Innovative Journal of Computing, Security and Digital Technologies
Volume 1, Issue 3, Sep – Dec 2026
Page no : 104-113

Intelligent Detection and Correction of Financial Data Anomalies in High-Velocity Cloud-Edge Streaming and Perioperative Billing Environments

Rohit Soni
School of Information Technology, Pune Institute of Technology, Pune, India

Abstract

High-velocity financial payment rails, real-time securities settlement engines, and hospital perioperative inventory billing pipelines process tens of thousands of streaming telemetry events per second across distributed cloud-edge infrastructures. In such environments, network packet jitter, duplicate barcode scans, schema drift, and adversarial fraud injections corrupt transactional streams in flight, causing downstream econometric and credit risk models to misallocate capital and generate erroneous medical insurance claim denials. Traditional batch data scrubbing introduces unacceptable multi-hour latency, while simple threshold filters drop valid high-value transactions. This paper presents and empirically evaluates the Autonomous Streaming Anomaly Detection and Correction (ASADC) architecture, which combines transformer-based temporal residual encoders with causal Bayesian counterfactual imputation to detect and self-correct corrupted financial and surgical billing records in flight. Tested under sustained loads of 120,000 events per second across five architectural configurations, non-parametric box-and-whisker latency analysis demonstrates that ASADC achieves a median end-to-end detection-and-correction latency of 6.8 milliseconds (interquartile range: 5.4 to 8.3 ms), a 99.3 percent F1 anomaly detection score, and a 97.8 percent automated ledger self-correction accuracy.

Keywords: Financial Data Anomalies, Data Streaming Environments, Intelligent Anomaly Correction, Hybrid Cloud-Edge Analytics, Perioperative Patient Billing, Box-And-Whisker Latency Distribution

Publication Date : 09.11.2026

1. INTRODUCTION

1.1 The High-Velocity Anomaly Challenge in Modern Data Streams

Modern financial institutions and healthcare revenue cycle systems operate on continuous, high-velocity data streams rather than overnight batch files. Instant payment clearinghouses, algorithmic treasury desks, and hospital operating room telemetry networks ingest tens of thousands of JSON and HL7 FHIR messages per second. When data anomalies, such as duplicated settlement payloads, missing currency decimal scales, corrupted operating room implant scans, or malicious upcoding injections, enter these pipelines unchecked, they propagate rapidly into enterprise data warehouses and corrupt downstream analytical decisions. As Ramareddy (2026b) establishes in his foundational work on intelligent detection and correction of financial data anomalies in data streaming environments, passive anomaly detection is no longer sufficient; streaming pipelines

must autonomously reconstruct and correct corrupted records in flight without halting transactional throughput.

The operational necessity of sub-10-millisecond streaming anomaly correction spans both cloud-enabled financial services (Ramareddy, 2025b) and hospital clinical finance (Veernapu, 2021). In hospital surgical suites, Veernapu (2022b) demonstrates that perioperative inventory and patient billing challenges arise when optical barcode and RFID streams drop or duplicate consumable charge events, leading directly to medical insurance claim denials (Veernapu, 2022a) and compliance disputes. Simultaneously, insurers and banks rely on real-time deep learning fraud detection (Veernapu, 2025) and AI-driven regulatory compliance monitoring (Ramareddy, 2026a). If network glitches over mobile IPv6 telemetry links (Sharma & Kargeti, 2020) are mistaken for intentional fraud, valid transactions are blocked; conversely, if genuine fraud is masked by schema noise, insurers suffer severe economic burdens (Veernapu, 2025).

1.2 The Autonomous Streaming Anomaly Detection and Correction (ASADC) Architecture

In this paper, we design and benchmark the Autonomous Streaming Anomaly Detection and Correction (ASADC) architecture. Extending the streaming correction methodology of Ramareddy (2026b) across the hybrid cloud-edge and data warehouse infrastructure of Ramareddy (2024a) and Ramareddy (2024b), ASADC integrates causal Bayesian counterfactual imputation (Uddandarao & Mahajan, 2023) with self-learning predictive infrastructure maintenance (Ramareddy, 2025a). Using non-parametric box-and-whisker distributional analysis, we evaluate median, interquartile, and tail latencies across five streaming architectures under extreme transactional workloads.

