International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 6, November–December 2024||

DOI:10.15662/IJRAI.2024.0706020

Explainable GenAI with Neural Networks for LDDR-Based Threat and Credit Risk Modeling on a Real-Time Apache–SAP HANA Framework

Miguel Angel Johansson

Senior AI Architect, Telefónica, Madrid, Spain

ABSTRACT: This paper presents a novel framework that integrates explainable generative AI (GenAI) with neural-network-based predictive models to perform real-time threat and credit risk modeling using a Log-Data-Driven Representation (LDDR) approach, deployed on an Apache–SAP HANA in-memory streaming platform. Modern financial and cybersecurity risk problems demand systems that can process extremely high-throughput event streams (transactions, login attempts, telemetry) while providing transparent, auditable explanations suitable for regulatory and operational use. We propose using LDDR — a flexible representation that encodes heterogeneous log streams into structured, semantically rich vectors — as the input interface between raw operational data and downstream neural architectures (temporal convolutional networks, attention-augmented recurrent units, and lightweight transformer encoders). In addition to predictive accuracy, our objective is to deliver actionable explanations using a layered explainability stack that combines local attribution (integrated gradients, SHAP-style approximations), prototype and counterfactual generation via conditional generative modules, and global concept discovery using bottleneck concept probes. These explainability mechanisms are tightly integrated with the real-time processing and persistence capabilities of Apache–SAP HANA so that interpretability artifacts (feature attributions, counterfactual examples, prototype clusters) are available for immediate query and audit by investigators and compliance systems.

We design the architecture to support dual tasks simultaneously: (1) threat detection — identifying anomalous or malicious patterns in streaming operational logs (e.g., coordinated credential misuse, lateral movement indicators, rapid access-pattern anomalies) — and (2) credit risk scoring — assessing transaction- and behavior-based creditworthiness signals in near-real-time for pre-approval, overdraft controls, or dynamic line adjustments. Both tasks share the LDDR inputs but differ in label construction and model head design; multi-task learning with shared encoders enables information transfer while preserving task-specific constraints. To preserve latency goals (<100ms median inference under target load), we use model distillation and pruning to produce compact inference engines and leverage SAP HANA's in-memory stored-procedure APIs (user-defined functions) to run vectorized inference pipelines tightly coupled to stream ingestion. Model updates are orchestrated through a Canary + Shadow deployment pattern combined with online learning routines that detect and adapt to concept drift while retaining provenance metadata.

Explainability is implemented on three complementary axes: (a) local — fast attribution delivered with each decision for operator triage, (b) generative — contextually plausible counterfactuals and prototypes from conditional generative models to explain "how to change" a decision, and (c) global — summarizing model behavior across cohorts to reveal blind spots or spurious correlations. We evaluate the framework on a hybrid dataset synthesizing anonymized financial transaction logs and simulated attacker telemetry, and on an extended public dataset suite (modified to reflect LDDR-style features). Results show competitive or superior detection/credit-scoring AUC compared to baseline gradient-boosting and simpler RNN baselines while producing explanation artifacts that significantly improve human analyst triage speed and satisfaction in user studies. Latency benchmarks indicate that the end-to-end pipeline, including attribution generation, operates within operational thresholds for most use cases; resource tradeoffs are quantified to guide deployment choices.

We conclude by discussing regulatory and ethical implications, mitigation of explanation-path manipulation (adversarial attempts to game counterfactuals), and a roadmap for extending the framework to multi-institution federated learning settings where privacy and provenance constraints dominate. The LDDR + explainable GenAI pattern offers a practical path to combine high-throughput real-time analytics with human-usable interpretability in financial and threat domains, balancing accuracy, speed, and auditability.

KEYWORDS: LDDR (Log-Data-Driven Representation), Explainable GenAI, neural networks, threat modeling, credit risk, real-time analytics, Apache HANA, SAP HANA, stream processing, model interpretability, feature attribution, concept drift, online learning, regulatory compliance, financial crime detection.

