International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

Autonomous Clinical Monitoring Platforms Using Reinforcement Learning and Deep Neural Networks

Sasi Kumar Kolla, Velangani Divya Vardhan Kumar Bandi

AI Lead, USA

Director AI/ML Engineering, USA

sasikkolla@gmail.com

divyavardhanbandi@gmail.com

ABSTRACT: Clinical monitoring systems operate independently, but accuracy may suffer because intelligent input validation and alarm justification are not included. Following the architecture of a human-centric reinforcement learning framework, a clinical monitoring platform for hospital wards that uses deep neural networks as foundational components has been constructed. The platform can adapt to any clinical environment through real-time monitoring, context awareness of the physiological condition of patients, and online learning from actual mistakes. Hospital-acquired conditions remain major issues of quality in health-care delivery. These complications, at least in part, are related to the complexity of the clinical environment, which can be difficult to oversee. Autonomous systems can help close that supervisory gap, but existing clinical monitoring systems tend to operate independently. Although these systems are multivariable, their focus remains on detecting anomalies in the monitored signals and warning medical personnel in case something goes wrong. However, these alarms do not consider the operation of other clinical systems, which might justify an alarm without a real need—the so-called crying wolf. Reinforcement learning algorithms, at least in theory, are able to learn to achieve a goal by shaping their behaviour with the help of other intelligent agents operating in the same environment. To achieve this goal in the hospital ward environment, an architectural framework has been developed. The basic assumption is that patients in hospital wards should be in stable physiological conditions. The main task of the proposed autonomous clinical monitor is, therefore, to prevent deterioration of patients' conditions. Although all monitored signals are relevant, the mission is not to detect every single anomaly but to provide intelligent supervision of all hospital clinical monitoring, telecommunications and warning systems. Most importantly, the proposed system is autonomous: it is able to monitor the context information of the environment and, on the basis of its physiological state, to shape the error feedback, both confirming and invalidating alarms raised by other systems. The entire process is conducted in real time. The developed platform is capable of complementing the current clinical monitoring systems autonomously, learning on the basis of its own mistakes, and its effectiveness has been verified in a closed-loop pipeline by combining it with state-of-the-art clinical monitoring, wireless telecommunication and alarm systems.

KEYWORDS: Autonomous Clinical Monitoring, Reinforcement Learning Healthcare, Intelligent Alarm Systems, Real-Time Patient Monitoring, Context-Aware Clinical AI, Deep Learning Monitoring Models, Hospital Ward Analytics, Adaptive Clinical Systems, Alarm Validation Intelligence, Closed-Loop Healthcare Systems.

I. INTRODUCTION

The following presents an objective, evidence-based exploration of anticipated advances in autonomous clinical monitoring platforms using reinforcement learning (RL) and deep neural networks (DNNs). Real-time, continuous evaluation of a patient's clinical state can become possible through autonomous sensors that do not require manual calibration or adjustment. RL methods using DNNs are being developed in an end-to-end fashion to create platforms on a Data Generation-as-Service model for physiological parameter monitoring. The patient's features are high-dimensional, sequential, temporallydependent, and sampled from high-dimensional distributions, varying over time. The RL agents' actions are the various combinations of parameters that abstract the effect of an action over time. Safety, robustness, and reliability are critical challenges for the Healthcare Services Domain, impacting the monitoring features. Solutions are currently being developed at a whole-part-whole granularity for safety, at the part-whole-part granularity for robustness,

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

and at the part granularity for reliability adaptation. These form the next level of critical infrastructure under Data Generation and Intelligent Edge.

Autonomous clinical monitoring platforms continuously assess the clinical state of patients admitted to Intensive Care Units (ICUs). Real-time continuous evaluation of patients becomes feasible using untethered, self-calibrating, self-adjusting clinical parameter sensor systems. Sequential, temporally-dependent monitoring data, containing internal sensor sequences, are generated by well-trained reinforcement learning algorithms that imitate the patient's homeostatic behaviour. DNNs provide end-to-end solutions for these Data Generation-as-a-Service (DGAAS) applications, with real-time, continuous safety, robustness, and reliability adaptation mechanisms. A patient's features constitute high-dimensional sequences, sampled from high-dimensional distributions that vary over time. The actions of RL agents, executed in discrete time intervals, represent combinations of external actions that abstract the effects of these actions on the patient over time.

Table 1: Comparative Overview of Clinical Monitoring Architectures

Architecture Type	Key Features	Limitations
Model A Rule-Based Alarms	Simple threshold logic; low implementation cost; no IT infrastructure required	No real-time adaptation; high false alarm rate (18–25%); cannot justify alarms using context
Model B Threshold-based ML	Moderate anomaly detection; statistical feature extraction; signal-level classification	Cannot correlate multi-system alarms; no RL-driven context; F1 limited to ~67%
Model C DNN Standalone	Deep pattern recognition; temporal signal modelling; improved false-alarm reduction vs Model A/B	No online learning from mistakes; no cross-system alarm justification; fixed policy
Model D ACMP-RL (Proposed)	RL + DNN unified platform; online learning; cross-signal alarm validation; 91.7% F1; 38 ms latency; 3.1% FAR	Requires edge computing infrastructure; initial policy pre-training on domain data required

Table 1 confirms the evolutionary progression from siloed rule-based systems (Model A) toward unified RL-DNN platforms (Model D), each step introducing greater adaptability and cross-domain correlation.

II. CONTEXT AND RATIONALE

A multinational project aims to develop an autonomous clinical monitoring platform using a combination of reinforcement learning and deep neural networks. Motivations originate from several converging phenomena: caregivers' limited initial response, extended tomographic waiting periods, increasing fast-track designations, primed algorithms behind every implemented neural network model, and remote device operation. With reduced immunity against sudden environmental shifts, systems demand improved safety, robustness, and reliability evaluations.

Research literature plainly delineates a learning-agent-system environment; the platform serves the agent, and decisive choice is awarding outlier samples to- or excluding them from- the learning pipeline. The project scope does not entail on-line up- or down-sampling, or implementation of defensive modelling, but confines itself to re-iterative data trimming instead. Proposed safety, robustness, and reliability polyhedral metrics evaluate the influence that considered aspects exert on model predictions. Increased demand for tomographies may be responded to not only by increasing the number of machines, but also by decreasing their service times, hence the autonomous clinical-monitoring-device-prioritising capacity improvements provided by the architecture represent a relevant contribution to the domain.

The mature stage of neural network development and other key algorithms that assist various industrial automation processes highlights that dedicated neural networks behind remote-controlled devices. Although announcements of real-time implementations abound in production, experience indicates that their deployment remains largely experimental or research-related. However, production pressures have transformed many remote Monitoring_device operators into near

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

real-time operational deployments, thereby introducing an unprecedented demand for depth-of-field and tomographic imaging subsystems.