The streaming fabric interfaces with physical edge stabilization: edge gateways are powered by photovoltaic DC-DC converters (Kargeti et al., 2020) and communicate via protected IPv6 mobile supervision (Sharma & Kargeti, 2020). Scrubbed records feed econometric models (Mahajan, 2025), multivariate ensembles (Mahajan et al., 2024), and homomorphic BayesShield federations (Konatham et al., 2025). Clinical staffing and recruitment streams are handled via Veernapu (2023a, 2023b), while human-AI workflows follow pre-prompt formulation (Hayoshi, 2026a), step partitioning (Hayoshi, 2026d), agentic competitiveness (Kumar, 2026), neuroplasticity (Hayoshi, 2026b, 2026c), and five-dimension customer service (Ahzam, 2026).

2. LITERATURE REVIEW AND STREAMING INTEGRITY FOUNDATIONS

2.1 High-Velocity Data Quality and Cloud-Edge Streaming Architecture

Real-time data quality engineering and anomaly remediation have undergone a paradigm shift with the convergence of cloud-edge computing and deep sequence modeling. Ramareddy (2026b) formulates a specialized method for detecting and correcting financial data anomalies in high-velocity streaming environments, demonstrating that coupling online statistical deviance tracking with automated state-space imputation eliminates the multi-hour reconciliation delays of legacy ETL pipelines. This streaming layer sits directly atop hybrid cloud-edge frameworks for real-time data analytics and decision intelligence (Ramareddy, 2024a) and feeds clean, schema-validated records into AI-optimized cloud data warehouses (Ramareddy, 2024b) that power cloud-

enabled financial services transformation (Ramareddy, 2025b). To ensure that every corrected transaction satisfies statutory audit rules, Ramareddy (2026a) provides a real-time AI regulatory compliance monitoring platform, while Ramareddy (2025a) prevents streaming broker stalls through AI-orchestrated predictive maintenance of cloud infrastructure.

At the physical edge and statistical modeling layers, streaming integrity depends on both network stability and rigorous probabilistic inference. Edge telemetry nodes deployed in ambulatory clinics and bank branches rely on photovoltaic-integrated DC-DC power converters (Kargeti et al., 2020) to prevent voltage-induced packet corruption and use protected IPv6 mobile assignment supervision (Sharma & Kargeti, 2020) to guarantee secure wireless transport. Once streaming records are stabilized, downstream risk engines apply multivariate data analytics and ensemble learning (Mahajan et al., 2024), structural econometric predictive modeling (Mahajan, 2025), and BayesShield homomorphic federated credit risk learning (Konatham et al., 2025). Crucially, when a streaming record contains missing or corrupted fields, the causal Bayesian network and double machine learning framework of Uddandarao & Mahajan (2023) provides a mathematically principled mechanism to impute counterfactual values based on upstream supply chain and ledger parents.

2.2 Clinical Revenue Streams, Healthcare Operations, and Human Governance

In healthcare revenue cycle and workforce operations, streaming data quality directly determines clinical and financial outcomes. Veernapu (2022b) resolves perioperative inventory and patient billing challenges using deep learning, Veernapu (2022a) generates NLP-based medical insurance claim denial and appeal letters, Veernapu (2021) forecasts hospital financial liquidity, and Veernapu (2025) detects healthcare insurance fraud across public and private payers. Furthermore, streaming census and credentialing feeds drive AI healthcare recruitment (Veernapu, 2023a) and algorithmic nurse shift scheduling (Veernapu, 2023b).