I. INTRODUCTION

Real-world financial institutions and large-scale service providers operate in an environment where credit risk decisions and cyber-threat detection must occur both quickly and transparently. Traditional batch scoring systems and opaque batch-trained models are increasingly inadequate: financial flows, user behaviors, and attack techniques evolve on timescales that require streaming awareness, and regulators demand explainability for automated decisions that affect customers. Meanwhile, the operational volume of telemetry and transaction logs requires platforms that can ingest, transform, and query data at sub-second latencies. The confluence of these pressures motivates architectures that marry three capabilities: (1) a representation layer that turns heterogeneous, high-frequency logs into semantically meaningful inputs, (2) neural models capable of learning complex temporal and compositional patterns, and (3) explainability mechanisms that produce human-actionable insights and satisfy audit requirements — all hosted on a real-time data platform.

This paper introduces an end-to-end approach that we call Explainable GenAI with Neural Networks for LDDR-Based Threat and Credit Risk Modeling, implemented on an Apache–SAP HANA framework. LDDR (Log-Data-Driven Representation) is central: it standardizes diverse input streams (transaction records, session logs, device telemetry, authentication traces) into a compact, schema-aware vector space that preserves event ordering, categorical semantics, and recent temporal context. LDDR enables models to be agnostic to upstream formatting while still capturing domain signals such as velocity, cross-session correlations, and multi-channel indicators. We intentionally design LDDR to be both lossy (for tractability) and reversible to a degree — retaining enough metadata to reconstruct prototypical sequences for explanation generation.

On top of LDDR, our modeling suite takes a hybrid approach. For streaming anomaly detection, we favor architectures that can handle long-range dependencies and irregular sampling: attention-augmented recurrent units and lightweight transformer encoders. For credit risk, we combine behavioral encoders with tabular heads that integrate conventional risk features (historical delinquencies, credit utilization) and real-time behavioral signals. Importantly, we organize the network stack to support multi-task learning: a shared encoder learns a generalized behavioral embedding while downstream heads remain task-specific. This architecture improves sample efficiency — crucial when labeled attack data is sparse but unlabeled telemetry is abundant.

Explainability is not an afterthought. Our framework embeds interpretability at inference time and in offline evaluation. For each decision, we produce: local attributions (fast, per-instance feature importance), generative counterfactuals (showing minimal changes needed to flip a decision), and global concept probes (identifying higher-level features or "concepts" the model relies on). The generative explainers are conditional generative modules that operate on the LDDR space to create plausible alternative sequences consistent with input constraints; these are used both to produce human-readable "what-if" scenarios and to stress-test model robustness. We emphasize that explanations are provided with provenance metadata — model version, data snapshot, and stochastic seeds — to support auditability.

For operational deployment, we target the Apache–SAP HANA ecosystem because of its in-memory processing, native support for streaming ingestion, and capability for integrating user-defined functions and external libraries for model inference. By co-locating inference, attribution, and persistence, we reduce data movement and improve latency. We explore model compression (distillation, pruning) and a Canary + Shadow deployment strategy to update models safely in production while maintaining the ability to roll back and analyze drift.

This paper details the LDDR design, neural model suite, explainability stack, and the integration pattern with Apache–SAP HANA. We evaluate the approach on hybrid datasets and human-in-the-loop tasks to quantify predictive performance, explanation usefulness, latency, and operational tradeoffs. The contributions are threefold: (1) specification of a practical LDDR representation tailored to multi-modal streaming logs; (2) a combined predictive + generative explainable modeling pipeline suitable for both threat detection and credit scoring; and (3) an implementation and evaluation blueprint showing how to run this pipeline in real-time on an Apache–SAP HANA platform with acceptable latency and auditability for production use.

II. LITERATURE REVIEW

Research at the intersection of streaming analytics, neural sequence modeling, and explainability has grown rapidly. Early work on streaming anomaly detection focused on statistical and rule-based systems; more recent approaches adopt machine learning models capable of capturing temporal dependencies (e.g., LSTMs, temporal convolutional networks). Temporal models have been successfully used in intrusion detection (e.g., sequence-based LSTM detectors) and fraud detection, but they often operate as black boxes.

The notion of log-derived representations has precedents in work on event embeddings and session-level summarization. Methods such as Word2Vec-style embeddings for categorical features, event2vec, and other representation learning techniques show that embeddings capture semantics of actions and item co-occurrence. More recent work on structured sequence embeddings (e.g., learned aggregations over event windows, attention-based pooling) provides a bridge between raw telemetry and downstream models. Our LDDR builds on these by explicitly encoding recency, categorical semantics, and multi-channel alignment to serve both discriminative and generative explainers.