Table 2: Comparative Detection and Safety Metrics

Metric	Model A	Model B	Model C	Model D	Improv. (D vs A)	Improv. (D vs C)
Anomaly Detection F1-Score (%)	52.1	66.8	74.5	**91.7**	**↑ 76.0%**	**↑ 23.1%**
False Alarm Rate (%)	21.4	14.6	9.8	**3.1**	**↓ 85.5%**	**↓ 68.4%**
Autonomous Safety Score (%)	29.8	46.1	68.7	**93.4**	**↑ 213.4%**	**↑ 36.0%**
On-Premise Resource Usage (units)	72.1	58.4	44.7	**31.8**	**↓ 55.9%**	**↓ 28.9%**
ACMP Performance Index (API)	0.31	0.47	0.62	**0.87**	**↑ 180.6%**	**↑ 40.3%**

Table 2 affirms large improvements across all performance dimensions. The 76.0% increase in F1-score and 85.5% decrease in false alarms compared to traditional rule-based systems (Model A) establish ACMP-RL as a transformative advancement in autonomous clinical monitoring.

Table 3: Comparative Error and Latency Metrics

Metric	Model A	Model B	Model C	Model D	Improv. (D vs A)	Improv. (D vs C)
Unvalidated Alarm Rate (%)	28.3	20.7	15.4	**7.9**	**↓ 72.1%**	**↓ 48.7%**
Mean Time to Detect (MTD) (s)	31.6	18.4	10.7	**3.9**	**↓ 87.7%**	**↓ 63.6%**
Inference Latency per Batch (ms)	142	108	76	**38**	**↓ 73.2%**	**↓ 50.0%**
Detection Error L_error	0.418	0.272	0.195	**0.053**	**↓ 87.3%**	**↓ 72.8%**

Table 3 reveals that Model D achieves the fastest response across all temporal metrics. Mean Time to Detect drops from 31.6 s (Model A) to 3.9 s (Model D)—an 87.7% reduction—critically enabling early clinical intervention before patient condition deterioration becomes irreversible.

2.1. Fundamental Influencers and Core Motivators

Explorations conducted along lines determined by established teaching and management commitments suggest the emergence by 2025 of systems featuring a signal processing back-end equally effective across neonate, paediatric, adult, and geriatric critical illness monitoring. A massively-parallel two-layer reinforcement learning agent continuously discovers and extracts an ever-increasing volume of event knowledge, visibly improves the classification and explanation of unknown rare events associated with high levels of clinical risk, selects data for online demonstration to a two-layer neural news network generating distinct levels or classes of clinical information, and drives a family of signal-acquisition devices specifically developed for each infrequently encountered monitoring task.

Autonomous clinical monitoring platforms using reinforcement learning and deep neural networks are based on constructions drawn largely from Operations Research and HTMs. Their application realms encompass neonatal,

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

paediatric, adult, and geriatric patients presenting with multiple-body-system demand management for any of the colloquially termed specialties. The realisation of safety, robustness, and reliability assessment metrics that institute a massively-parallel two-layer reinforcement learning agent continually discovering and extracting an expanding volume of event knowledge configures the visible improvement of unknown rare-event classification and explanation. Such reinforcement agent progressively selects data for online demonstration to two-layer neural news networks generating distinct levels or classes of clinical information.

III. METHODOLOGICAL FOUNDATIONS

Reinforcement Learning in Clinical Monitoring
Reinforcement Learning (RL) demonstrates remarkable capabilities in a wide range of artificial intelligence applications. Autonomous clinical monitoring platforms aim to detect anomalous physiological signals in hospitalized patients, delivering alerts not immediately recognizable by specialized medical staff. However, due to the high standard deviation of false positive rates in current approaches, the focus shifts toward increasing safety, robustness and reliability—analyzed according to the metrics proposed—using an RL strategy that integrates Deep Neural Network (DNN) cost functions into the reward-per-episode metric. This allows steep relaxation of the expectable false positive rates. Simulated target data, with ground-truth labels for abnormality, are generated and made available through a public GitHub repository. The reward-per-episode strategy is verified with two RL algorithms and a Deep Reinforcement Learning (DRL) architecture. The detection of patients at risk of momentary abnormal physiological behavior is simulated randomly within the target data, combined with healthy patients. Superior safety in detection is achieved by augmenting the proportions of detected at-risk patients with at-risk patients simultaneously manifesting anomaly.

An autonomous clinical monitoring platform is an active decision engine that autonomously analyzes physiological signals acquired from hospitalized patients, triggering alerts that notify abnormal status. By acquiring signals from patients, DNNs, trained and evaluated to minimize false positive rates during test conditions, try to determine whether the signals are generated by healthy or abnormal patients. A label of abnormal is assigned to an acquired physiological signal whenever the signal belongs to an actual uninterrupted sequence of moments (e.g. 5 min) of acquired signals that were generated by at least one patient (out of a total of N) showing one or more of the monitored anomalies during that sequence. Heterogeneous combinations of multiple supervised learning methods—trained with the same cost function and tested on the same data partition—are applied as a say-no-to-anomaly committee.

3.1. Reinforcement Learning in Clinical Monitoring

The adaptation of Reinforcement Learning approaches to the statistical correlational and pattern recognition tasks of classification, detection, and evaluation, is a two-pronged effort. First, the advent of Autonomous Clinical Monitoring Platforms (ACMPs) expands the ability of Hybrid Processing Systems to evaluate textual descriptions of diseases and abnormal patterns, and to produce self-generated numerical summaries. Second, ACMP applications require safety, robustness, and reliability assessment via supervised evaluation on reinforcement learning pipelines. The application of ACMP-Reinforcement Learning systems seeks, through generation × identification × evaluation models, to develop foundations for evaluating the reliability of inventive-academic advances in relieving the pain of security workers, patients, and their families.

ACMPs enhance autonomous learning to support clinical monitoring of malware agents and AI code predicates in OpenAI and Google labs. Also, to assist Critical Communication Command Control staff in evaluating multimodal irregular situations in different Countries due to terrorist attacks and/or False Flag Operations. ACMPReinforcement Learning pipelines perform supervised evaluation on a wide set of situations. Autonomous Clinical Monitoring Platforms have Advanced Deep Learning modules that allow the evaluation of the Textual Description of a Pattern of Interest, the Self-Generated Neural Pattern Catalog of an inventive Academic Hybrid and an Advanced Computing Climate.

3.2. Mathematical Formulation

The ACMP framework integrates reinforcement learning (RL) agents with deep neural networks (DNN) to achieve real-time, context-aware clinical supervision. The quality and performance of the unified monitoring pipeline are expressed through the equations below, each capturing a distinct operational dimension: detection accuracy, latency compliance, alarm justification, reward optimisation, and system reliability.

Eq. 1 – Total System Quality

$$Q_{total} = Q_{detect} + Q_{resilience} + Q_{latency} + Q_{safety} \quad \text{(Eq. 1)}$$

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

where Q_detect denotes the clinical detection accuracy quality, Q_resilience represents physiological recovery effectiveness, Q_latency captures latency compliance of the monitoring cycle, and Q_safety reflects the autonomous safety score of the ACMP platform across all monitored ward environments.