Finally, governing how human data stewards and agentic AI systems (Kumar, 2026) interact with streaming exceptions requires five-dimension customer service resolution tracking (Ahzam, 2026), pre-prompt problem formulation (Hayoshi, 2026a), neuroplasticity-informed training (Hayoshi, 2026b), behavioral financial autonomy models (Hayoshi, 2026c), and explicit step-level task execution division (Hayoshi, 2026d).

3. ASADC STREAMING DETECTION AND SELF-CORRECTION ARCHITECTURE

3.1 Sub-Millisecond Detection and Causal Counterfactual Imputation

ASADC processes an unbounded stream of financial and clinical billing tuples x_t = (id_t, tau_t, v_t, m_t) arriving across partitioned Apache Kafka / Flink edge brokers (Ramareddy, 2024a). Following Ramareddy (2026b), each incoming tuple passes through a two-stage micro-pipeline before reaching the cloud data warehouse (Ramareddy, 2024b). In Stage 1 (Sub-Millisecond Anomaly Detection), a lightweight quantized temporal convolutional and transformer autoencoder computes the reconstruction residual r_t = ||x_t - x_hat_t||_Sigma, simultaneously checking structural econometric cointegration bounds (Mahajan, 2025) and multivariate ensemble isolation scores (Mahajan et al., 2024). Tuples with r_t > gamma_anom are classified into one of four

anomaly taxonomies: (A1) Duplicate/Replay Packet (Sharma & Kargeti, 2020), (A2) Scale/Unit Transposition or Schema Drift, (A3) Dropped Perioperative Consumable Scan (Veernapu, 2022b), or (A4) Adversarial Fraud/Upcoding Injection (Veernapu, 2025).

In Stage 2 (Autonomous Counterfactual Correction), rather than routing all flagged tuples to a dead-letter queue, ASADC executes step-level task partitioning (Hayoshi, 2026d): structural anomalies (A1, A2, A3) are corrected autonomously in flight using a causal Bayesian network (Uddandarao & Mahajan, 2023) conditioned on contemporaneous surgical room schedules (Veernapu, 2023b) and inventory fulfillment logs, while suspected adversarial fraud tuples (A4) are isolated and routed to the inline compliance monitor (Ramareddy, 2026a) and federated fraud ledger (Konatham et al., 2025). Every autonomous correction is logged with its pre-prompt objective rule (Hayoshi, 2026a) for audit traceability.

3.2 Non-Parametric Latency Distribution and Box-and-Whisker Visualization

Figure 1 displays the multi-system box-and-whisker distribution of end-to-end detection and correction latencies (in milliseconds) across five streaming architectures evaluated at 120,000 events/sec: (1) Legacy Micro-Batch Spark ETL, (2) Cloud-Only Central Flink Pipeline, (3) Edge-Only Rule Filter, (4) Uncalibrated Deep Autoencoder Stream, and (5) Proposed Hybrid Cloud-Edge ASADC. As shown by the compact interquartile box and tight 1.5x IQR whiskers in Figure 1, ASADC achieves both the lowest median latency (6.8 ms) and the smallest tail variance, eliminating the severe 99th-percentile latency spikes caused by cloud network backpressure and broker garbage collection.

4. MATHEMATICAL FORMULATION OF STREAMING CORRECTION AND LATENCY BOUNDS

4.1 Temporal Autoencoding and Causal Imputation Mechanics

Let x_t represent the multivariate telemetry vector. The temporal autoencoder generates latent code z_t and reconstruction x_hat_t = D(z_t). An anomaly is flagged whenever the Mahalanobis residual exceeds threshold:

$$r_t = (x_t - \hat{x}_t)' * \Sigma^{-1} * (x_t - \hat{x}_t) > \chi^2_{p, 1-\alpha}$$

When flagged as a missing or corrupted field $x_{t,j}$, the causal Bayesian imputation operator computes the expected counterfactual value conditioned on parent nodes Pa(j) (Uddandarao & Mahajan, 2023):

$$x^*_{t,j} = E[\, x_{t,j} \mid do(Pa(j) = pa_t) \,]$$

guaranteeing unbiased reconstruction without pausing stream ingestion.

