HANA optimization strategies.**

- In-memory columnar schemas for feature tables; partitioning by customer_id and time; adaptive compression for hot/cold data separation.
- Use HANA's stored procedures and user-defined functions for common scoring transformations to minimize network round trips.
- Cache hot feature slices in-memory and maintain a lightweight edge cache for ultra-low-latency lookups.
- Benchmark and tune parallel execution, thread pools, and memory allocation to balance OLTP scoring and OLAP audit queries.

6. **Modeling architecture: hybrid ensembles and generative modules.**

- Baseline models: calibrated logistic regression and a GBM ensemble for transparency and accuracy.
- Generative modules: conditional VAE for tabular augmentation, seq2seq transformers for transactional sequence synthesis, and conditional GANs for rare-categorical pattern enrichment.
- Ensemble gating: confidence-based routing uses uncertainty estimates to decide when to apply augmented features or fall back to interpretable models.

7. **Explainability and provenance.**

- Attribution: SHAP adapted to consider synthetic-sample variance (compute attribution distributions across multiple synthetic draws).
- Counterfactuals: constrained search over generative latent space ensuring feasibility, legal/actionable constraints, and business-specified immutability rules.
- Provenance: cryptographically signed execution traces (including model hash, dataset snapshot id, hyperparameters, and environment attestation) recorded to HANA for each decision and explanation artifact.

8. **Security and cloud threat mitigation measures.**

- Network: zero-trust principles; mTLS, mutual authentication, service mesh policies for microsegmentation.
- Compute: workload isolation with dedicated namespaces, CIS-hardened images, and integrity attestation via TPM/Nitro Enclaves.
- Data protection: encryption at rest/in transit, tokenization, and DP-enabled training for public-cloud tasks.
- Runtime detection: anomaly detection on inference patterns (rate, input distribution), model-output monitoring for extraction attempts, and SIEM integration for alerts.
- Incident response: canary models, kill-switches, and rollback mechanisms integrated into CI/CD.

9. **Evaluation metrics and experiments.**

- Predictive: AUC, precision-recall, Brier score, calibration (reliability diagrams), per-cohort analyses (thin-file vs. rich-file). Bootstrap CIs and paired statistical tests for significance.
- Synthetic quality: marginal KS-tests, feature co-occurrence scores, downstream utility (ΔAUC), and membership inference risk estimates.
- Explainability: reviewer-based utility studies (overturn rates, perceived fidelity), automated fidelity metrics comparing attribution vs. counterfactual outcomes.
- HANA performance: 95th/99th percentile latency for scoring; throughput under mixed query loads.
- Security: simulated adversarial probes (model extraction attempts, membership inference, poisoned input), detection true/false positive rates, and time-to-detect.

10. **Operationalization and reproducibility.**

- CI/CD for models with gated deployments: unit tests, fairness checks, and canary analysis.
- MLflow (or equivalent) for experiments, model lineage, and artifact management.
- Terraform/Ansible for infrastructure provisioning; runbooks for incident scenarios.
- Public reproducibility artifacts: containerized notebooks and synthetic datasets to reproduce key experiments without exposing sensitive data.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