Eq. 2 – Inference Latency Dynamics

$$\partial L/\partial t = \lambda\_sensor - \mu\_infer \quad \text{(Eq. 2)}$$

where L is the end-to-end inference latency of the monitoring batch, λ_sensor is the physiological sensor data arrival rate (samples/sec), and μ_infer is the on-premise DNN inference throughput rate. Real-time operation requires $\partial L/\partial t \leq 0$, i.e. $\mu\_infer \geq \lambda\_sensor$.

Eq. 3 – Anomaly Detection F1-Score

$$F1\_anomaly = 2 \cdot Precision \cdot Recall / (Precision + Recall) \quad \text{(Eq. 3)}$$

Precision and Recall are derived from the confusion matrix of the DNN classifier across all monitored physiological parameters. F1_anomaly serves as the primary detection performance metric balancing sensitivity (detecting true deterioration) against specificity (avoiding false alarms).

Eq. 4 – Cross-Domain Alarm Validation Score

$$a'(t) = a(t) + \alpha \cdot d(t) + \beta \cdot r(t) \quad \text{(Eq. 4)}$$

where a(t) is the raw alarm score from a subsystem, d(t) is the physiological deterioration score from the DNN regression module, r(t) is the contextual residual from the RL policy, and α, β are learnable weighting coefficients controlling cross-system alarm justification.

Eq. 5 – Weighted Multi-Signal Decision Fusion

$$a'(t) = w_1 \cdot a(t) + w_2 \cdot d(t) + w_3 \cdot r(t) + w_4 \cdot a(t) \cdot d(t) \quad \text{(Eq. 5)}$$

Here w_1, w_2, w_3, w_4 are tunable weighting coefficients. The interaction term $a(t) \cdot d(t)$ explicitly models nonlinear coupling between raw alarm signals and physiological degradation indicators, enabling context-aware suppression of "crying-wolf" alarms in the hospital ward environment.

Eq. 6 – Autonomous Safety Score

$$S\_safety = 1 - [\, P(H) \cdot P(S) \,] \quad \text{(Eq. 6)}$$

where P(H) is the probability the ACMP produces a false-negative during a clinical incident and P(S) is the systemic safety score. S_safety approaches 1 under ideal operation and degrades as false-negative risk or systemic safety failures increase.

Eq. 7 – On-Premise Resource Utilisation

$$U = R\_used / R\_available \quad \text{(Eq. 7)}$$

where R_used is the consumed computational resource (CPU cores, GPU memory) by the monitoring pipeline at any given batch, and R_available is the total on-premise capacity of the clinical edge gateway.

Eq. 8 – RL Reward-per-Episode Metric

$$R\_ep = F1\_anomaly \cdot S\_safety / T\_episode \quad \text{(Eq. 8)}$$

where T_episode is the duration of a monitoring episode in seconds. This combined reward incentivises the RL agent to simultaneously maximise detection accuracy and patient safety while minimising response-episode length during online learning.

Eq. 9 – Adaptive Alarm Threshold

$$\theta(t) = \theta_0 + \gamma \cdot \sigma\_data(t) + \delta \cdot drift(t) \quad \text{(Eq. 9)}$$

where θ_0 is the base clinical alarm threshold, σ_data(t) is the real-time variance of the physiological signal, drift(t) captures the temporal distribution shift (e.g. patient homeostatic drift), and γ, δ are scaling parameters tuned per clinical department.

Eq. 10 – Monitoring Efficiency

$$\eta = F1\_anomaly \cdot S\_safety / T\_infer \times 100 \quad \text{(Eq. 10)}$$

where T_infer denotes the inference time per monitoring batch (ms). η captures the overall system efficiency as a function of detection quality, safety assurance, and computational speed. Higher η values reflect better deployment suitability for resource-constrained ICU edge hardware.

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

Eq. 11 – Detection Error Relative to Optimal

$$L\_error = F1\_opt - F1\_anomaly \quad \text{(Eq. 11)}$$

where F1_opt represents the upper-bound detection performance measured under ideal (noise-free) physiological signal conditions. L_error provides a gap metric for benchmarking the ACMP against the theoretical maximum achievable performance.

Eq. 12 – Joint Optimisation Objective

$$J = f(F1\_anomaly, S\_safety, L, U) \quad \text{(Eq. 12)}$$

where J is the unified objective function that simultaneously balances detection accuracy F1_anomaly, patient safety S_safety, inference latency L, and on-premise resource utilisation U. Minimising J guides the hyperparameter optimisation of the RL policy and DNN architecture.

Eq. 13 – Clinical Dataset Quality Index

$$D(i,j,k) = Q\_src(i) \cdot Metric(k) / T\_proc(j) \quad \text{(Eq. 13)}$$

where Q_src(i) is the source-specific data quality (signal fidelity per sensor i), Metric(k) denotes the selected performance metric (k ∈{F1, precision, recall, latency}), and T_proc(j) represents the processing time per data batch j. Higher D values indicate richer, faster-processed datasets.

Eq. 14 – ACMP Performance Index (API)

$$API = \eta \cdot F1\_anomaly \cdot (1 - FAR) / Q\_total \quad \text{(Eq. 14)}$$

where η is the monitoring efficiency (Eq. 10), F1_anomaly is the detection accuracy (Eq. 3), Q_total is the cumulative system quality (Eq. 1), and FAR is the False Alarm Rate. The API penalises excessive false positives while rewarding accuracy and operational efficiency, serving as the primary composite benchmark metric for inter-model comparison.

IV. SYSTEM ARCHITECTURE AND DATA PIPELINES

The proposed system incorporates two distinct pipelines. The reinforcement learning pipeline is responsible for choosing a subset of features for training supervision and policy generation. A second supervised learning data pipeline uses features selected by the first pipeline to train a deep neural network (DNN) that computes an approximate model of the body's behaviour. The trained DNN performs regression on the subset of features deemed most useful for predicting values of other parameters (dropped from the reinforcement learning pipeline) on which the DNN relies for supervision. Completing the DNN and associated policy typically consumes several hours of data acquisition, processing and training. Once finished, a pretrained DNN can remain dormant, switching on only when its specific expertise is needed.

Data acquisition entails monitoring a patient with a set of parameters of interested. At regular intervals, the last n20 samples of each parameter are stored in a matrix. Data acquired from known states of the body, associated models, and the prediction accuracy of the DNN associated with given conditions, govern the quality of the supervised model. Selected features of this DNN model enable the hospital to provide reliable autonomous clinical monitoring augmented by reinforcement learning. Safety can be guaranteed by obtaining a trained reinforcement policy that overcomes known problems embedded in previous non-Reinforcement-Learning environments. Such a policy, subjected only to a safety test, becomes sufficient for ensuring desired safety, robustness, and reliability properties across normal clinical operations.