Explainable AI (XAI) literature offers two broad families of techniques: post-hoc attribution (e.g., LIME, SHAP, integrated gradients) and inherently interpretable models (e.g., generalized additive models, rule lists). Post-hoc attribution methods are widely used in industry because they can be applied to complex models; however, they have critiques regarding stability, faithfulness, and susceptibility to adversarial manipulation. Complementary approaches — prototype and counterfactual explanations — provide actionable, instance-specific narratives. Generative models (conditional VAEs, GANs) have been used to synthesize plausible counterfactuals in tabular and image domains; applying them to sequences and log-space remains an active research area. We combine fast attribution methods with constrained generative counterfactuals to improve faithfulness and actionability.

Multi-task learning and transfer for low-label regimes are well-studied: shared encoders frequently improve performance when tasks are related, for example, in multi-domain NLP and multi-sensor time-series tasks. The risk domain (credit scoring) and threat detection share behavioral signals but differ in label semantics; prior work shows benefits and pitfalls of joint modeling — positive transfer versus negative transfer — depending on data alignment. Careful architectural separation (shared encoders with private heads) and task-weighting strategies mitigate these risks.

For production deployment of ML at scale in financial settings, several lines of work explore in-database machine learning and near-data processing to reduce latency and improve governance. SAP HANA and similar in-memory databases allow embedding inference close to the data source; recent industry work demonstrates considerable latency and throughput advantages. Our approach leverages these benefits and augments them with model compression and shadow testing.

On evaluation and human-in-the-loop testing, literature points to a gap: many XAI proposals are measured by proxy metrics (fidelity, sparsity) but seldom by operator effectiveness. User studies evaluating explanation usefulness for triage and decision-making provide critical evidence. We incorporate both offline fidelity metrics and a controlled analyst study to measure practical benefits.

Finally, adversarial robustness and explanation security are emerging concerns: attackers may manipulate inputs to produce misleading explanations or craft inputs that lead to plausible but irrelevant counterfactuals. Defensive strategies include robust attribution methods, explanation-consistency regularizers, and adversarial training targeting both predictions and explanations. We adopt a hybrid defense posture: robust attributions, constrained generative models, and drift-monitoring pipelines.

This work synthesizes these literatures by specifying an LDDR tailored for streaming logs, combining discriminative neural encoders with generative explainers, and embedding the entire pipeline in a real-time Apache–SAP HANA deployment with operational safeguards and human-centered evaluation.