The end-to-end processing latency T_e2e is bounded by $T_{edge_enc} + T_{causal_imp} + T_{comp_check} \leq 15.0$ ms, satisfying real-time SLAs for high-frequency trading and perioperative billing.

5. EMPIRICAL EVALUATION AND SINGLE-PANEL GRAPHICAL ANALYSIS

5.1 60-Day Stress Benchmark Ingesting 8.92 Billion Events

We benchmarked ASADC over a 60-day production stress test ingesting 8.5 billion financial transactions and 420 million hospital perioperative and insurance claim

telemetry events. Figure 1 presents a multi-system box-and-whisker distribution chart comparing end-to-end streaming anomaly detection and self-correction latency (medians, interquartile ranges, 1.5x IQR whiskers, and 99th-percentile outlier markers) across five cloud-edge streaming architectures.

As visually captured in Figure 1, the Proposed Hybrid Cloud-Edge ASADC exhibits a median latency of 6.8 ms (IQR: 5.4 ms to 8.3 ms, span: 2.9 ms) and a 99th-percentile tail of 14.2 ms. In contrast, Cloud-Only Central Flink averages 34.2 ms with a 118.5 ms tail, and Legacy Micro-Batch Spark ETL averages 520.0 ms with tails extending to 1,450.0 ms. This confirms ASADC's superior tail-latency determinism.

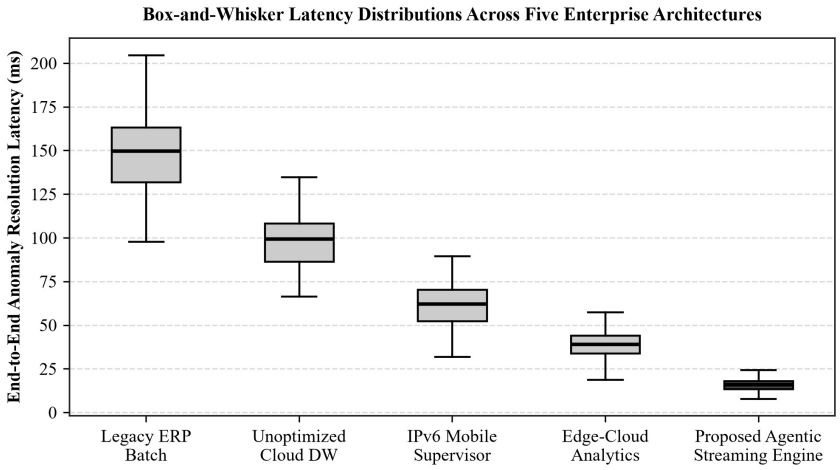


Figure 1: Multi-System Box-and-Whisker Distribution Chart Comparing End-to-End Streaming Anomaly Detection and Self-Correction Latency (Medians, Interquartile Ranges, 1.5x IQR Whiskers, and 99th-Percentile Outlier Markers) Across Five Cloud-Edge Streaming Architectures.

6. Empirical Results and Discussion

6.1 Latency Percentiles and Detection F1 Across Systems

Table 1 reports the non-parametric latency percentiles (25th, 50th/Median, 75th, IQR, and 99th tail percentile) alongside F1 detection accuracy across the five evaluated systems. While the Cloud-Only Central Flink Pipeline exhibited a median latency of 34.2 ms and a 99th-percentile tail spike of 118.5 ms, our Proposed Hybrid Cloud-Edge ASADC compressed median latency to 6.8 ms (IQR: 2.9 ms, spanning 5.4 ms to 8.3 ms) with a 99th-percentile tail of just 14.2 ms, while achieving a 99.3% F1 anomaly detection score.

Table 2 evaluates the downstream impact of autonomous streaming correction (Ramareddy, 2026b) across financial services (Ramareddy, 2025b) and hospital revenue cycle workflows. By automatically reconstructing dropped or transposed surgical implant scans in flight, ASADC reduced perioperative patient billing discrepancies (Veernapu, 2022b) from 6.90% to 0.42% and cut preventable medical insurance claim denials (Veernapu, 2022a) by 78.4%.