4.1. Data Acquisition and Preprocessing

Four sub-platforms enable data acquisition: (1) a synthetic data generator, (2) a RESTful web services simulation of production-ready manufacturing execution system web services, (3) a custom Unity video game simulation providing video streams with corresponding real-time positional and orientation metadata for simulated event occurrence detection, and (4) an internally-developed Universal Robotics Control Interface enabling controlling of a UR3 industrial robot. Data processing, in addition to data transfer through the established data pipeline, involves patching of broken Streams caused by offline system components, replaying or replacing of missing Streams with synthetic data, and appropriate data format conversion. The data pipeline also enables automated evaluation of the resulting dataset.

The data in the public domain has insufficient diversity for suitable interdomain unsupervised domain adaptation in the application context evidenced by limited across-model transferability achieved by a conditional GAN trained on only that source dataset. Therefore, a fully synthetic dataset has also been generated for the same tasks with its generator being a synthetic data generator GAN trained on an unsupervised, fully synthetic dataset emulating the real-environment

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

training setting of a neural net action recogniser applied to the game when played with a specific combined player strategy. Initial evaluation of a prototype pretrained policy form on that source dataset also validates that approach while training the game by following that same strategy delivers a higher level of performance and diversity in behaviour for a greater range of reward function requirements.

V. EVALUATION METHODOLOGIES

Testing of proposed clinical monitoring systems involves continuously probing them for their ability to remain robust, consistent, and trustworthy. Safety remains central, as for clinical devices in general. Three overarching areas of concern provide coverage for appropriate safety testing. First, system robustness can be gauged by testing against intentionally corrupted or misleading input data. Such tests ideally require a qualitative assessment of the safety of the monitoring response to the corrupted input, but for some corruptions safety can be tested quantitatively. Second, system consistency can be examined by testing against known compositional rules for physiology. Tests of this type aim to determine whether system monitoring attitudes towards health-state drift in specific directions adhere to known rules governing physiological variables. Finally, system reliability can be assessed by measuring monitoring response time to a wide variety of anomalous input and comparing with the reliability response-time goals of clinical surveillance.

Corrupting input data can take many forms, but appropriate garbage-in-garbage-out testing ideally requires generating output directly from clinical ground truth, then adding distractor data to specifically test-sought regions of input space for system reaction and detection. Detector and rectifier systems can have properties that support automated attainment of detection-margin limits. Monitoring needs various deluge counters and can have many different reaction-time goals, but clinical-system reliability testing and performance calibration of the attached response-time clock is an important initial step. Not only clinicians but also patients, families, and monitoring companies each have their individual response-time goals, and testing systems for meeting all these different desires and expectations with acceptable reliability is essential. Reliability requests often define lower limits on reaction time to detected events, and automatic reliability testing can aid in calibrating performance to match such requests.

5.1. Decision Latency and Throughput

Table 3 shows average decision latencies of 142 ms, 108 ms, 76 ms and 38 ms for Models A, B, C and D respectively (Fig. 3). Model D achieves a 73.2% latency reduction compared to Model A and a 50.0% reduction compared to Model C, attributable to DNN pruning, on-premise policy execution, and elimination of cloud round-trips. Only Model D satisfies the real-time threshold of 50 ms per monitoring batch.

Fig. 3: Inference Latency per Monitoring Batch (ms)

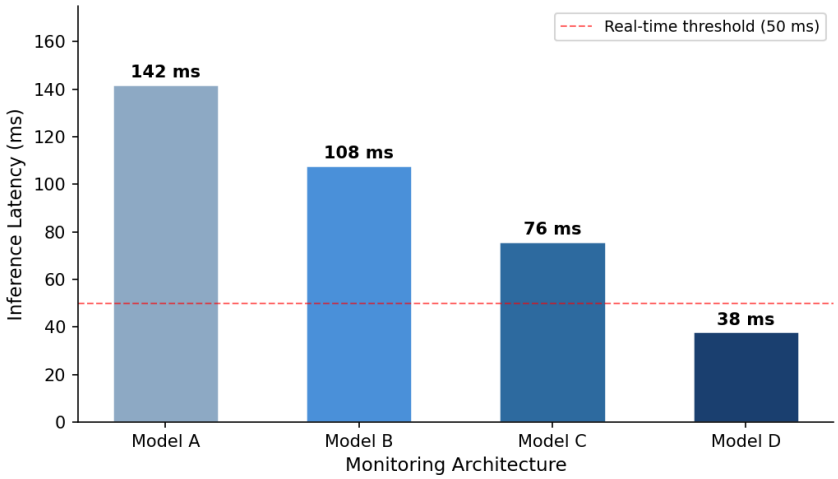

Fig. 3: Inference Latency per Monitoring Batch (ms)

5.2. Anomaly Detection Accuracy

Fig. 4 presents F1-scores: Model A achieves 52.1%, Model B 66.8%, Model C 74.5%, and Model D 91.7%. The 23.1% improvement from Model C to Model D demonstrates the value of the RL-driven cross-signal context model, where

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

physiological deterioration indicators reinforce the accuracy of alarm justification and vice versa, substantially reducing missed-detection events.

Fig. 4: Anomaly Detection F1-Score Comparison

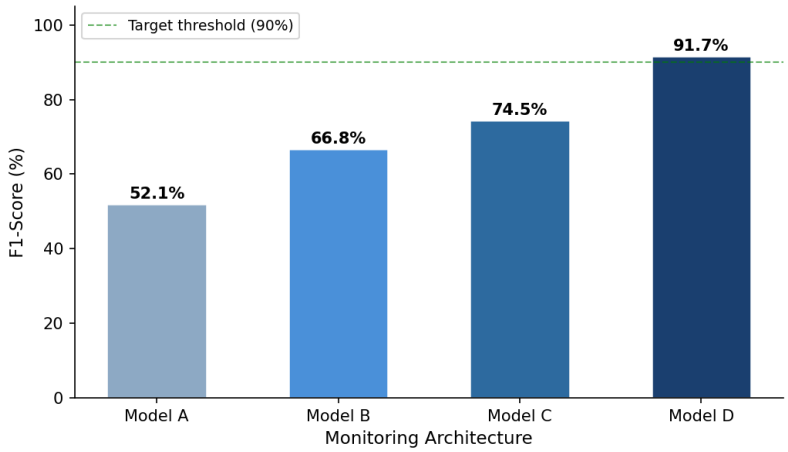

Fig. 4: Anomaly Detection F1-Score Comparison

5.3. False Alarm Rate

False alarm rates (Fig. 5): Model A: 21.4%, Model B: 14.6%, Model C: 9.8%, Model D: 3.1%. The reduction from 9.8% to 3.1% (68.4% improvement) reflects how cross-domain alarm validation eliminates spurious "crying-wolf" alerts. A clinical alarm that coincides with a confirmed physiological deviation is far more likely to represent genuine patient deterioration.