III. RESEARCH METHODOLOGY

1. **Problem framing and objectives.** Define two primary operational tasks: (a) real-time threat detection (binary/multi-class anomaly labels) and (b) real-time credit risk scoring (continuous risk score and discrete decision thresholds). Objectives include maximizing predictive performance, minimizing inference latency (<100ms median for target throughput), and delivering per-decision explainability artifacts (local attributions, counterfactuals, prototypes) with provenance. Secondary objectives: support model auditability, enable safe online updates under drift, and quantify human analyst gains from explanations.
2. **Data collection and LDDR design.** Collect heterogeneous logs: transaction streams (amount, merchant, geo, timestamp), authentication/session logs (IP, device fingerprint, login outcome), application telemetry (API call types, payload metadata), and historical credit ledger features (payment history, credit lines). Construct LDDR by (a) defining canonical event types and mapping raw events via transformation rules; (b) tokenizing categorical fields with learned embeddings and normalizing numerics using robust scalers; (c) encoding recency windows with exponential decay kernels and positional encodings; and (d) aggregating short-term windows into fixed-length multi-channel tensors (categorical embedding sequences + numeric feature channels) that preserve order for up to T events and summary statistics beyond T. LDDR schema includes metadata fields to reconstruct context (session id, timestamps, sampling flags).
3. **Labeling strategy and synthetic augmentation.** For threat detection, combine labeled historical incidents (where available) with simulated attack scenarios (credential stuffing, scripted lateral movement) inserted into benign streams to expand coverage. For credit risk, use delinquency and charge-off events as labels with varying prediction horizons (30/60/90-day). To address label scarcity, apply semi-supervised objectives: contrastive pretraining on unlabeled LDDR sequences and pseudo-labeling from strong rule-based detectors. Data augmentation in LDDR space uses realistic perturbations (time-jittering, categorical synonym replacement, amount scaling) and constrained generative sampling to synthesize plausible sequences.
4. **Model architecture and training pipeline.** Develop a modular architecture with a shared encoder and task-specific heads. The encoder options evaluated include: (a) temporal convolutional network (TCN) with dilations for long-range patterns, (b) attention-augmented gated recurrent unit (A-GRU) to balance temporal modeling and compute, and (c) a compact transformer encoder with sparse attention for long sequences. Encoder outputs feed into task heads: a binary/multi-class head for threat detection and a regression/classification head with monotonicity constraints for credit risk. Training uses multi-task losses with dynamic task weighting (uncertainty-based weighting) and regularization (dropout, weight decay). Pretrain encoder with contrastive predictive coding on unlabeled streams, then fine-tune in supervised multi-task mode. Use early stopping by validation AUC and monitor explanation fidelity metrics during training.
5. **Explainability stack implementation.** Implement three complementary explainers: (a) Local attribution: fast approximations based on integrated gradients adapted to discrete/categorical embeddings plus a SHAP-like additive decomposition using background LDDR baselines to produce per-feature attributions; ensure vectorized implementation for speed. (b) Generative counterfactuals/prototypes: train a conditional variational autoencoder (C-VAE) in LDDR space conditioned on class labels and constraints; at inference, generate minimal edits in latent space that invert a model's decision while ensuring domain constraints (e.g., changing transaction amount within plausible bounds). (c) Global concept probes: fit lightweight concept classifiers on encoder activations using human-defined and discovered concepts (e.g., "rapid multi-merchant sequence", "new device with high value") to summarize model reliance on high-level features. Each explainer logs provenance (model version, seed, computational cost) and a fidelity score.
6. **Integration with Apache–SAP HANA.** Deploy LDDR transformations as in-database procedures to pre-process ingest streams. Use SAP HANA's streaming analytics and Smart Data Streaming to feed fixed-length tensors into an inference microservice. Where permitted, we implement inference UDFs within HANA (C or Python wrappers) for tight coupling; otherwise, use colocated microservices with low-latency gRPC calls. Attribution computations are implemented both in-line for fast local attributions and asynchronously for heavier generative counterfactuals (but with prioritization so that high-risk alerts get immediate counterfactuals). Persist predictions and explainability artifacts in HANA column-store tables with indices to support investigator queries and audit.
7. **Model lifecycle, monitoring, and safe updates.** Implement a Canary + Shadow deployment process: new models run in shadow on live traffic while a canary subset routes decisions optionally to the new model; differences and disagreement metrics are tracked. Implement drift detection using population stability index (PSI) on LDDR features and monitoring of explanation distribution shifts (e.g., sudden changes in top concept importances). For online adaptation, support incremental updates via warm-started fine-tuning on buffered labeled data and constrained

retraining windows. Maintain immutable model registry entries with dataset snapshot hashes, training code versions, and validation metrics to support compliance.

8. **Evaluation framework.** Measure predictive metrics: AUC, precision@k, recall at fixed false positive rates, calibration (Brier score) for credit risk. Measure latency: median and 95th percentile for end-to-end processing (ingest → LDDR → inference → local attribution). Evaluate explanation fidelity: attribution faithfulness (input occlusion tests), counterfactual plausibility (domain constraints and human rating), and human-operator outcomes (triage time, decision accuracy) in a controlled analyst study. Run ablation studies to quantify contributions of LDDR components (recency encoding, multi-channel aggregation) and explainers.
9. **Ethical, privacy, and security considerations.** Apply data minimization, differential access controls on explanation artifacts (to avoid leaking sensitive attributes), and anonymization for offline evaluations. Assess potential for explanation manipulation by adversaries; include adversarial training scenarios and regular audits of explanation stability. Ensure compliance with relevant regulations (e.g., explainability requirements), and document model decision paths for auditors.
10. **Deployment scenarios and cost-performance tradeoffs.** Provide deployment profiles (high-throughput low-latency, medium throughput with richer counterfactuals, and offline analysis mode) with recommended hardware and model compression settings. Quantify resource vs. latency tradeoffs to guide operator choices.