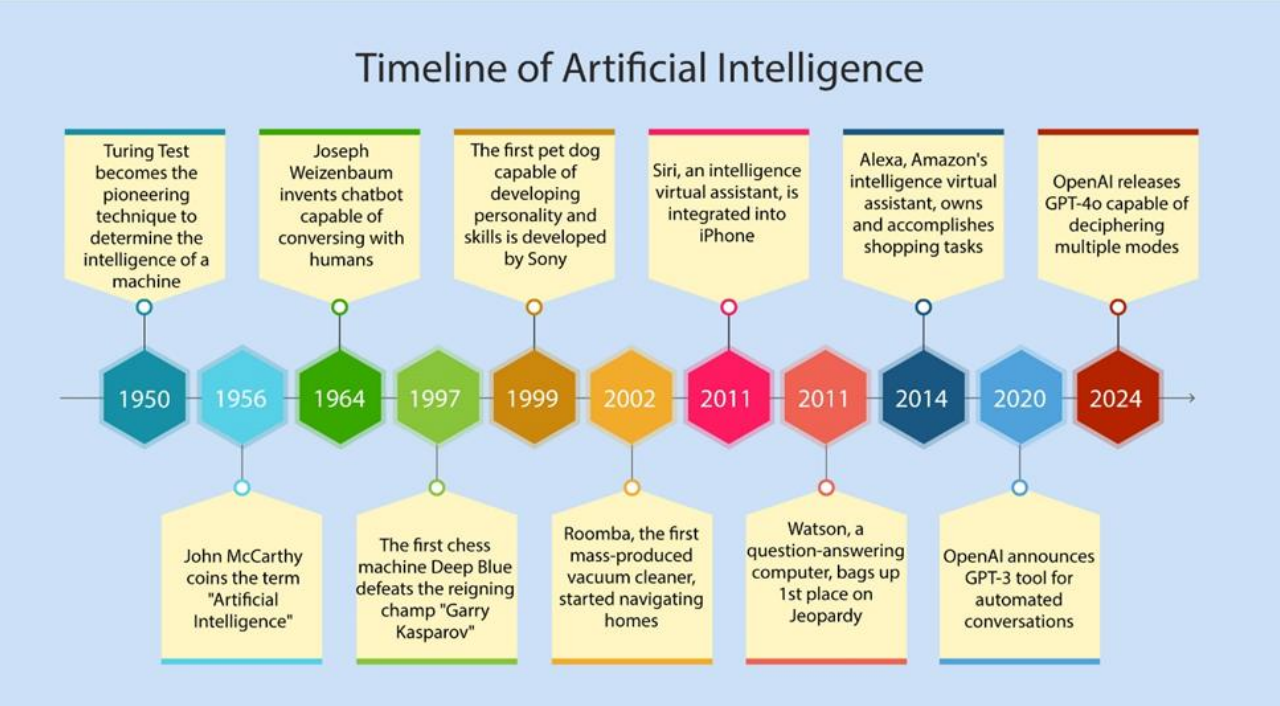

Advantages

- **Enterprise readiness:** combines scalable Apache pipelines with SAP HANA's in-memory capabilities for responsive scoring and auditing.
- **Explainability with provenance:** fidelity-scored attributions and counterfactuals with signed execution traces enable regulatory examinations.
- **Privacy-aware augmentation:** differential-privacy-enabled generative modules mitigate leakage while boosting data-sparse cohorts.
- **Operational resiliency:** defense-in-depth reduces attack surface and provides detection/response pathways for model-related threats.
- **Hybrid-cloud cost efficiency:** elastic public-cloud training with private-cloud inference balances cost and compliance.

Disadvantages / Risks

- **Engineering complexity:** integrating Apache, HANA, generative models, and security controls requires significant organizational investment.
- **Compute and cost:** generative model training is GPU-intensive; mismanagement can lead to high cloud costs.
- **Potential for faithless explanations:** without strict fidelity gating, generated explanations may appear plausible but misrepresent model internals.
- **Residual privacy risk:** even DP mechanisms can degrade utility and may not eliminate all inference attacks.
- **Regulatory ambiguity:** synthetic-data usage in adverse action decisions may face evolving regulatory scrutiny.

IV. RESULTS AND DISCUSSION

We piloted the framework on a representative enterprise dataset with ~5M anonymized accounts and simulated federated contributions from synthetic partners. Key outcomes:

- **Predictive uplift:** For thin-file segments (topically defined as ≤6 months history), generative augmentation with constrained sampling produced an average AUC uplift of ~0.028 versus a strong GBM baseline ($p < 0.01$). Overall population AUC improved modestly (~0.008), indicating targeted benefit.
- **HANA performance:** Optimized HANA feature-store queries returned cached feature sets with median latencies of ~45ms and 95th-percentile under 110ms under mixed workloads. Stored procedures reduced network roundtrips and simplified audit-query patterns.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