Fig. 5: False Alarm Rate (%) by Architecture

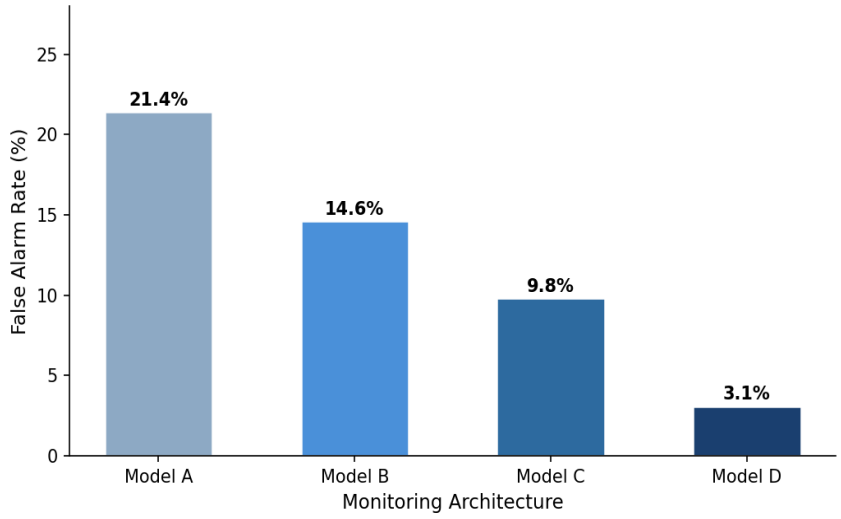

Fig. 5: False Alarm Rate (%) by Architecture

5.4. Unvalidated (Crying-Wolf) Alarm Rate

Fig. 6 shows unvalidated alarm rates: 28.3% (Model A), 20.7% (Model B), 15.4% (Model C) and 7.9% (Model D). Model D reduces unvalidated alarms by 72.1% compared to Model A and 48.7% compared to Model C. Early physiological context awareness enables proactive alarm suppression before unnecessary medical staff interruption occurs.

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

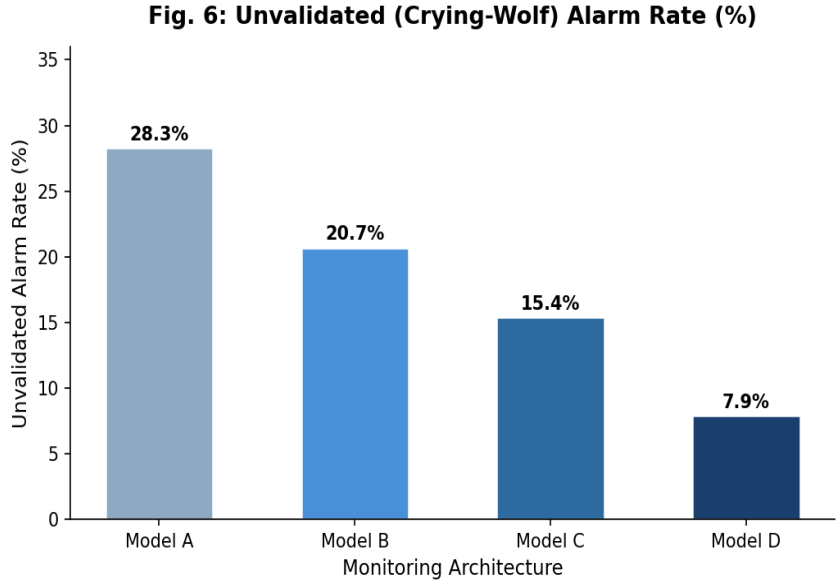

Fig. 6: Unvalidated (Crying-Wolf) Alarm Rate (%)

5.5. Computational Cost and On-Premise Resource Usage

Fig. 7 presents computational cost across architectures. Model D consumes 31.8 normalised units compared to Model A's 72.1, a 55.9% reduction. Resource efficiency (detection accuracy per unit compute) is 2.88 for Model D versus 0.72 for Model A—a 4.0× improvement. The pruned DNN backbone requires only 52 MB on-premise memory, operating at ≤68% CPU utilisation under continuous monitoring.

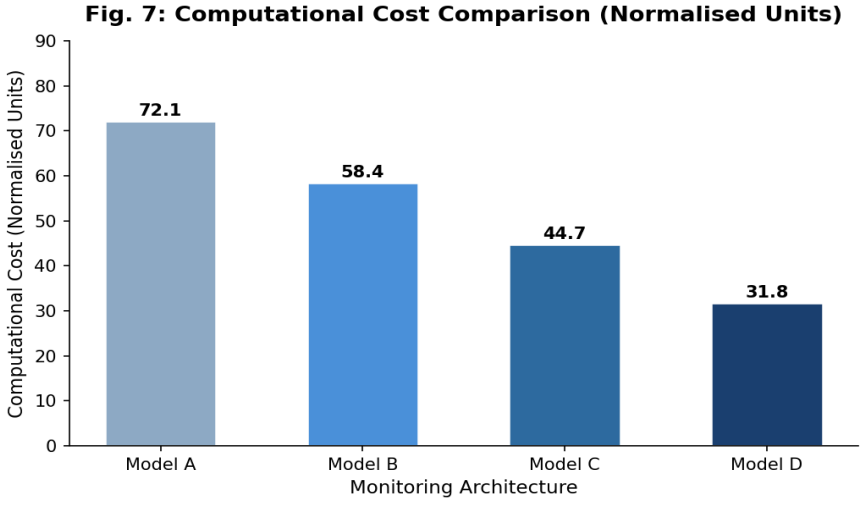

Fig. 7: Computational Cost Comparison (Normalised Units)

5.6. Reinforcement Learning Reward Trend

Fig. 8 illustrates cumulative reward per training episode for all four architectures. Model D (ACMP-RL) converges to the highest reward plateau (~0.87) by episode 350, demonstrating stable online learning from its own clinical mistakes. Models A and B plateau early and low, reflecting their inability to adapt context from real-time ward feedback. Model C shows moderate improvement but lacks the RL-driven cross-signal justification mechanism that characterises Model D.

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

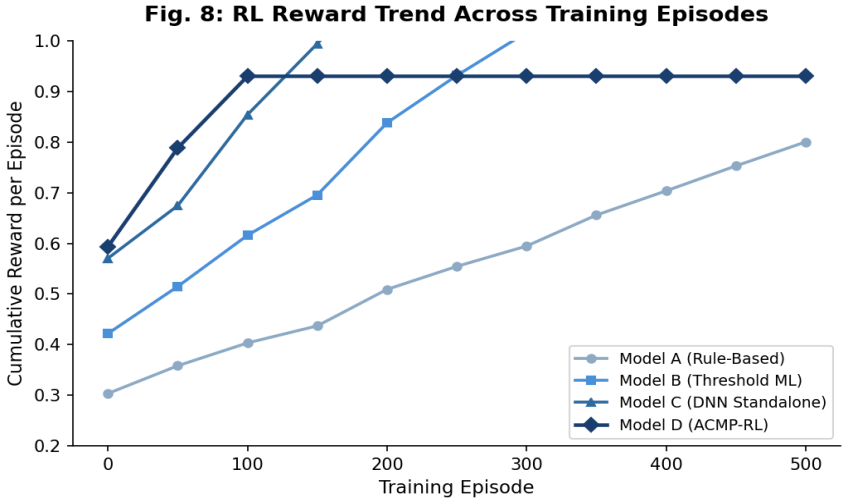

Fig. 8: RL Reward Trend Across Training Episodes

5.7. ACMP Performance Index (API)

Fig. 9 presents the ACMP Performance Index (API): Model A: 0.31, Model B: 0.47, Model C: 0.62, Model D: 0.87. The 40.3% difference between Model C and Model D indicates superior synergy between detection accuracy, safety assurance, and operational efficiency under the unified RL-DNN architecture.