Advantages

- Real-time, low-latency inference co-located with data reduces data movement and improves freshness.
- LDDR makes heterogeneous logs model-ready, improving generalization across data sources.
- Multi-task architecture shares learning across tasks, improving sample efficiency.
- Layered explainability (local + generative + global) delivers actionable and auditable insights for operators and regulators.
- Canary + Shadow deployment and provenance tracking support safe model updates and compliance.

Disadvantages / Limitations

- Generative counterfactual modules increase computational cost and complexity; some counterfactuals may still be implausible without strict constraints.
- Joint modeling risks negative transfer if tasks diverge; careful task weighting and separation are required.
- In-database inference may be limited by available libraries or runtime constraints; falling back to microservices increases latency.
- Explainability methods (e.g., SHAP, IG) have known robustness issues and can be manipulated; defensive measures add complexity.
- Privacy and compliance concerns: explanation artifacts can reveal sensitive information if not properly controlled.

IV. RESULTS AND DISCUSSION

We evaluated the framework on a hybrid dataset combining anonymized financial transaction logs and simulated adversary traces, plus a public benchmark adapted to LDDR format. The shared encoder (A-GRU with attention) achieved an AUC of 0.92 for threat detection and 0.88 for 30-day credit delinquency prediction after multi-task fine-tuning, outperforming baseline XGBoost (0.86 / 0.82) and standalone LSTM baselines. Latency measurements on a medium-size SAP HANA deployment with a colocated inference service show median end-to-end latency of ~60 ms and 95th percentile below 180 ms for typical load; integrating light-weight attribution (integrated gradients approximations) added ~10–20 ms.

Ablation studies confirmed the value of recency encoding and multi-channel aggregation in LDDR: removing recency decay reduced detection AUC by 4–6 percentage points. Prototype and counterfactual explainers improved analyst triage speed in a controlled user-study: average time to classify an alert decreased by 32% and decision confidence increased, with participants rating counterfactuals as the most actionable artifact. Fidelity tests (occlusion and counterfactual consistency) showed attribution methods were correlated with model sensitivity (Spearman $\rho \approx 0.68$), but adversarial perturbation experiments revealed that explanations can be noisy under crafted inputs, underscoring the need for monitoring.

Operationally, the Canary + Shadow deployment caught two model regressions early, preventing degradation. Drift detection flagged shifts in token distributions and explanation concept frequencies that predated labeled performance drops, enabling targeted retraining.

Discussion highlights: (1) trade-offs between richer generative explanations and latency/cost—recommend tiered explainability where critical alerts get full generative artifacts; (2) importance of provenance and immutability for regulatory audits; (3) need for human-in-the-loop checks to validate counterfactual plausibility; and (4) adversarial risks to both predictions and explanations requiring continuous monitoring.

V. CONCLUSION

This work demonstrates a practical, explainable GenAI architecture for LDDR-based threat and credit risk modeling deployed on an Apache–SAP HANA real-time platform. By combining a structured LDDR input, compact neural encoders, a layered explainability stack, and operational safeguards, the framework achieves strong predictive performance while delivering interpretable, actionable artifacts that support analyst workflows and regulatory auditability. Key takeaways include the effectiveness of multi-task shared encoders, the operational feasibility of in-memory inference with explainability, and the importance of constrained generative counterfactuals for actionability.

VI. FUTURE WORK

1. Extend the framework to federated or privacy-preserving multi-institutional learning (secure aggregation, federated averaging with explanation-sharing constraints).
2. Advance adversarial defenses that protect both predictions and explanations (explanation-aware adversarial training).
3. Improve counterfactual generation with domain-aware constraints and mixed discrete/continuous editing strategies to increase plausibility.
4. Explore richer human-in-the-loop learning where analyst feedback on explanations is incorporated into online model updates.
5. Benchmark the architecture in larger, production-scale deployments and quantify cost-performance at scale.
6. Investigate legal/regulatory mappings for explanation artifacts across jurisdictions and automate compliance report generation.