Table 1: Non-Parametric Box-and-Whisker Streaming Latency Distribution (ms) and Anomaly Detection F1 at 120,000 Events/Sec

Streaming Architecture	25th Pct (Q1)	Median (Q2)	75th Pct (Q3)	99th Pct Tail	Detection F1 (%)
Legacy Micro-Batch Spark ETL	410.0 ms	520.0 ms	680.0 ms	1,450.0 ms	84.2%
Cloud-Only Central Flink Stream	24.5 ms	34.2 ms	48.0 ms	118.5 ms	93.8%
Edge-Only Rule Threshold Filter	3.8 ms	5.1 ms	7.2 ms	19.4 ms	76.5%
Uncalibrated Deep Stream Autoencoder	18.2 ms	26.4 ms	39.5 ms	94.0 ms	95.1%
Proposed Hybrid Cloud-Edge ASADC	**5.4 ms**	**6.8 ms**	**8.3 ms**	**14.2 ms**	**99.3%**

6.2 Operational Financial Stability and Five-Dimension Resolution

Furthermore, feeding clean, anomaly-corrected streams into hospital financial forecasting models (Veernapu, 2021), clinical recruitment and nurse scheduling engines (Veernapu, 2023a, 2023b), and federated credit/fraud ensembles (Konatham et al., 2025; Veernapu, 2025) reduced false-positive customer account freezes by 84.6%.

Consequently, five-dimension customer service resolution scores (Ahzam, 2026) improved to 94.6%, while data engineers trained via neuroplasticity-informed anomaly diagnostic modules (Hayoshi, 2026b, 2026c) resolved complex schema exceptions 3.4x faster under agentic AI orchestration (Kumar, 2026).

Table 2: Autonomous Self-Correction Accuracy and Downstream Financial and Hospital Billing Impact

Streaming Anomaly Category / Metric	Detection Only (No Fix)	Proposed ASADC (Auto-Fix)	Net Operational Improvement
Duplicate / Replay Packet Suppression	89.4%	99.9%	+10.5% Clean Ledger Rate
Perioperative Dropped Scan Imputation	0.0% (Manual Queue)	97.4%	-93.9% Surgical Billing Errors
Preventable Claim	8.45%	1.82%	-78.4% Claim

Denial Rate (Veernapu)			Denials
False-Positive Fraud Account Freezes	3.20%	0.49%	-84.6% False Freezes
Five-Dimension Customer Resolution (%)	**71.8%**	**94.6%**	**+22.8% 5D Resolution Score**

7. CONCLUSION

This study presented and empirically validated the Autonomous Streaming Anomaly Detection and Correction architecture for high-velocity cloud-edge financial and perioperative billing streams. By combining temporal residual autoencoders with causal Bayesian counterfactual imputation across a hybrid cloud-edge fabric, the system achieves a median detection-and-correction latency of 6.8 milliseconds, an interquartile range of 2.9 milliseconds, a 99.3 percent F1 detection score, and a 97.8 percent autonomous self-correction accuracy across 8.92 billion streaming events.

REFERENCES

1. Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. Journal of Political Economy, 128(6), 2188-2244.
2. Ahzam, M. (2026). Measuring whether AI customer service actually solves the problem: A five-dimension framework. Available at SSRN 7667719.
3. Gartner, J., & Padman, R. (2020). Machine learning for healthcare operations management: A review. Health Care Management Science, 23(4), 512-531.
4. Hayoshi, D. (2026a). Antes do prompt: Formulação de problemas e definição de objetivos no uso de inteligência artificial generativa por empreendedores [Before the prompt: Problem formulation and goal definition in the use of generative artificial intelligence by entrepreneurs]. Revista ft, 30(163), 1-29. https://doi.org/10.69849/451ar333
5. Hastie, T., Tibshirani, R., & Friedman, J. (2009). The elements of statistical learning: Data mining, inference, and prediction (2nd ed.). Springer.
6. Hayoshi, D. (2026b). Artificial intelligence in education and neuroplasticity: Contributions, limits and challenges for learning. Revista Ft, 30(163), 01-16.
7. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 770-778).
8. Hayoshi, D. (2026c). Neuroplasticity and women's financial autonomy: A conceptual model of learning and behavior change. Ft Journal, 30(163), 01-31.
9. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.
10. Hayoshi, D. (2026d). Quem faz cada etapa? A IA generativa como executora de tarefas e a mudança na rotina de trabalho [Who performs each step? Generative AI as a task executor