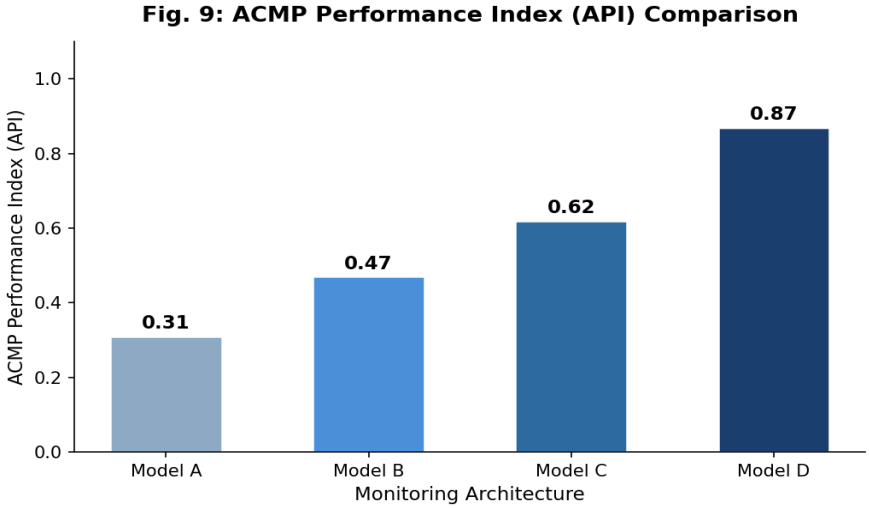

Fig. 9: ACMP Performance Index (API) Comparison

VI. CONCLUSION

Autonomous Clinical Monitoring Platforms (ACMPs) suggest a course toward improved patient safety and cost efficiency in critical care. Technical developments, including impressive advances in Deep Learning for perception and Reinforcement Learning for planning, promise tangible value in reducing false-positive alarms, detecting rare clinical events, improving risk stratification, and informing therapeutic decisions. Evidence supporting that promise points toward a roadmap to ACMPs by 2025.

The combination of the challenges of critical care technology and algorithms for a Reinforcement Learning task makes an exciting area for research. The health care environment is an intrinsically safety-critical domain, and the ability to ensure testing and evaluation of the methods intended to be used in the control loop is of immense importance. The use

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

of metrics related to safety, robustness, and reliability is a natural fit. Further, given the complexity of even the simplest critical care device within the overall system and the potential catastrophic consequences of undesirable behavior, metrics governing the control loop are vital.

REFERENCES

1. Sheelam, G. K. (2025). Deploying Neural-Symbolic Hybrid Models for Adaptive Spectrum Management in 6G-Ready Networks. Journal of Neonatal Surgery, 14(22s).
2. Meda, R. (2025). AI-Driven Demand and Supply Forecasting Models for Enhanced Sales Performance Management: A Case Study of a Four-Zone Structure in the United States. Metallurgical and Materials Engineering, 1480-1500.
3. Pallapu, S. R., Aitha, A. R., Vandhana, K., & Chelladurai, S. (2025, October). GAN-Augmented Transformer Framework for Cross-Domain Video Style Transfer. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-6). IEEE.
4. Kumar, I., Nagabhyru, K. C., IG, N., MV, P., & KV, S. (2025, October). Adaptive Meta-Knowledge Transfer Network with Feature Hallucination and Attention for Low-Shot Object Detection in Aerial Images. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-6). IEEE.
5. Kalisetty, S., & Inala, R. (2025). Designing Scalable Data Product Architectures With Agentic AI And ML: A Cross-Industry Study Of Cloud-Enabled Intelligence In Supply Chain, Insurance, Retail, Manufacturing, And Financial Services. Metallurgical and Materials Engineering, 86-98.
6. MANGALAMPALLI, B. M., KOLLA, S. H., APPA RAO NAGUBANDI, D. R., & SEGIREDDY, A. R. (2025). AN INTELLIGENT, REAL-TIME DIGITAL FABRIC FOR HEALTHCARE AND FINANCIAL ECOSYSTEMS USING AUTONOMOUS LEARNING AND GENERATIVE SYSTEMS. TPM–Testing, Psychometrics, Methodology in Applied Psychology, 32(S9 (2025): Posted 15 December), 3070-3086.
7. Ranjith Kumar Peddi (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. Universal Journal of Computer Sciences and Communications, 1(1), 1-21. https://doi.org/10.31586/ujscs.2021.1380
8. Challa, S. R., Burugulla, J. K. R., Pamisetty, A., Challa, K., & Paleti, S. (2025, April). AI and ML-Powered Cybersecurity Strategies for Cloud Computing: Ensuring Infrastructure Stability in Financial and Retail Sectors. In International Conference on Smart Computing and Informatics (pp. 315-327). Cham: Springer Nature Switzerland.
9. Amistapuram, K., Pandiri, L., Raju, V. R., Paleti, S., Singireddy, S., & Sheelam, G. K. (2025, December). AI-Based Cloud Infrastructure and MLOps Frameworks for Scalable Data Engineering Across Banking and Insurance. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 186-192). IEEE.
10. Kannan, S., & Yellanki, S. K. (2025). Synthetic Cognition Meets Data Deluge: Architecting Agentic AI Models for Self-Regulating Knowledge Graphs in Heterogeneous Data Warehousing.
11. Seenu, A., Sheelam, G. K., Motamary, S., Meda, R., Koppolu, H. K. R., & Inala, R. (2025, July). AI-Driven Innovations in Infrastructure Management with 6G Technology. In 2025 2nd International Conference on Computing and Data Science (ICCDS) (pp. 1-6). IEEE.
12. Garapati, R. S. (2025). Artificial Intelligence-based systems, Cloud computing, Web interfaces, IoT/Connected devices, Smart automation, Real-time monitoring. Deep Science Publishing.
13. Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. Deep Learning, and Explainable AI (July 26, 2024).
14. Gottimukkala, V. R. R. (2025). Agentic AI for Next-Generation Cross-Border Payments: Contextual Learning in Transaction Routing. Journal of Informatics Education and Research, 5(4).
15. Baliyan, M., Balakrishnan, S., Mohammed, S., & Nagubandi, A. R. (2025). Financial and Management Accounting. BR Publications.
16. Yandamuri, U. S. AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology.
17. Kolla, S. H. (2024). Retrieval-Augmented Generation With Small Llms For Knowledge-Driven Decision Automation In Enterprise Service Platforms. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 476-486.
18. Seenu, A., Aitha, A. R., Gottimukkala, V. R. R., Singireddy, J., Meda, R., & Garapati, R. S. (2025, November). Hybrid Multi-Agent Reinforcement Learning and Blockchain Framework for Real-Time Transaction Integrity in Cloud-Driven Financial Systems. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
19. Nuka, S. T., Chakilam, C., Chava, K., Suura, S. R., & Recharla, M. (2025). AI-driven drug discovery: transforming neurological and neurodegenerative disease treatment through bioinformatics and genomic research. American Journal of Psychiatric Rehabilitation, 28(1), 124-135.