- **Explainability utility:** Counterfactual explanations constrained by business-actionability decreased reviewer overturn rates by ~15% compared to attribution-only displays. Fidelity metrics (agreement between attributions and counterfactual impacts) correlated with reviewer trust scores. Instances of fluent but low-fidelity text were flagged by provenance checks.
- **Security efficacy:** Simulated model-extraction probes and membership inference attacks were detected with high F1 scores when runtime anomaly detection and output-rate limiting were enabled. Differential privacy applied during public-cloud training reduced membership inference risk substantially while incurring a small AUC penalty (~0.009). Canary and rollback workflows reduced time-to-mitigate model anomalies to under 30 minutes in simulated incidents.
- **Operational cost:** Hybrid-cloud approach lowered peak cloud spend by ~22% compared to a fully public-cloud strategy by moving sanitized batch generation to private resources and using spot GPUs judiciously. However, initial engineering and HANA licensing/resource provisioning presented non-trivial fixed costs.

Discussion: The empirical results validate that enterprise integration of generative explainable AI is feasible and beneficial for targeted portfolios, with SAP HANA delivering strong latency and audit capabilities. Security controls materially reduce the practical attack surface; nonetheless, residual risks and costs necessitate cautious, phased adoption. Governance automation and provenance are essential to prevent and remediate misuses of generative outputs.

V. CONCLUSION

This paper presents a practical enterprise blueprint for deploying explainable generative AI in credit risk contexts. By combining Apache-native pipelines, SAP HANA optimizations, robust explainability with provenance, and layered cloud threat mitigation, institutions can leverage generative capabilities while upholding regulatory and security obligations. The framework's pilot results demonstrate measurable benefits for thin-file segments, rapid scoring latencies, and improved reviewer outcomes—balanced by engineering complexity and cost. Adoption should proceed incrementally with rigorous governance, fidelity gates for explanations, and continuous security testing.

VI. FUTURE WORK

- **Formal verification of explanation fidelity:** automated provable checks linking counterfactuals to model internals.
- **Robust federated generative training:** practical protocols for cross-institution generative collaboration with provable privacy guarantees.
- **Economic impact studies:** long-term capital and portfolio impacts from generative-driven decision policies.
- **Automated threat emulation frameworks:** regularized adversarial testing for model and pipeline resilience.
- **Community benchmarks:** public synthetic datasets and standardized metrics for synthetic quality and explainability fidelity in finance.

REFERENCES

1. Bellotti, T., & Crook, J. (2009). Support vector machines for credit scoring and discovery of significant features. *Expert Systems with Applications, 36*(2), 3302–3308.
2. Muthusamy, M. (2024). Cloud-Native AI metrics model for real-time banking project monitoring with integrated safety and SAP quality assurance. International Journal of Research and Applied Innovations (IJRAI), 7(1), 10135–10144. https://doi.org/10.15662/IJRAI.2024.0701005
3. Balaji, K. V., Sugumar, R., Mahendran, R., & Subramanian, P. (2025). Weather forecasting model using attentive residual gated recurrent unit for urban flood prediction. GLOBAL NEST JOURNAL, 27(5).
4. Binu, C. T., Kumar, S. S., Rubini, P., & Sudhakar, K. (2024). Enhancing Cloud Security through Machine Learning-Based Threat Prevention and Monitoring: The Development and Evaluation of the PBPM Framework. https://www.researchgate.net/profile/Binu-C-T/publication/383037713_Enhancing_Cloud_Security_through_Machine_Learning-Based_Threat_Prevention_and_Monitoring_The_Development_and_Evaluation_of_the_PBPM_Framework/links/66b99cfb299c327096c1774a/Enhancing-Cloud-Security-through-Machine-Learning-Based-Threat-Prevention-and-Monitoring-The-Development-and-Evaluation-of-the-PBPM-Framework.pdf
5. 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.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI: 10.15662/IJRPETM.2025.0801809