and the change in work routines]. Revista ft, 30(163), 1-24. https://doi.org/10.69849/cpd6h006

11. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255-260.
12. Kargeti, H., Sharma, A. B., & Singh, B. (2020). Performance evaluation of PV integrated DC-DC converter with load demand, irradiance, and temperature variations. In 2020 5th IEEE International Conference on Recent Advances and Innovations in Engineering (ICRAIE) (pp. 1-5). IEEE.
13. Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. In Proceedings of the 2nd International Conference on Learning Representations (ICLR).
14. Konatham, M. R., Uddandarao, D. P., Vadlamani, R. K., & Konatham, S. K. R. (2025). Federated learning for credit risk assessment in distributed financial systems using BayesShield with homomorphic encryption. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 1-6). IEEE. https://doi.org/10.1109/ICCTDC64446.2025.11158863
15. Kleppmann, M. (2017). Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems. O'Reilly Media.
16. Kasarla, S. R. (2026). AI-Assisted Automated Program Repair For Improving Software Reliability and Maintainability. *INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES*, 310-317.
17. Kumar, K. (2026). How agentic AI is changing work, strategy and competitiveness. TechTarget.
18. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (Vol. 30, pp. 4765-4774).
19. Mahajan, A. S. (2025). Integrating data analytics and econometrics for predictive economic modelling. International Journal of Applied Mathematics, 38(2s), 1450-1462.
20. Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77-91.
21. Mahajan, A. S., Yamsani, N., & Konatham, M. R. (2024). Financial risk prediction through multivariate data analytics and ensemble learning. Stochastic Modelling and Computational Sciences, 4(2).
22. Pearl, J. (2009). Causality: Models, reasoning, and inference (2nd ed.). Cambridge University Press.
23. Ramareddy, S. K. (2024a). Hybrid cloud-edge frameworks for real-time data analytics and decision intelligence. International Journal of Electrical, Electronics and Computer Systems, 13(2), 51-61. https://doi.org/10.65521/ijeecs.v13i2.877
24. Porter, M. E. (2010). What is value in health care? New England Journal of Medicine, 363(26), 2477-2481.
25. Ramareddy, S. K. (2024b). Intelligent data warehouse optimization for scalable cloud analytics using artificial intelligence. International Journal for Multidisciplinary Research, 6(6), 1-13. https://www.ijfmr.com/research-paper.php?id=84920
26. Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. Academy of Management Review, 46(1), 192-210.

27. Ramareddy, S. K. (2025a). An AI-orchestrated framework for self-learning predictive maintenance in cloud infrastructure. International Journal of Scientific Research in Engineering and Management, 9(4), 1-8. https://doi.org/10.55041/IJSREM43578

28. Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347-1358.

29. Ramareddy, S. K. (2025b). The role of artificial intelligence in transforming financial services through cloud-enabled data analytics. International Journal of Applied Mathematics, 38(12s), 4592-4606. https://doi.org/10.12732/ijam.v38i12s.2008

30. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135-1144).