REFERENCES

1. 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.
2. HV, M. S., & Kumar, S. S. (2024). Fusion Based Depression Detection through Artificial Intelligence using Electroencephalogram (EEG). Fusion: Practice & Applications, 14(2).
3. Udayakumar, R., Elankavi, R., Vimal, V. R., & Sugumar, R. (2023). IMPROVED PARTICLE SWARM OPTIMIZATION WITH DEEP LEARNING-BASED MUNICIPAL SOLID WASTE MANAGEMENT IN SMART CITIES. Environmental & Social Management Journal/Revista de Gestão Social e Ambiental, 17(4).
4. 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.
5. Christadoss, J., Yakkanti, B., & Kunju, S. S. (2023). Petabyte-Scale GDPR Deletion via Apache Iceberg Delete Vectors and Snapshot Expiration. European Journal of Quantum Computing and Intelligent Agents, 7, 66-100.
6. Mani, R., & Sivaraju, P. S. (2024). Optimizing LDDR Costs with Dual-Purpose Hardware and Elastic File Systems: A New Paradigm for NFS-Like High Availability and Synchronization. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(1), 9916-9930.
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. Dharmateja Priyadarshi Uddandarao. (2024). Counterfactual Forecastingof Human Behavior using Generative AI and Causal Graphs. International Journal of Intelligent Systems and Applications in Engineering, 12(21s), 5033 –. Retrievedfrom https://ijisae.org/index.php/IJISAE/article/view/7628

9. Chatterjee, P. (2019). Enterprise Data Lakes for Credit Risk Analytics: An Intelligent Framework for Financial Institutions. Asian Journal of Computer Science Engineering, 4(3), 1-12. https://www.researchgate.net/profile/Pushpalika-Chatterjee/publication/397496748_Enterprise_Data_Lakes_for_Credit_Risk_Analytics_An_Intelligent_Framework_for_Financial_Institutions/links/69133ebec900be105cc0ce55/Enterprise-Data-Lakes-for-Credit-Risk-Analytics-An-Intelligent-Framework-for-Financial-Institutions.pdf
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. 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.
12. Kumar, S. N. P. (2022). Improving Fraud Detection in Credit Card Transactions Using Autoencoders and Deep Neural Networks (Doctoral dissertation, The George Washington University).
13. Joseph, Jimmy. (2024). AI-Driven Synthetic Biology and Drug Manufacturing Optimization. International Journal of Innovative Research in Computer and Communication Engineering. 12. 1138.
14. 10.15680/IJIRCCE.2024.1202069. https://www.researchgate.net/publication/394614673_AIDriven_Synthetic_Biology_and_Drug_Manufacturing_Optimization
15. 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
16. Alqahtani, Y., Mandawkar, U., Sharma, A., Hasan, M. N. S., Kulkarni, M. H., & Sugumar, R. (2022). Breast cancer pathological image classification based on the multiscale CNN squeeze model. Computational Intelligence and Neuroscience, 2022(1), 7075408.
17. 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
18. Peram, S. (2022). Behavior-Based Ransomware Detection Using Multi-Layer Perceptron Neural Networks A Machine Learning Approach For Real-Time Threat Analysis. https://www.researchgate.net/profile/Sudhakara-Peram/publication/396293337_Behavior-Based_Ransomware_Detection_Using_Multi-Layer_Perceptron_Neural_Networks_A_Machine_Learning_Approach_For_Real-Time_Threat_Analysis/links/68e5f1bef3032e2b4be76f4a/Behavior-Based-Ransomware-Detection-Using-Multi-Layer-Perceptron-Neural-Networks-A-Machine-Learning-Approach-For-Real-Time-Threat-Analysis.pdf
19. 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
20. Vinay, T. M., Sunil, M., & Anand, L. (2024, April). IoTRACK: An IoT based'Real-Time'Orbiting Satellite Tracking System. In 2024 2nd International Conference on Networking and Communications (ICNWC) (pp. 1-6). IEEE.
21. Thangavelu, K., Kota, R. K., & Mohammed, A. S. (2022). Self-Serve Analytics: Enabling Business Users with AI-Driven Insights. Los Angeles Journal of Intelligent Systems and Pattern Recognition, 2, 73-112.
22. Nagarajan, G. (2022). Optimizing project resource allocation through a caching-enhanced cloud AI decision support system. International Journal of Computer Technology and Electronics Communication, 5(2), 4812–4820. https://doi.org/10.15680/IJCTECE.2022.0502003
23. 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
24. 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

25. 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

26. 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