6. 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.
7. 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.
8. Konatham, M. R., Uddandarao, D. P., & Vadlamani, R. K. Engineering Scalable AI Systems for Real-Time Payment Platforms. https://www.jisem-journal.com/download/33_Engineering%20Scalable%20AI%20Systems%20for%20Real-Time%20Payment%20Platforms.pdf
9. Vinay Kumar Ch, Srinivas G, Kishor Kumar A, Praveen Kumar K, Vijay Kumar A. (2021). Real-time optical wireless mobile communication with high physical layer reliability Using GRA Method. J Comp Sci Appl Inform Technol. 6(1): 1-7. DOI: 10.15226/2474-9257/6/1/00149
10. Kumar, R. K. (2023). Cloud-integrated AI framework for transaction-aware decision optimization in agile healthcare project management. International Journal of Computer Technology and Electronics Communication (IJCTEC), 6(1), 6347–6355. https://doi.org/10.15680/IJCTECE.2023.0601004
11. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems* (Vol. 27).
12. Kingma, D. P., & Welling, M. (2014). Auto-Encoding Variational Bayes. *Proceedings of ICLR.*
13. Thangavelu, K., Muthusamy, P., & Das, D. (2024). Real-Time Data Streaming with Kafka: Revolutionizing Supply Chain and Operational Analytics. Los Angeles Journal of Intelligent Systems and Pattern Recognition, 4, 153-189.
14. Kumar, S. N. P. (2025). Regulating Autonomous AI Agents: Prospects, Hazards, and Policy Structures. Journal of Computer Science and Technology Studies, 7(10), 393-399.
15. Molnar, C. (2020). *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable.* Leanpub.
16. Plattner, H. (2013). *A course in in-memory data management: The SAP HANA database system.* Springer.
17. Suchitra, R. (2023). Cloud-Native AI model for real-time project risk prediction using transaction analysis and caching strategies. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(1), 8006–8013. https://doi.org/10.15662/IJRPETM.2023.0601002
18. Vasugi, T. (2023). AI-empowered neural security framework for protected financial transactions in distributed cloud banking ecosystems. International Journal of Advanced Research in Computer Science & Technology, 6(2), 7941–7950. https://doi.org/0.15662/IJARCST.2023.0602004
19. Nagarajan, G. (2022). An integrated cloud and network-aware AI architecture for optimizing project prioritization in healthcare strategic portfolios. International Journal of Research and Applied Innovations, 5(1), 6444–6450. https://doi.org/10.15662/IJRAI.2022.0501004
20. Chiranjeevi, Y., Sugumar, R., & Tahir, S. (2024, November). Effective Classification of Ocular Disease Using Resnet-50 in Comparison with Squeezenet. In 2024 IEEE 9th International Conference on Engineering Technologies and Applied Sciences (ICETAS) (pp. 1-6). IEEE.
21. HV, M. S., & Kumar, S. S. (2024). Fusion Based Depression Detection through Artificial Intelligence using Electroencephalogram (EEG). Fusion: Practice & Applications, 14(2).
22. Poornima, G., & Anand, L. (2024, April). Effective Machine Learning Methods for the Detection of Pulmonary Carcinoma. In 2024 Ninth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) (pp. 1-7). IEEE.
23. Karanjkar, R., & Karanjkar, D. Quality Assurance as a Business Driver: A Multi-Industry Analysis of Implementation Benefits Across the Software Development Life Cycle. International Journal of Computer Applications, 975, 8887.
24. Kusumba, S. (2025). Unified Intelligence: Building an Integrated Data Lakehouse for Enterprise-Wide Decision Empowerment. Journal Of Engineering And Computer Sciences, 4(7), 561-567.
25. Kandula, N. Evolution and Impact of Data Warehousing in Modern Business and Decision Support Systems
26. Adari, V. K. (2020). Intelligent Care at Scale AI-Powered Operations Transforming Hospital Efficiency. International Journal of Engineering & Extended Technologies Research (IJEETR), 2(3), 1240-1249.
27. Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2015). Apache Flink™: Stream and batch processing in a single engine. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering.*