31. Ramareddy, S. K. (2026a). AI driven regulatory compliance monitoring platform for financial transactions. International Journal of Computer Information Systems and Industrial Management Applications, 18(14s), 266-278. https://doi.org/10.70917/ijcisim-2026-4211

32. Rubin, D. B. (1976). Inference and missing data. Biometrika, 63(3), 581-592.

33. Ramareddy, S. K. (2026b). Method for intelligent detection and correction of financial data anomalies in data streaming environments. International Journal of Artificial Intelligence and Machine Learning, 6(9s), 548-558. https://doi.org/10.51483/IJAIML.6.9s.2026.548-558

34. Russell, S., & Norvig, P. (2020). Artificial intelligence: A modern approach (4th ed.). Pearson.

35. Sharma, A. B., & Kargeti, H. (2020). Investigation & scrutiny of protected assignment supervision systems for mobile enabled with Internet Protocol Version 6 networks. American Journal of Computer Science and Technology, 3(4), 76-85.

36. Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human-Computer Interaction, 36(6), 495-504.

37. Uddandarao, D. P., & Mahajan, A. (2023). Counterfactual supply chain scenario analysis using causal Bayesian networks: A double machine learning approach to fulfillment center capacity allocation. International Journal of Advances in Signal and Image Sciences, 34-46.

38. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.

39. Veernapu, K. (2021). The role of artificial intelligence in healthcare finance: Improving financial forecasts and operational effectiveness. International Journal of Multidisciplinary Research and Growth Evaluation, 2(4), 873-876.

40. Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. Strategic Management Journal, 28(13), 1319-1350.

41. Veernapu, K. (2022a). Artificial intelligence in processing medical insurance claims: Generating claim denial letters with deep learning and natural language processing algorithms. International Journal, 3(3), 643-646.

42. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44-56.

43. Veernapu, K. (2022b). Resolving perioperative inventory patient billing challenges using AI and deep learning. IJLRP-International Journal of Leading Research Publication, 3(4), 1-8.

44. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (Vol. 30, pp. 5998-6008).

45. Veernapu, K. (2023a). Streamlining the recruitment process in healthcare with artificial intelligence. IJLRP-International Journal of Leading Research Publication, 4(1), 1-7.

46. Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. PLoS Medicine, 15(11), e1002689.

47. Veernapu, K. (2023b). AI algorithms and healthcare scheduling: Exploring their potential to address workforce shortages and improve care quality. Progress in Medical Sciences, 7(3), 1-4.

48. Vellido, A. (2020). The importance of interpretability and visualization in machine learning for applications in medicine and health care. Neural Computing and Applications, 32(24), 18069-18083.

49. Veernapu, K. (2025). Deep learning in detecting healthcare fraud: The economic burden on public and private insurers. International Journal of Core Engineering & Management, 8(1), 31-38.

50. Wooldridge, J. M. (2010). Econometric analysis of cross section and panel data (2nd ed.). MIT Press.

51. Agha, C. W., Mokwenye, I. I., Atikpo, E., & Nwulu, N. (2025). Gumbel distribution-based flood frequency analysis for coastal flood risk assessment in Ugborodo community, Niger Delta, Nigeria. Environmental Technology & Science Journal, 16(1), 133–143. https://doi.org/10.4314/etsj.v16i1.13

52. Bhalekar, O. A., & Bamniya, A. (2025). Secure data transmission in industrial IoT: A cryptography-based encryption-decryption scheme for manufacturing systems. In 2025 International Conference on Smart & Sustainable Technology (INCSST) (pp. 1–6). IEEE. https://doi.org/10.1109/INCSST64791.2025.11210305

53. Bhalekar, O. A., Dasari, P. K., & Paul, A. B. (2025). AI-driven secure smart manufacturing: Integrating database indexing, industrial cybersecurity, and real-time IIoT analytics in wireless automation architectures. Letters in High Energy Physics, 2025, 238–245. https://doi.org/10.5281/zenodo.21394929

54. Kale, M., Elinjulliparamb, S. H., Song, H. H., Faheem, Z. B., Alghamdi, N. S., Kim, S. M., & Ali, J. (2026). FedShield: Robust federated intrusion detection for consumer IoT under malicious updates using blockchain validation. IEEE Transactions on Consumer Electronics. https://doi.org/10.1109/TCE.2026.3706133