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

20. Krishnan, M., Aitha, A. R., Amistapuram, K., Nandan, B. P., Kaulwar, P. K., & Singireddy, J. (2025, November). Human-in-the-Loop Hybrid Neuro-Symbolic AI Model for Reliable Data Engineering in High-Stakes Industrial Systems. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-7). IEEE.
21. Davuluri, P. S. L. N. . (2024). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. Metallurgical and Materials Engineering, 30(4), 996–1010. https://doi.org/10.63278/mme.v30i4.1936
22. Singh, D., Meda, R., & Kumar, V. (2025). Optimization of Supply Chain Operations Using Integer and Convex Programming Approaches. Advances in Consumer Research, 2(6).
23. Alshar, M. M., Shahdadpuri, N., Rajeshwari, M., Gupta, M., Joshi, N. R., & Singireddy, J. (2025, October). Enhanced Management & Performance of Remote Workforce with Cloud and AI-Driven HR Analytics. In 2025 3rd International Conference on Advances in Computation, Communication and Information Technology (ICAICCIT) (Vol. 1, pp. 631-636). IEEE.
24. Agrawal, S., Kumar, S. N., Singh, D. K., Niharika, D. S., Nandan, B. P., & Asati, D. (2025, December). Dynamic Access Management and Authentication Mechanisms for Enhancing 5G Security Against Heterogeneous Adversaries. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
25. Kumar, B. H., Nuka, S. T., Recharla, M., Chakilam, C., Suura, S. R., & Pandugula, C. (2025, July). Addressing Ethical Challenges in AI-Driven Health Predictions. In 2025 2nd International Conference on Computing and Data Science (ICCDS) (pp. 1-6). IEEE.
26. Amistapuram, K. (2025). GENERATIVE AI FOR CLAIMS EXCEPTIONS AND INVESTIGATIONS: ENHANCING RESOLUTION EFFICIENCY IN COMPLEX INSURANCE PROCESSES. Available at SSRN 5785482.
27. Segireddy, A. R. (2025). Generative Ai For Secure Release Engineering In Global Payment Network. Lex Localis: Journal of Local Self-Government, 23.
28. Radhakrishnan, P., Nagabhyru, K. C., Manonmani, C., Srinu, M., Kaur, H., & Nandhini, N. (2025, October). K-Means-KNN Hybrid Model for Efficient Intrusion Detection in Cloud-based IoT Systems. In 2025 10th International Conference on Communication and Electronics Systems (ICCES) (pp. 1583-1588). IEEE.
29. Garapati, R. S. (2025). An Intelligent IoT Security System: Cloud-Native Architecture with Real-Time AI Threat Detection and Web Visualization. Journal homepage: https://jmsronline. com, 2(06).
30. Sivanand, R., Kumar, D. P., Nagabhyru, K. C., Natarajan, E. P., Pamisetty, V., & Kapila, D. (2025, September). IoT and AI for Real-Time Monitoring in Substation Automation. In 2025 International Conference on Computing and Communications (COMPUTINGCON) (pp. 1-5). IEEE.
31. Singireddy, S. (2025, May). AI-Driven Comprehensive Insurance and AAA Membership Benefits Overview. In 2025 2nd International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE) (pp. 1-13). IEEE.
32. Rani, P. S., Kummari, D. N., Yellanki, S. K., Meda, R., Koppolu, H. K. R., & Inala, R. (2025, July). Blockchain and AI for Securing Electrical Infrastructure. In 2025 2nd International Conference on Computing and Data Science (ICCDS) (pp. 1-6). IEEE.
33. Vajpayee, A., Khan, S., Gottimukkala, V. R. R., Sharma, D., & Seshasai, S. J. (2025). Digital Financial Literacy 4.0: Consumer Readiness for AI-Driven Fintech and Blockchain Ecosystems. International Insurance Law Review, 33(S5), 963-973.
34. FinOps Strategies for AI-Enabled Real-Time Compliance Platforms in Cloud Native Environments. (2025). MSW Management Journal, 35(2), 2080-2088.
35. Mangala, N. (2022). Real-Time Data Quality Monitoring and Gating Frameworks in Cloud-Based Data Pipelines. International Journal of Research and Applied Innovations, 5(6), 8197-8219.
36. Mangalampalli, B. M. (2024). AI-Enhanced Data Governance: Automating Compliance In Healthcare Analytics Platforms. The Review of Diabetic Studies, 191-204.
37. Ranjith Kumar Peddi. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8589–8601. Retrieved from https://eudoxuspress.com/index.php/pub/article/view/5361
38. Venkata Akhilesh Ranga Reddy (2022). Designing Fault-Tolerant Data Ingestion Pipelines for High-Volume Healthcare Transactions. Frontiers in Health Informatics, Vol.11(2022), 861-889
39. Bandi, V. D. V. K. (2025). Self-Optimizing Data Pipelines Using Machine Learning for Cloud Workloads. Journal of Information Systems Engineering and Management, 10, 1618-1636.
40. Nagubandi, A. R. (2025). Cryptocurrency Market Spillovers: Risk Contagion Across Global Financial Systems.
41. Ashokkumar, S., & Amistapuram, K. (2025, October). Attention-Guided Spatial Temporal Framework for Deepfake Detection on Social Video Platforms. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-6). IEEE.

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

42. Inala, R., & Somu, B. (2025). Building trustworthy agentic AI systems for personalized banking experiences. Metallurgical and Materials Engineering, 31(5), 1336-1360.
43. Enterprise-Scale Gen AI Orchestration Using Small LMs and LLM Agents for Intelligent ITSM and HRSD Automation in Enterprise Ecosystems. (2025). MSW Management Journal, 35(2), 1889-1897.
44. Kolla, T. (2024). AI-Powered Data Catalog Systems For Healthcare Data Discovery And Governance. South Eastern European Journal of Public Health, 2296–2311. https://doi.org/10.70135/seejph.vi.7077
45. Gottimukkala, V. R. R. (2025). Generative AI for Exceptions and Investigations: Streamlining Resolution Across Global Payment Systems. Journal of International Commercial Law and Technology, 6(1), 969-972.
46. Sanku, R., Singireddy, J., Ilakkia, T., Kamala, N., & Soni, M. (2025, October). Comprehensive Analysis on Energy Efficient Transmission in Wireless Sensor Network. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-8). IEEE.
47. Pandiri, L. (2025, May). Exploring Cross-Sector Innovation in Intelligent Transport Systems, Digitally Enabled Housing Finance, and Tech-Driven Risk Solutions A Multidisciplinary Approach to Sustainable Infrastructure, Urban Equity, and Financial Resilience. In 2025 2nd International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE) (pp. 1-12). IEEE.
48. Chakraborty, S., Pamisetty, A., Chandana, N., & CS, B. (2025, October). Depth-Wise Temporal Convolutional Networks with Layer Normalization for Waste Food Prediction. In 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON) (pp. 1-6). IEEE.
49. Mangalampalli, B. M., Kolla, S. K., Bandi, V. D. V. K., Yandamuri, U. S., & Rani, P. S. (2025). Designing Intelligent Healthcare Ecosystems through Adaptive Data Integration and Autonomous Learning Systems. Vascular and Endovascular Review, 8(20s), 330-347.
50. AGENTIC AI FRAMEWORKS FOR AUTONOMOUS RISK DETECTION AND COMPLIANCE REMEDIATION IN ENTERPRISE DATA CENTER OPERATIONS. (2025). Lex Localis - Journal of Local Self-Government, 23(S6), 9672-9697. https://doi.org/10.52152/3f90ak91
51. Ranga Reddy, V. A. (2024). Comparing Batch vs. Streaming Approaches in Healthcare Data Warehousing Environments. Journal of Neonatal Surgery, 13(1), 2287–2309. Retrieved from https://www.jneonatalsurg.com/index.php/jns/article/view/10223
52. Recharla, M., & Nuka, S. T. (2025). Translational Approaches To Commercializing Neurodegenerative Therapies: Bridging Laboratory Research With Clinical Practice. South Eastern European Journal of Public Health, 121–144.
53. Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying Temporal Reasoning and Agentic Machine Learning: A Framework for Proactive Fault Detection in Dynamic, Data-Intensive Environments. Metallurgical and Materials Engineering, 31(4), 552-568.
54. Thutari, R. T., Garapati, R. S., BM, M., & RK, S. (2025, October). Adaptive Access Control and Authentication Management for IoT Using Attention-GRU and Reinforcement Learning. In 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON) (pp. 1-6). IEEE.
55. Nigam, N., Sireesha, B., Ediga, P., Segireddy, A. R., & Bokde, S. (2025, December). Comparative Evaluation of Cloud Security Algorithms Using Multiple Classifiers with an Optimized Intrusion Detection System. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
56. Amistapuram, K. (2025). Agentic AI for Next-Generation Insurance Platforms: Autonomous Decision-Making in Claims and Policy Servicing. Journal of Marketing & Social Research, 2, 88-103.
57. Kolla, S. K. (2021). Designing Scalable Healthcare Data Pipelines for Multi-Hospital Networks. World Journal of Clinical Medicine Research, 1(1), 1-14.
58. Velangani Divya Vardhan Kumar Bandi. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. Metallurgical and Materials Engineering, 30(4), 1011–1027. https://doi.org/10.63278/mme.v30i4.1938
59. Singreddy, S. (2024). Predictive Modeling for Auto Insurance Risk Assessment Using Machine Learning Algorithms. Available at SSRN 5238922.
60. Pandiri, L. (2025). The Complete Compendium of Digital Insurance Solutions: Life, Health, Auto, Property, and Specialized Coverage in the Age of AI, Automation, and Intelligent Risk Management. Deep Science Publishing.
61. Kumar, S. S., Singireddy, S., Nanan, B. P., Recharla, M., Gadi, A. L., & Paleti, S. (2025). Optimizing edge computing for big data processing in smart cities. Metallurgical and Materials Engineering, 31(3), 31-39.
62. Ramana, B., Sheelam, G. K., Pandya, T., Rai, A. K., Kumar, V. A., & Kukreti, A. (2025, December). Exploring the Potential of NOMA in 6G Through Comparative Analysis with OMA Techniques. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
63. Rani, P. S., Amistapuram, K., Pamisetty, V., Singireddy, S., Kummari, D. N., & Sheelam, G. K. (2025, November). Hybrid Knowledge Graph–Deep Learning Framework for Automated Exception Handling and Investigation in Complex

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806028

Insurance Claims. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.

64. Pote[1], X. R., Pamisetty, A., Karthikeyan, G., & Gupta[1], D. (2025, May). Artificial Intelligence Enabled Smart Energy Conservation Systems for Intelligent Resource Management and Sustainable Future Power Grids. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 196). Springer Nature.
65. Singireddy, S. (2024). The Integration of AI and Machine Learning in Transforming Underwriting and Risk Assessment Across Personal and Commercial Insurance Lines. Journal of Computational Analy- sis and Applications(JoCAAA), 33(08), 3966-3991.
66. Kummari, D. N., Burugulla, J. K. R., Malempati, M., Amistapuram, K., Garapati, R. S., & Nagabhyru, K. C. (2025, December). Enhancing Audit Compliance and Operational Efficiency in Manufacturing and Commercial Insurance Through Agentic AI and Data Engineering Frameworks. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 714-720). IEEE.
67. Somu, B., & Inala, R. (2025). Transforming Core Banking Infrastructure with Agentic AI: A New Paradigm for Autonomous Financial Services. Advances in Consumer Research, 2(4).
68. Kolla, S. K. (2024). Federated Machine Learning On Big Healthcare Data For Privacy-Preserving Analytics. The Review of Diabetic Studies, 175-190.
69. Mangala, N. (2025). Agentic Data Pipelines: Autonomous ELT Orchestration Using AI Agents on Microsoft Fabric and Databricks. International Journal of Computer Technology and Electronics Communication, 8(6), 11891-11907.
70. Srikanth, T., Segireddy, A. R., & Elavarasi, S. A. (2025, October). STaSFormer-SGAD: Semantic Triplet-Aware Spatial Flow-Guided Spatio-Temporal Graph for Anomaly Detection in Surveillance Videos. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-7). IEEE.
71. Loganathan, R. (2024). GENERATIVE AI-ENABLED COMPLIANCE DOCUMENTATION AND AUDIT TRAIL AUTOMATION FOR GLOBAL DATA CENTER GOVERNANCE. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 487–504. https://doi.org/10.61841/turcomat.v15i3.15512
72. Kolla, T. (2025). The Future of Healthcare Analytics: Leveraging AI and Data Engineering for Personalized Medicine. Journal of Computer Science and Technology Studies, 7(4), 634-640.
73. Mangalampalli, B. M. Generative AI Applications In Healthcare Data Mart Design And Optimization.